Title
A model of cat primary visual cortex and its thalamic input

Author
Krukowski, Anton E.

Publication Date
2000

Peer reviewed|Thesis/dissertation

A Model of Cat Primary Visual Cortex and its Thalamic Input

by

Anton E. Krukowski

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Biophysics

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA SAN FRANCISCO

[redacted]

Date University Librarian

Degree Conferred: .....

To Mom and Dad and Eva.

Acknowledgements

While research is often a lonely undertaking, most of what I have learned during my (oh so) many years at UCSF has been from the people around me. The structural biology group was a wonderfully rich intellectual world to be immersed in when I first got to UCSF, barely knowing any biology at all. David Agard was a terrific research advisor during the year before I started graduate school. He and David Baker patiently taught me a huge amount of biology, and were a pleasure to work with. Ken Dill was a great teacher and research advisor. I feel lucky to have had the chance to work with him and Hue-Sun Chan, who were models for me in how to do rigorous mathematical and computational research. I am grateful to have tasted a number of different fields in my rotation projects with Robert Stroud, Peter Walter, Ken Dill and Juan Korenbrot. Juan gave me my first glimpse of the enticing world of neuroscience, which led to my decision to switch over to Ken Miller's laboratory. I am grateful to the biophysics program for the flexibility that allowed me to make the switch; I am particularly grateful to Tack Kuntz for helping me to realize it was what I wanted to do, and to Julie Ransom for many supportive conversations and logistical help both at that time and throughout my time here.

Once again, I found myself in an intellectually rich atmosphere, as a relative naive. The Keck Center has been a wonderful place to learn neuroscience and to grow as a scientist. It is hard to list everyone who has been a part of making this such a stimulating environment for me. Most of the neuroscience I learned was from the graduate students of the neuroscience program, who were also a warm and supportive cohort, and welcoming to me as a biophysical intruder. I want especially to thank Michele Solis, Michael Silver, Kimberly Tanner, Henry Mahncke and Charlotte Boettiger for their guidance and friendship. I am also grateful that the Sloan program brought such a varied group of theoretically minded, bright, interesting, witty, and warm post-docs to the greater HSE 806, tri-closet area. I have learned a lot from Frederic Theunissen, Mark Kvale, Hagai Attias, Kamal Sen, Robert Liu, KT Moortgat, Suman Kumar, and Brian Wright; I'd like especially to thank Zeph Landau, Carlos Cassanello, Leslie Osborne, and Ginny de Sa for their warmth and friendship, and the stimulating conversations about both science and Keck Center gossip.

And then there is Ken's lab itself. Ken is a remarkable scientist. His energy and enthusiasm for science in general, and for the work of the lab, has seemed boundless. Over the years there have been countless times that I have been rejuvenated by his excitement for our work together. He is also one of the most generous people I know. He has been endlessly giving of his time, as a teacher, as a research mentor, and as a friend. Todd Troyer has been more than a long term research partner and collaborator on a lot of the work presented in this dissertation. He has also been a second mentor for me, along with Ken. He and Ken have both taught me much of what I know about how to "do science": how to zero in on key issues, how to think clearly, and how to present ideas. I have learned a lot from working with Nicholas Priebe and Andy Kayser, and it was wonderful to have them as fellow students in (and around) the modeling end of the lab. My recent collaboration with

Thomas Lauritzen has been stimulating and fruitful, and it is gratifying to know that he will be continuing to work on issues related to my dissertation research. Brian Wright has been a generous and tireless support with the computers, whether he was actually in Ken's lab or not, and it's been a pleasure to share an office with him. Paul Bush taught me a lot about the visual system and a wide range of computational methods, and I had numerous interesting conversations with him while he was in the lab. I have also enjoyed my many conversations with Ed Erwin, Sergei Rebrik and Al Emondi over the years.

Finally, I'd like to thank the members of my thesis and orals committees for many helpful and stimulating conversations and suggestions on my work: Michael Stryker, Christoph Schreiner, Allison Doupe, Stephen Lisberger, Ken Dill, and, of course, Ken Miller.

My friend Mark once told me about a man he had interviewed who was a survivor of the Holocaust, and had become a sculptor as an adult. Mark said that the sculptor was one of the most content people he had ever met, despite the fact that he had lived through unspeakable trauma as a child, and had lost his entire family. The sculptor described to him a large piece that he worked on for a long time. It was placed in the middle of a square, where children played on it, posters were stuck on it, and pigeons defaced it every day. Mark asked him if it bothered him that it was not treated with more respect. It did not bother him at all, he said. He had poured a piece of his soul into an object that was a part of people's daily life, and it would be for years and years to come. They did not need to react to it as a work of art for it to affect them. The sculptor loved his work. Creating things that were permanent made him extremely happy.

As I am preparing to laser print onto acid-free, 100% cotton paper this document that is meant to represent many years of my life, I feel a strong urge to pour my soul onto the relatively permanent page. The acknowledgements section seems like the place to try such a thing, since the description of the research that follows feels more like an expression of my sweat than of my inner life. And an acknowledgment of the people who have shaped me in my life outside of work is the stuff of what moves me the most. But this is not, in fact, an artform. I do not get the benefit of partially camouflaging my difficult and complex feelings in intricate metaphor, or twisted bronze. Without that cover, I feel too naked to put my feelings into words. I hope the people who have been closest to me can glean from these clumsy phrases the intensity of the feelings that underlie them.

Statement from Research Advisor on Role of Anton Krukowski in Co-authored Work

The first chapter of Anton Krukowski's thesis is the following published paper:

- Troyer, T.W., A.E. Krukowski, N.J. Priebe and K.D. Miller (1998). "Contrast-Invariant Orientation Tuning in Visual Cortex: Thalamocortical Input Tuning and Correlation-Based Intracortical Connectivity". *Journal of Neuroscience* **18**, 5908-5927

This paper was published with a footnote stating that the first two authors (Troyer and Krukowski) contributed equally to the work. All of the studies of the "computational model" in that paper were

carried out by Krukowski; the simulations and analysis of the "conceptual model" were carried out by Troyer. The three of us worked closely together and ideas flowed readily in all directions.

This work was initially carried out by Krukowski and the research advisor Miller. Krukowski spent substantial time (around a year) writing a simulator to study the effect of a hypothesized correlation-based circuit for layer 4 of cat V1 on the contrast dependence of orientation tuning of cortical cells. Nick Priebe was a rotation student who wrote the simulator for the LGN inputs to the cortical model. When the simulations were finally run, it was discovered that the hypothesized circuit in fact rendered the orientation tuning contrast invariant. We quickly came to an intuitive understanding of this result in terms of a simple four-cell circuit. Todd Troyer, a postdoc in the lab, then took this four-cell picture and turned it into a mathematical model that could be used to demonstrate the workings of the circuit in the simplest possible form. Troyer, Krukowski and Miller all worked on the the first version of the manuscript, which was submitted to Science, and Krukowski took a leave of absence at the time of submission. That was not accepted, at which point Troyer took first author responsibilities and along with Miller began writing the manuscript for Journal of Neuroscience. Before the current manuscript was completed, Krukowski returned to the project and was heavily involved in all remaining stages of the writing of the work.

In summary, the paper comprising the first chapter represents the outcome of a very fruitful collaboration initiated by Krukowski's work and to which, in the end, Troyer and Krukowski contributed equally. The remainder of the thesis represents Krukowski's continued development of the model first studied in that paper.

Kenneth D. Miller

Abstract

A Model of Cat Primary Visual Cortex and its Thalamic Input

Anton E. Krukowski

There are several response properties that distinguish cortical simple cells of cat primary visual cortex (V1) from the lateral geniculate nucleus (LGN) cells that provide their input. Understanding the origin of these differences is a central problem for understanding cerebral cortical circuitry. This dissertation uses a semi-detailed model of the excitatory and inhibitory cells in layer 4 of V1, and their synaptic connections both within the cortex and from their LGN inputs, to explore what patterns of circuitry and what physiological mechanisms are consistent with the observed response properties of the cortex.

In the first chapter, we characterize the input to cat simple cells generated by a oriented arrangement of LGN afferents as proposed by Hubel and Wiesel. We demonstrate that this type of input is not sufficient on its own to account for the invariance of the width of cortical orientation tuning with respect to stimulus contrast. We then describe a simple intracortical circuit that is expected to develop under correlation-based rules of synaptic plasticity, and demonstrate that it can robustly account for contrast-invariant orientation tuning.

In the second chapter, we show that the same model circuit unexpectedly accounts for the low-pass shift in the temporal frequency tuning curves of cortical cells relative to geniculate cells. This arises when NMDA-mediated excitatory conductances are included in the geniculocortical synapses, at levels observed experimentally in thalamocortical synapses. The longer time course of the NMDA conductance, as compared to AMPA conductances, selectively preserves the modulation of lower frequency gratings, while it dampens out the modulation of higher frequencies. Since cortical firing is predicted to occur only with temporal modulation strong enough to overcome the inhibitory mean component, the NMDA-induced dampening of higher frequency modulations prevents cortical response to higher-frequency inputs.

In the third chapter we run the model in a series of simulations over a range of parameter space in order to explore how the main model parameters affect a number of experimental constraints. In the final chapter we explore how certain parameters effect predictions of the orienation and temporal frequency tuning of the mean firing rate of the model inhibitory cells.

Contents

List of Tables

List of Figures

Chapter 1

Contrast-Invariant Orientation Tuning

Abstract

The origin of orientation selectivity in visual cortical responses is a central problem for understanding cerebral cortical circuitry. In cats, many experiments suggest that orientation selectivity arises from the arrangement of LGN afferents to layer 4 simple cells. However, this explanation is not sufficient to account for the contrast invariance of orientation tuning.

To understand contrast invariance, we first characterize the input to cat simple cells generated by the oriented arrangement of LGN afferents. We demonstrate that it has two components: a spatial-phase-specific component (*i.e.*, one that depends on receptive field spatial phase), which is tuned for orientation; and a phase-nonspecific component, which is untuned. Both components grow with contrast.

Second, we show that a correlation-based intracortical circuit – in which connectivity between cell pairs is determined by the correlation of their LGN inputs – is sufficient to achieve well-tuned, contrast-invariant orientation tuning. This circuit generates both spatially opponent, "anti-phase" inhibition ("push-pull"), and spatially matched, "same-phase" excitation. The inhibition, if sufficiently strong, suppresses the untuned input component and sharpens responses to the tuned component at all contrasts. The excitation amplifies tuned responses. This circuit agrees with experimental evidence showing spatial opponency between, and similar orientation tuning of, the excitatory and inhibitory inputs received by a simple cell. Orientation tuning is primarily input driven, accounting for the observed invariance of tuning width after removal of intracortical synaptic input, as well as for the dependence of orientation tuning on stimulus spatial frequency.

The model differs from previous push-pull models in requiring dominant rather than balanced inhibition, and in predicting that a population of layer 4 inhibitory neurons should respond in a contrast-dependent manner to stimuli of all orientations, although their tuning width may be similar to that of excitatory neurons. The model demonstrates that fundamental response properties of cortical layer 4 can be explained by circuitry expected to develop under correlation-based rules of

1.1 Introduction

Thirty-five years ago, Hubel and Wiesel discovered that cells in cat primary visual cortex (V1) are tuned for the orientation of light/dark borders [Hubel and Wiesel, 1962]. The inputs to V1 come from the lateral geniculate nucleus (LGN), whose cells are not significantly orientation selective [Hubel and Wiesel, 1961]. The origin of orientation selectivity in visual cortex has been one of the most thoroughly investigated questions in neuroscience, and serves as a model problem for understanding how the cortex processes and represents information.

In cats, orientation selective responses appear in cortical layer 4. Cat layer 4 is composed of *simple cells* [Hubel and Wiesel, 1962, Gilbert, 1977, Bullier and Henry, 1979]: cells with receptive fields (RFs) composed of oriented subregions, each giving exclusively ON or OFF responses (response to light onset/dark offset or light offset/dark onset). Hubel and Wiesel proposed that the orientation selectivity of these cells derives from an oriented arrangement of inputs from the LGN: ON-center LGN inputs have RF centers aligned over the simple cell's ON-subregions, and similarly for OFF-center inputs [Hubel and Wiesel, 1962]. Such an input arrangement has been confirmed experimentally Reid and Alonso [1995], Tanaka [1983]. Because the total LGN input grows with increasing contrast for stimuli of all orientations, this model by itself is insufficient to explain the invariance of orientation tuning under change in stimulus contrast [Sclar and Freeman, 1982, Skottun et al., 1987]. A threshold for spiking responses might narrow the tuning at any one contrast, but higher contrast would require a higher threshold to prevent broadening of tuning.

Two major approaches to achieving contrast invariance have been proposed. Many authors have suggested that responses in simple cells are approximately linear, *i.e.* the response can be predicted by linear summation of stimulus luminance (relative to background), weighted by the cell's RF Movshon et al. [1978a], Tolhurst and Dean [1990], Heeger [1992], Carandini and Heeger [1994], Carandini et al. [1998, 1997], Albrecht and Geisler [1991], Glezer et al. [1982]. Contrast change in such a model simply multiplies responses by a constant; contrast-invariant tuning follows automatically. It has been proposed that linear responses might be achieved by a balanced "push-pull" arrangement of inputs, in which an ON-subregion shows equal excitation ("push") to light stimuli as inhibition ("pull") to dark stimuli, and conversely for OFF-subregions [Glezer et al., 1982, Tolhurst and Dean, 1990, Carandini and Heeger, 1994, Carandini et al., 1998, 1997]. However, there are two problems with achieving linear response in an actual neural circuit. First, spike thresholds are nonzero, and therefore oriented stimuli that at low contrast give positive but sub-threshold input would yield spike responses at higher contrast. Second, at contrasts above about 5%, LGN responses increase more than they decrease, since spike rates cannot decrease below zero (*i.e.* responses "rectify"). This input nonlinearity alters the balance between push and pull.

Other authors have proposed that orientation tuning emerges from orientation-specific short-range excitation and longer-range inhibition in cortex [Ben-Yishai et al., 1995, Somers et al., 1995], despite evidence that in cat layer 4, excitation and inhibition show similar orientation tuning [Fer-

ster, 1986]. The width of orientation tuning in these models is an emergent property of intracortical circuitry, and so doesn't depend on the parameters of the stimulus, including stimulus contrast. These proposals appear inconsistent with the fact that orientation tuning widths in cats do depend on at least one stimulus parameter, the spatial frequency of sinusoidal grating stimuli [Vidyasagar and Sigüenza, 1985, Webster and De Valois, 1985, Jones et al., 1987, Hammond and Pomfrett, 1990].

We propose a new model for cat layer 4 cortical circuitry that yields contrast-invariant orientation tuning. Our model addresses two basic questions. First, what is the nature of the thalamocortical input to cortical simple cells? We assume thalamocortical connectivity can be modeled by a Gabor function: a two-dimensional Gaussian multiplied by a sinusoid [Jones et al., 1987, Reid and Alonso, 1995]. The *spatial phase* of the sinusoid determines the location of ON and OFF subregions within the thalamocortical RF. Using a simple model of LGN responses, we show that the total LGN input has two components: a spatial-phase-specific component (a component that varies with the spatial phase of a cell's RF) that is tuned for orientation, and a phase-nonspecific component that is entirely untuned. Both components grow with contrast. Separating these input components helps clarify the debate over whether the LGN input to simple cells is well or poorly tuned. In response to drifting gratings, the phase-specific component corresponds to the temporally modulated input component, which Ferster et al. [1996] recently demonstrated to be tuned. However, the total input includes the phase-nonspecific, temporally unmodulated component; this should be untuned, and was not measured by Ferster et al. [1996]. Separating the input components also clarifies the problem that cortical circuitry must solve to achieve contrast-invariant orientation tuning: eliminate the untuned component of the LGN input in a contrast-dependent manner, while extracting and sharpening the tuned component.

Second, what patterns of intracortical connectivity are sufficient to yield contrast-invariant orientation tuning? We arrive at a surprisingly simple answer: "correlation-based" connectivity yields contrast invariance. By correlation-based connectivity, we mean that intracortical connection strengths between two cells are fixed based on the correlation in their thalamocortical RFs. Thus, inhibitory connections occur between cells with anticorrelated RFs, while excitatory connections occur between cells with correlated RFs. The "anti-phase" inhibition eliminates the untuned input component and sharpens responses to the tuned component, while "same-phase" intracortical excitation amplifies the tuned response. As a result, our model achieves contrast-invariant tuning in the presence of positive thresholds and LGN rectification.

Our model uses a form of "push-pull" circuitry, but differs from other such models in that inhibition dominates rather than balances excitation and responses are not linear. Furthermore, we predict that a population of inhibitory neurons in cat layer 4 should respond in a contrast-dependent manner to stimuli of all orientations, although they may be tuned for orientation. The model has both developmental and functional implications for understanding the layer 4 cortical

circuit, and suggests a general means of separating stimulus intensity (here represented by contrast) from stimulus form (represented by orientation).

An abstract of this work has appeared previously [Krukowski et al., 1996].

1.2 Methods

We study both a very simple ("conceptual") model and a more realistic ("computational") model. We first present the elements common to both, then present each model.

1.2.1 Elements Common to Both Conceptual and Computational Models

LGN Model. Our model was based on cat V1 at approximately 5^o eccentricity. LGN spatial RFs were center-surround difference of Gaussians, with cells responding either to light onset (ON cells) or light offset (OFF cells) in their RF centers. LGN spatial filter parameters ($\frac{17}{\sigma^2_{\text{center}}} e^{-x^2/\sigma^2_{\text{center}}} - \frac{16}{\sigma^2_{\text{surround}}} e^{-x^2/\sigma^2_{\text{surround}}}$; $\sigma_{\text{center}} = 15'$, $\sigma_{\text{surround}} = 1^o$) were taken from Peichl and Wassle [1979] and Linsenmeier et al. [1982]. Firing rates in response to sinusoidal gratings were calculated on the assumption of linear rectified responses (unrectified firing rate was a sinusoid of the same temporal frequency as the stimulus; negative rates were then set to zero), using contrast response curves fromCheng et al. [1995] (Fig. 1.1). Assuming background firing rates of 10 Hz (ON cells) and 15 Hz (OFF cells) (modified from Kaplan et al. [1987], considering the lower mean luminance of 20 cd/m^2 used in Cheng et al. [1995])
we calculated the sinusoidal amplitude that would lead to the reported values of the first harmonic (F1) after rectification. (Throughout, we will use F1 to denote the amplitude of the sinusoidal component at the frequency of the grating stimulus, although this value is twice as large as the value obtained using the Fourier transform normalized so that the F0 or DC component is the mean level. [Skottun et al., 1991] The amplitudes were then fit to $R = R_{max}(\frac{C^n}{C_{50}^n + C^n})$, where R is response amplitude and C is contrast (ON cells: $R_{max} = 53.0$ Hz, $n = 1.20$, $C_{50} = 13.3\%$; OFF cells: $R_{max} = 48.6$ Hz, $n = 1.29$, $C_{50} = 7.18\%$). LGN responses for gratings of non-optimal spatial frequencies were calculated by reducing modulation amplitudes by the factor predicted from the application of LGN spatial filters. ON and OFF cells had temporal phases offset by 180^o. To calculate the firing rates in response to moving bars, LGN cell spatiotemporal RFs were used. Temporal filters were taken from the central RF pixel in reverse correlation data from 100% contrast M-sequences (supplied by R.C. Reid); center and surround temporal filters were assumed equal for simplicity.

Cortical Receptive Fields. Cat cortical layer 4 simple cell RFs were modeled as Gabor functions (Fig. 1.2). A Gabor function is a 2D Gaussian, here with peak value 1, multiplied by a sinusoid. Positive regions of the Gabor correspond to ON subregions and yield connections from ON-center LGN cells, negative regions correspond to OFF subregions and yield OFF-center inputs;

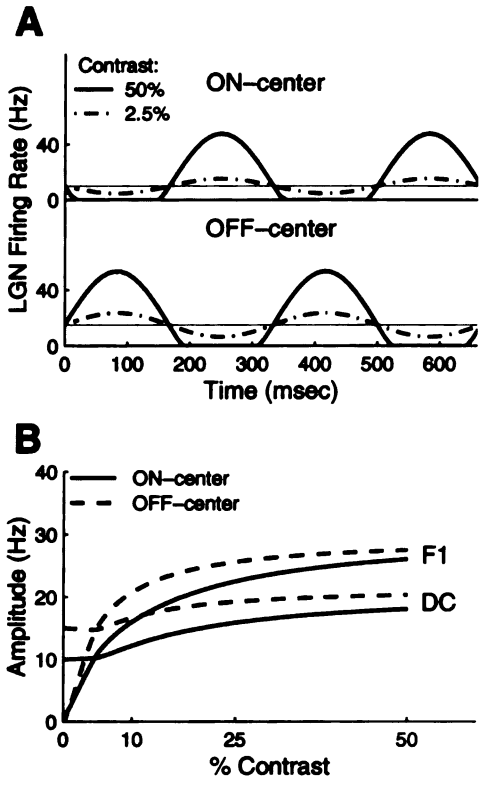

Figure 1.1:

LGN cell responses to 3 Hz, 0.8 cycles/degree moving gratings. **A.** Instantaneous firing rate. Straight line: background. **B.** Contrast response functions. Upper: amplitude of first harmonic (F1); lower: mean (DC) firing rate. The mean rate increases at contrasts greater than 5%, due to rectification as seen in (A). Data modified from Cheng et al. [1995]) (see Methods).

the strength of the connection depends on the magnitude of the Gabor. The number of subregions is defined as the ratio of the width of the Gaussian envelope (at 5% of peak) to the width of a half-cycle of the sinusoid. The aspect ratio of a single subfield is defined as the ratio of the Gaussian envelope length to the sinusoid half-cycle width. Two sets of Gabor parameters were used. "Default" parameters were the mean values for simple cell physiological RFs reported in Jones and Palmer [1987]: 2.65 subregions and an aspect ratio of 4.54. (Care must be taken when comparing these numbers to other experimental estimates – *e.g.*, using a 10% cutoff for the Gaussian reduces these numbers by nearly 1/4). All RFs have 0.625^o half-cycle width, corresponding to a preferred spatial frequency of 0.8 cycles/degree, the approximate mean preferred spatial frequency of cortical cells at 5^o eccentricity [Movshon et al., 1978a]. Gaussian 5% envelope length and width are equal to 2.84^o and 1.65^o respectively. The measurements of Ferster et al. [1996] suggest that the net LGN input to a simple cell has broader orientation tuning than results from the default parameters (see Results). To model this broader tuning, we used a second set of Gabor parameters, identical to those above except that the Gaussian envelope was compressed by a factor of .7 in both length and width. This yields 1.85 subfields, a subfield aspect ratio of 3.18 and a 5% envelope length and width of 1.99^o and 1.15^o respectively.

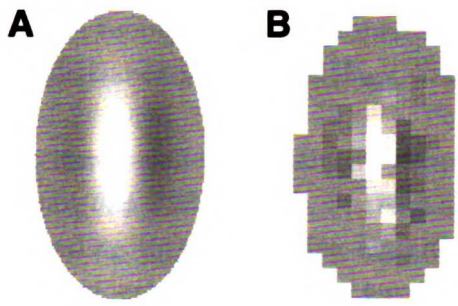

Figure 1.2:

Gabor-shaped cortical RFs. Lighter grays to white: positive values of Gabor function, corresponding to weights of ON-center LGN cells with centers at corresponding spatial positions; darker grays to black: negative values of Gabor function, corresponding to weights of OFF-center cells. **A**. A full Gabor function, used to determine LGN inputs to a cortical cell in the conceptual model. **B**. Typical LGN inputs to a cortical cell in the computational model, after probabilistic sampling from the full Gabor (see Methods). These receptive fields are 'typical'; different cortical cells may have different preferred orientations, spatial phase (relative locations of ON or OFF subregions), spatial location and, in the computational model, different outcomes of the probabilistic sampling. Spatial frequency of sinusoid in Gabor function is 0.8 cycles/degree.

1.2.2 Conceptual Model

To explore the basic concepts underlying our results, we constructed a "conceptual model" designed to be as simple as possible. The model contains two "rate-coded" cortical neurons, one excitatory and one inhibitory; the inhibitory cell inhibits the excitatory cell. The activity of each cell is represented by a scalar value corresponding to average firing rate. The LGN was modeled as a uniform sheet of cells, approximated as a dense lattice (lattice spacing = $.05^o$). The two cortical RFs were determined by Gabor RFs with identical Gaussian shape and location but having sinusoids of opposite spatial phase (thus, the inhibitory cell provides "anti-phase" inhibition).

For computational convenience in obtaining orientation tuning curves, rather than showing many gratings to one pair of cells, we showed one grating to many independent cell pairs. Thus, we constructed multiple pairs of cortical RFs with identical retinotopic positions, and with orientation and spatial phases spaced at 10^o and 20^o intervals respectively.

For each time step, we first calculated the LGN input to each RF by summing LGN firing rates, weighted by the Gabor function, to give the excitatory input $A(\theta, \phi)$ to the cell of orientation θ and phase ϕ. The net input to an excitatory cell with parameters (θ, ϕ) was the weighted sum $A(\theta, \phi) - wA(\theta, \phi + 180^o)$; $A(\theta, \phi + 180^o)$ is the LGN input to the (inhibitory cell) RF having the same orientation but opposite (180^o difference) spatial phase. The inhibitory gain factor, w, is

unitless and represents the transformation from LGN excitatory current to inhibitory spike rate to inhibitory current in the excitatory cell. w is the only free cortical parameter in this model and controls the width of orientation tuning (Fig. 1.5). A match to experimental tuning widths of $\approx 20^o$ is given by $w = 1.5$ for default Gabor parameters (Figs. 1.4, 1.7), and $w = 4.5$ for broadly tuned Gabor parameters.

The output rate of an excitatory cell was obtained by thresholding the net input, *i.e.* spike rate is proportional to $[A(\theta,\phi) - wA(\theta,\phi+180^o) - \xi]^+$. For each set of Gabor parameters, the threshold ξ was set automatically according to the following algorithm (thus, ξ is not a free parameter). For a given level of inhibition w, orientation tuning curves were constructed by determining the peak input over a stimulus cycle for cells of each orientation preference, averaged over cells of all spatial phases. Such curves were obtained for gratings of 5%, 10%, 25% and 50% contrast. Linear interpolation was used to sample these tuning curves at 0.1^o intervals, and the orientation that gave the smallest variance in peak input across contrasts was determined (see Fig. 1.7). The threshold $\xi(w)$ was then set to the average across contrasts of the peak input for that orientation and level of inhibition. The excitatory cell's total response was determined by integrating its activity (calculated every 10 *msec*) over the course of one cycle. A single stimulus cycle was sufficient since the conceptual model is completely deterministic.

The inhibition level w_{best} that gave a best match to experimental tuning widths ($w = 1.5$ or $w = 4.5$ depending on Gabor parameters, as just described) was determined by constructing tuning curves for a range of w. Note that, by the procedure just described, each value of w yields a different threshold $\xi(w)$. To test the robustness of the model to variations in w (Fig. 1.5), for each set of Gabor parameters, we fixed ξ to the level appropriate for w_{best}, and calculated all responses using this fixed threshold.

1.2.3 Computational Model

The majority of simulations were carried out in a "computational model" incorporating details of cortical cells and maps.

Computational LGN Model. For the computational model, a realistically dense lattice of LGN cells was used. We restricted our attention to LGN X-cells, which dominate central cat V1 physiology [Ferster, 1990b]. At 5^o of eccentricity, $1mm^2 = 5^o \times 5^o$ of visual field [Bishop et al., 1962], and retinal ganglion X-cells (X-RGCs) have density $1000/mm^2$ [Peichl and Wassle, 1979], including both ON and OFF cells. We assume each X-LGN cell receives input from a single X-RGC, and each X-RGC projects to 4 X-LGN cells (as in Wörgötter and Koch [1991]; this value is intermediate between values from Sherman [1985] and Peters and Yilmaz [1993]). We thus use 7200 LGN cells to cover 6.8^o x 6.8^o of the visual field, arranged in 4 overlying sheets of ON cells (30 x 30 cells each) and 4 sheets of OFF cells (30 x 30), with ON and OFF lattices offset by 1/2 lattice spacing. After calculating LGN spike rates as above, spikes were produced in a random (Poisson) fashion:

firing rates were converted into the probability of producing a spike in each simulated time step (.25 *msec*). To match data showing correlations among LGN cells with overlapping RFs [Alonso et al., 1996], overlaying cells had 25% correlations in their spike trains (each of 4 overlaying cells picked spikes with probability 1/4 from a common set of 4 Poisson processes). These correlations made no detectable difference in model behavior.

The connection strength to a given cortical cell from each LGN cell was determined by a repeated probabilistic sampling of the Gabor function describing the cortical RF (Fig. 1.2). LGN synaptic strengths were equal to $(\bar{g}^{\mathrm{ff}}_{\mathrm{ex}}/n^{\mathrm{ff}}_{\mathrm{pick}}) \sum_{i=1}^{n^{\mathrm{ff}}_{\mathrm{pick}}} p_i$ where $n^{\mathrm{ff}}_{\mathrm{pick}}=3$, $\bar{g}^{\mathrm{ff}}_{\mathrm{ex}}=0.89\ nS$, and $p_i=1$ with probability determined by the absolute value of the Gabor function and equals 0 otherwise. The number of picks, $n^{\mathrm{ff}}_{\mathrm{pick}}$, determines the degree of sampling of the Gabor function: for $n^{\mathrm{ff}}_{\mathrm{pick}} \to \infty$, the RF becomes a perfect Gabor function. A typical sampled RF is shown in Fig. 1.2. With this sampling, cortical cells received input from 125 ± 8 (mean $\pm$ standard deviation) LGN cells using the default Gabor. Using the more broadly tuned Gabor, cortical cells received input from 61 ± 5 LGN cells.

Cortical Model. Cortical cells were modeled as simple integrate-and-fire neurons as described in Troyer and Miller [1997b,a], with parameters matched to experimental data from McCormick et al. [1985]. Excitatory cells were fitted to responses from regular spiking cells, and inhibitory cells to responses from fast spiking neurons. Briefly, each cell is a single compartment with a capacitance C, leak conductance g_{leak}, resting potential V_{leak}, and two synaptic conductances: fast (AMPA) excitation, g_{ex} (reversal potential $V_{\mathrm{ex}}=0\ mV$); and fast (GABA-A) inhibition, g_{in} ($V_{\mathrm{in}}=-70\ mV$). Excitatory cells also have a spike-triggered adaptation conductance g_{adapt} ($V_{\mathrm{adapt}}=-90\ mV$). Each time varying conductance, g, is modeled as a difference of exponentials: $g(t)=\sum_{t_j<t} \bar{g}\left(e^{-(t-t_j)/\tau^{\mathrm{fall}}}-e^{-(t-t_j)/\tau^{\mathrm{rise}}}\right)$, where the sum is over spike times t_j (presynaptic spike times for $g_{\mathrm{ex}}, g_{\mathrm{in}}$; postsynaptic for g_{adapt}). When V crosses threshold, $V_{\mathrm{thresh}} = -52.5\ mV$, synaptic events are triggered after a delay (randomly chosen for each spike from a uniform distribution, $0.25\ msec \le t_{\mathrm{delay}} \le 2.25\ msec$), adaptation is triggered (exc. cells only), and V is set to V_{reset} and held there for t_{refract}. V_{reset} was fit to the experimentally measured DC gain of cortical cells (the curve of firing rate vs. level of DC injected current; Troyer and Miller, 1997a, 1997b). All cells receive non-thalamocortical background excitatory input (Poisson with a mean rate of 5800 *Hz* and synaptic conductances equal to $\bar{g}^{\mathrm{bg}}_{\mathrm{ex}}$). The magnitude of this input was set to give low mean background firing rates for excitatory cells (.16 Hz) at default values of the parameters; identical background input was given to inhibitory cells and resulted in mean background firing rates of 12.2 Hz. Parameters are, for excitatory cells: $C=500\ pF$, $g_{\mathrm{leak}}=25\ nS$, $V_{\mathrm{leak}}=-73.6\ mV$, $V_{\mathrm{reset}}=-56.5\ mV, t_{\mathrm{refract}}=1.5\ msec$; for inhibitory cells: $C=214\ pF$, $g_{\mathrm{leak}}=18.0\ nS$, $V_{\mathrm{leak}}=-81.6\ mV$, $V_{\mathrm{reset}}=-57.8\ mV$, $t_{\mathrm{refract}}=1.0\ msec$; for conductances: $\tau^{\mathrm{rise}}_{\mathrm{ex}}=0.25\ msec$, $\tau^{\mathrm{fall}}_{\mathrm{ex}}=1.75\ msec$, $\tau^{\mathrm{rise}}_{\mathrm{in}}=0.75\ msec$, $\tau^{\mathrm{fall}}_{\mathrm{in}}=5.25\ msec$, $\tau^{\mathrm{rise}}_{\mathrm{adapt}}=1\ msec$, $\tau^{\mathrm{fall}}_{\mathrm{adapt}}=83.3\ msec$, $\bar{g}_{\mathrm{adapt}}=3\ nS$, $\bar{g}^{\mathrm{bg}}_{\mathrm{ex}}=0.89\ nS$. $\bar{g}^{\mathrm{ctx}}_{\mathrm{ex}}$, $\bar{g}^{\mathrm{ff}}_{\mathrm{ex}}$, and $\bar{g}_{\mathrm{in}}$ were free parameters and set as described below.

The model contains 1600 excitatory and 400 inhibitory layer 4 simple cells, representing a 2/3mm x 2/3mm patch of cortex and $.75^o x .75^o$ in visual angle ($0.9mm = 1^o$ of visual field at 5^o eccentricity [Tusa et al., 1978]). A 20 x 20 grid of inhibitory cells was interspersed within a 40 x 40 grid of excitatory neurons, with each inhibitory RF center aligned with every other excitatory cell. Gabor-shaped RFs were defined by three parameters in addition to those described above: preferred orientation, determined by an optically measured cortical map from cat V1 (provided by Michael Crair and Michael Stryker – shown in Fig. 1.8A); retinotopic position, progressing uniformly across the sheet; and spatial phase, assigned randomly to each cell [DeAngelis et al., 1992, Ghose et al., 1993].

The probability that any two cortical cells were connected depended on the correlation between their RFs. The following scheme was used for both excitatory and inhibitory connections. Raw correlation $c'(a,b)$ between RFs of cortical cells a, b is $c'(a,b) = \sum_{i,j \in \mathrm{LGN}} w(a,i) w(b,j) c(i,j)$. Here, i, j are LGN cells, $w(a,i)$ and $w(b,j)$ are the thalamocortical weights from i to a and j to b, and $c(i,j)$ is the cross-correlation of the spatial RFs of i and j, where OFF spatial RFs are negative of ON. Correlation is then $c(a,b) = c'(a,b)/\sqrt{c'(a,a)c'(b,b)}$. A connectivity function $C(a,b)$ – roughly, the probability of a connection from a to b – is defined as $C(a,b) = [\mathrm{sgn}(a) c(a,b)^{n_{\mathrm{pow}}}]^+$ where $\mathrm{sgn}(a) = 1$ if a is excitatory, -1 if a is inhibitory; $[x]^+ = x$, $x > 0$, $[x]^+ = 0$ otherwise. n_{pow} is a parameter that determines connectivity strength as a function of correlation. Smaller values of n_{pow} lead to broader connectivity and more intracortical connections per cell; larger values have the opposite effect (see Fig. 1.8B). At the default value, $n_{\mathrm{pow}} = 6$, a cortical cell receives connections from 132 ± 38 (mean $\pm$ standard deviation) other cortical cells (80% from excitatory cells, 20% from inhibitory cells, on average). Just as the thalamocortical connections were sampled from the Gabor function, the intracortical connections were sampled from $C(a,b)$: the strength of intracortical connection from a to b, $g(a,b)$, is $g(a,b) = (\bar{g}/n^{\mathrm{ctx}}_{\mathrm{pick}}) \sum_{i=1}^{n^{\mathrm{ctx}}_{\mathrm{pick}}} p_i$, where $p_i = 1$ with probability $C(a,b)$ ($\bar{g} = \bar{g}^{\mathrm{ctx}}_{\mathrm{ex}}$ or $\bar{g} = \bar{g}_{\mathrm{in}}$; $n^{\mathrm{ctx}}_{\mathrm{pick}} = 10$). As $n^{\mathrm{ctx}}_{\mathrm{pick}} \to \infty$, the connectivity becomes exactly $\bar{g} C(a,b)$.

The main parameters controlling model behavior were the total synaptic strength for each type of connection: thalamocortical (LGN); intracortical excitation (e→{e,i}), and intracortical inhibition onto excitatory cells (i→e). The total synaptic strength is obtained by (i) assuming the cell is voltage clamped at threshold; (ii) for each synapse, integrating over time the synaptic current induced by one presynaptic spike; and (iii) summing over all synapses of the given type. Thus, total synaptic strength is expressed in units of $nA\ msec$. The parameters were chosen to satisfy various experimental constraints such as orientation tuning width. We used two different parameter sets: the "feedforward" set with LGN and intracortical inhibitory connections only, and the "full circuit" set which also included feedback intracortical excitation. For simplicity, inhibitory cells received only excitation; we have not yet explored the influence of inhibitory-to-inhibitory connections. For most simulations, the total intracortical excitatory synaptic strength onto each excitatory cell (e→e

connections) and onto each inhibitory cell (e→i connections) was identical. Some simulations were run with intracortical excitatory connections onto excitatory cells only (e→e, but no e→i). After determining the pattern of synaptic strengths by probabilistic sampling, synaptic conductances were multiplicatively scaled so that the total conductance from each synaptic type received by each cell was set to its respective mean across cells. This avoids large differences in the amount of input to different cells resulting from the unequal representation of orientations in our spatially limited sample of an orientation map. For the "feedforward" parameter set, (Figs. 1.3, 1.4, 1.7), total synaptic strengths received by a cell from each type of connection were 10 $nA\ msec$ (LGN), and 3.75 $nA\ msec$ (i→e), yielding mean values for unitary conductances of $\bar{g}_{\mathrm{ex}}^{\mathrm{ff}} = 2.1\ nS$, $\bar{g}_{\mathrm{in}} = 8.3\ nS$. For the "full circuit" parameter set, (Figs. 1.8-1.12), total synaptic strengths received by a cell from each type of connection were 5 $nA\ msec$ (LGN), 4.25 $nA\ msec$ (e→{e,i}), and 7.5 $nA\ msec$ (i→e), yielding mean values for unitary conductances of $\bar{g}_{\mathrm{ex}}^{\mathrm{ctx}} = 2.0\ nS$, $\bar{g}_{\mathrm{ex}}^{\mathrm{ff}} = 1.0\ nS$, and $\bar{g}_{\mathrm{in}} = 16.6\ nS$. Effects of varying these values were also explored (see Fig. 1.13). Note that we have realistic numbers of LGN cells but unrealistically small numbers of cortical cells, therefore intracortical connections are unrealistically strong relative to thalamocortical.

Simulations. A typical simulation consisted of 3 cycles of a 3 Hz sinusoidal grating. During each time step (.25 $msec$), values for time varying conductances were updated, and the membrane time constant and the equilibrium voltage for each cell were then calculated from the cell's conductances. Each cell's voltage was then adjusted according to an exponential decay. Finally, threshold crossings were detected and subsequent synaptic, adaptation and refractory events were registered. Simulations were written as C subroutines (mex files) in the MATLAB simulation environment. Initial conditions were determined by simulating 1 sec of model behavior at default parameter values and with LGN cells at background firing rates.

All orientation preferences are represented in the cortical network. Orientation tuning curves were constructed from the presentation of a single stimulus, by binning responses from all cells in the network according to their preferred orientation in 10^o bins. Most results used as a stimulus a grating oriented at 128^o. This orientation was chosen to avoid artifacts that might result from alignment of the stimulus with the axes of the LGN grid, but we saw no evidence of such behavior.

When displaying synaptic conductances and currents, we show "stimulus induced" curves in which we have subtracted the mean values of these conductances and currents at background. These mean values were determined by running "blank stimulus" trials in which LGN firing rates were unmodulated.

To reproduce the results of Nelson et al. [1994], we ran simulations in which the inhibition and adaptation currents were blocked in a single cell (Fig. 1.12). To accomplish this in a computationally convenient way, we ran a single simulation without any blockade, but monitored the behavior of an additional "blocked cell" for each cell in the network. The blocked cell made no connections. It received identical excitatory input as its unblocked partner cell, but had no inhibitory or adaptation

current and was injected with sufficient hyperpolarizing current to bring the background firing rates back to normal. Thus each blocked cell received input from a network in which all other cells were normal (unblocked), but did not itself affect any other cells in the network. Under the assumption that altering a single cell does not affect network behavior, this method allows us to simulate numerous experiments in which one cell undergoes intracellular inhibitory blockade.

1.3 Results

1.3.1 Modeling Approach

We pursued two parallel approaches to modeling contrast-invariant orientation tuning. In order to explore the basic ideas underlying such tuning, we constructed a "conceptual" model, designed to be as simple as possible. This model considered two cortical simple cells, one excitatory and one inhibitory, with a monosynaptic connection from inhibitory to excitatory. The RFs of the two cells had identical position and preferred orientation, but opposite spatial phase (see Methods). The neurons were "rate-coded": the average firing rate of each cell was determined by a linear thresholding operation applied to the weighted sum of input cell firing rates. For simplicity, the inhibitory threshold was set to zero (*i.e.* the inhibitory cell's response was a linear function of its input). The excitatory cell's threshold was set automatically to the level that best produced contrast-invariant tuning for contrasts of 5% and above (see Methods). Therefore, after determining the structure of the cortical receptive fields, the conceptual model had only a single free parameter: the strength of intracortical inhibition relative to the strength of thalamocortical excitation.

In order to study robustness of our ideas to the complexity of real cortical circuits, we also constructed a "computational" model that incorporated known details of cortical cells and maps. The cortical component of this model consisted of 1600 excitatory and 400 inhibitory layer 4 simple cells, arranged in a 2/3mm x 2/3mm cortical sheet. Preferred orientations were determined by a measured V1 map, and intrinsic connectivity was determined probabilistically based on correlations in input RFs. Excitatory and inhibitory cells were modeled as conductance-based integrate-and-fire neurons, with parameters matched to those measured in cortical regular-spiking and fast-spiking cells respectively, including a spike-rate adaptation current in the excitatory cells [McCormick et al., 1985, Troyer and Miller, 1997b,a] (details in Methods). We considered only the effects of fast synaptic conductances (AMPA and GABA-A); the role of slow conductances (NMDA and GABA-B) will be explored in future work.

1.3.2 LGN Input

We focused our research on the response to full-field sinusoidal gratings, as these are the only stimuli for which contrast dependence of orientation tuning has been studied [Sclar and Freeman, 1982, Skottun et al., 1987]. Our model was based on cat V1 at approximately 5^o eccentricity.

Circularly symmetric, center-surround LGN spatial receptive fields were used [Peichl and Wassle, 1979, Linsenmeier et al., 1982], and LGN firing rates were determined as rectified linear filterings of the input luminance using experimentally measured contrast gain curves (see Methods and Fig. 1.1B) [Peichl and Wassle, 1979, Cheng et al., 1995]. To determine whether our model would yield well-tuned responses to transient stimuli, we also modeled responses to moving bars.

LGN cells responded to sinusoidal grating stimuli with a sinusoidal modulation in firing rate (Fig. 1.1A). The temporal responses of ON-center and OFF-center cells with spatially overlapping RFs were 180^o out of phase. Increasing stimulus contrast resulted in a larger modulation of firing rate. At contrasts above about 5%, the spike rate modulation exceeds the background firing rate. For these contrasts, responses are no longer purely sinusoidal, because spike rate cannot be negative (Fig. 1.1A, solid lines); that is, LGN responses rectify. Once responses rectify, mean ("DC") firing rates increase with increasing contrast (DC curves in Fig. 1.1B), because peak firing rates continue to increase while minimal firing rates cannot decrease below zero. This contrast-dependent increase in mean LGN firing rates has important consequences for contrast-invariant orientation tuning that will be discussed in detail below.

The oriented arrangement of LGN inputs to simple cell RF subregions was modeled using a Gabor function – a two dimensional Gaussian multiplied by a sinusoid (Fig. 1.2A). In the conceptual model, the Gabor function directly determined the weights of geniculocortical connections: positive values corresponded to the weights of ON-center inputs, negative values to the weights of OFF-center inputs (Fig. 1.2A). In the computational model, geniculocortical synaptic strengths were determined by probabilistic sampling of the Gabor function from a realistically dense lattice of LGN cells (Fig. 1.2B).

We considered two different sets of Gabor parameters to describe geniculocortical connections. The first set was matched to RF parameters taken from physiological measurements of cat simple cells [Jones and Palmer, 1987]. The use of the Jones and Palmer parameters as a measure of LGN connectivity in simple cells is based on the experiments of Reid and Alonso [1995], which show that physiological RF parameters at least roughly correspond to the pattern of geniculocortical connections in cat layer 4. These will be used as our "default" parameters. We also consider a second set of parameters representing more broadly tuned LGN input, for several reasons. If cortical circuitry plays a significant role in sharpening simple cell orientation tuning, then the LGN input to a cell would have broader tuning than the cell's responses. Furthermore, the parameters of Jones and Palmer [1987] represent an average of simple cells from all layers, whereas layer 4 cells may on average be more broadly tuned for orientation than other layers [Tolhurst and Thompson, 1981]. We base our more broadly tuned parameter set on the experiments of Ferster et al. [1996], who cooled the cortex to largely eliminate cortical inputs. Using intracellular electrodes, they then measured the direct LGN input for gratings presented at 30^o intervals. The tuning of this input was quantified by measuring the first harmonic (F1) of the voltage response, as a function of stimulus

orientation. Although orientation was sampled only coarsely, the figures presented in Ferster et al. [1996] show average orientation tuning halfwidth at half height (HWHH) of approximately 35^o. This is significantly broader than the input F1 tuning under our default Gabor parameters, which we find to be 24^o. To mimic the broader tuning observed by Ferster *et al.*, we artificially shrunk the default RFs by a factor of 0.7, leaving the width of each subregion unchanged. This resulted in an input F1 tuning width of 34.8^o.

In the conceptual model, the excitatory and inhibitory cells had identical Gabor RFs except that their sinusoids were 180^o out of phase. In the computational model, a distribution of receptive fields was obtained from variations in three parameters: preferred orientation, determined by a measured cortical map (Fig. 1.8A); retinotopic position, progressing uniformly across the sheet; and spatial phase, assigned randomly to each cell [DeAngelis et al., 1992].

1.3.3 Tuning of the LGN Input to a Simple Cell

At the preferred orientation, the bright and dark portions of a sinusoidal grating stimulus align with the cortical cell's ON and OFF subregions simultaneously. Thus, all of the cortical cell's LGN inputs fire relatively synchronously and the temporal modulation of this input is large (Fig. 1.3A). At the null orientation, the inputs are stimulated asynchronously, so the temporal modulation of the total input is small. Note that the mean rate of LGN input does not depend on stimulus orientation. This follows from the assumption that LGN cells are untuned for orientation: since the mean LGN input received by a simple cell is the (weighted) sum of the mean rates of the LGN cells projecting to it, this mean input must also be untuned for orientation [Ferster, 1987]. Therefore, only the temporally modulated component of the LGN input is orientation tuned.

The untuned mean input presents the primary problem for a purely thalamocortical explanation of contrast-invariant orientation tuning. Mean LGN firing rates increase with increasing contrast due to LGN rectification (Fig. 1.1). This contrast-dependent increase in firing rate is sufficiently large that the *mean* LGN input at the null orientation at high contrasts exceeds the *peak* LGN input at the preferred orientation at low contrast (Fig. 1.3). No single spiking threshold level can yield well-tuned responses for stimuli of all contrasts.

Therefore, to achieve contrast-invariant orientation tuning in response to sinusoidal gratings, the cortex must cancel the untuned, mean input component in a contrast-dependent manner, while extracting the tuned, modulated component. We will show that this decomposition of the input into a tuned and an untuned component generalizes to stimuli such as flashed and moving bars.

1.3.4 Anti-Phase Inhibition Can Achieve Contrast-Invariant Orientation Tuning

The main result of this paper is to demonstrate that correlation-based intracortical inhibition can achieve contrast-invariant orientation tuning (the effects of correlation-based intracortical excitation

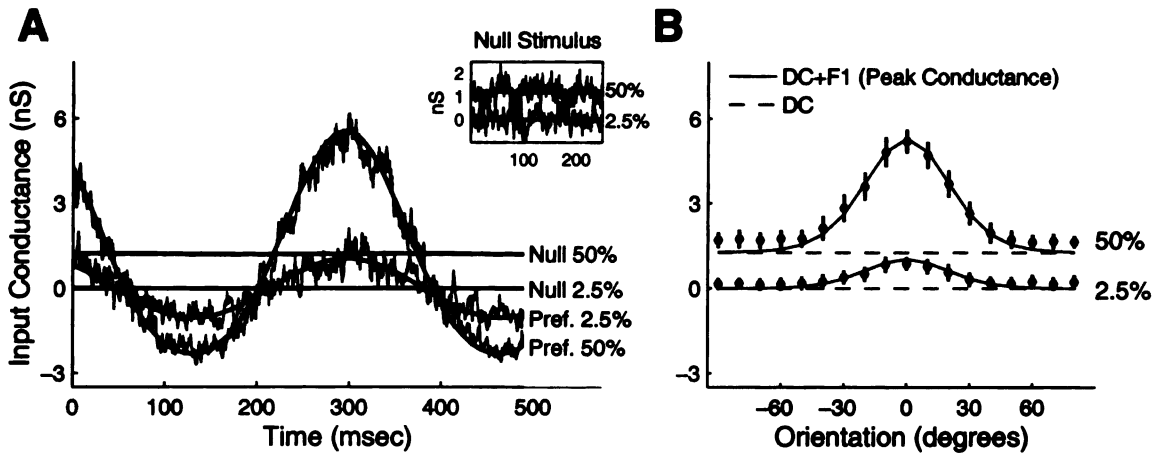

Figure 1.3:

Tuning of total LGN input. **A.** Input to cortical cells in response to gratings at the preferred and null (orthogonal to preferred) orientations. Upper traces: high (50%) contrast; lower traces: low (2.5%) contrast. Curved traces show input in response to preferred orientation: black traces: average input (40 presentations) from computational model, using a sampled Gabor RF (as in 2B); gray curves: input for conceptual model, using connections from the full Gabor function (2A). Gray straight lines show response in the conceptual model to a stimulus at the null orientation; in inset, these lines are repeated and compared to average input to null stimuli in computational model (black traces). Note that input to null stimulus at 50% contrast typically exceeds peak input to preferred stimulus at 2.5% contrast. Agreement of the two models for both preferred and null stimuli indicates RF sampling and Poisson firing of LGN inputs have little effect. **B.** Tuning of mean (dashed lines) and mean plus first harmonic (solid lines) of thalamic input conductance. Lines show results from the conceptual model; solid circles show results from the computational model; error bars represent ± 1 std. dev. Sum of mean plus first harmonic represents peak input during a cycle of the grating stimulus. Note that mean input is untuned for orientation, and mean input at high contrasts exceeds peak input to preferred orientation at low contrasts. Thus, although the first harmonic is well tuned, no single spike threshold can give tuned responses at both high and low contrasts. In this and subsequent figures showing orientation tunings, cells are grouped by preferred orientation in 10^o bins, and orientation axis represents difference of stimulus orientation from preferred. The plot shows line connecting means for each bin; error bars are over cells within a bin. See Methods for more details.

will be considered below). By correlation-based inhibition, we mean that the probability of a connection from an inhibitory cell to an excitatory cell is an increasing function of the degree of anticorrelation between their RFs, *i.e.* the strongest inhibitory connections are made between cells with the most anti-correlated RFs (see Methods). This implies that an excitatory cell receives the strongest inhibition from inhibitory cells with identical Hubel-Wiesel RFs but of opposite spatial phase. We will call such an inhibitory neuron the cell's "anti-phase partner". (By "spatial phase" of an RF, we refer to absolute position in visual space of the ON or OFF subregions, rather than to their position relative to each cell's Gabor function; thus, two RFs have "opposite spatial phase"

if the ON-subregions of one tend to overlap the OFF-subregions of the other in visual space). The existence of such "spatially opponent" or "anti-phase" inhibition in cat layer 4 is well-supported experimentally: at ON locations, where a light stimulus evokes excitation (EPSPs), dark stimuli evoke inhibition (IPSPs), and vice versa for OFF locations [Palmer and Davis, 1981, Ferster, 1988, Hirsch et al., 1995a]. Note that, since simple cells with orthogonal orientation preference have weakly- or un-correlated RFs, correlation-based connectivity results in little or no inhibition from cells with orthogonal tuning. Instead, inhibition comes from cells with similar preferred orientations.

Our model is not a developmental model: we first determined the pattern of LGN input to cortical cells, and then fixed the pattern of intracortical connections according to the above correlation-based rule. However, this pattern of inhibition would be expected to arise from a Hebb-type synaptic modification rule, generalized to apply to inhibitory synapses. Such a rule states that synaptic strengths grow more negative (more strongly inhibitory) when pre- and postsynaptic firings are anticorrelated; or equivalently, that synapses strengthen when they are effective, *i.e.* when the inhibitory presynaptic cell is active while the postsynaptic cell is inactive. Such generalization of Hebb-type learning rules to inhibitory synapses is only a hypothesis; plasticity of inhibitory synapses is not well understood (but see Komatsu [1996]). (This intracortical connectivity could also emerge without inhibitory synaptic plasticity: in models in which only thalamocortical synapses undergo correlation-based plasticity, presence of a fixed inhibitory connection from one cortical cell to another tends to cause the two to develop anticorrelated thalamocortical RFs [Miller, 1994].)

The sufficiency of correlation-based inhibition for contrast-invariant tuning is demonstrated in Fig. 1.4, which shows tuning curves for both the computational and conceptual models, for gratings of 2.5%, 5%, 10%, 25%, and 50% contrast. Both models display contrast-invariant orientation tuning above 5% contrast. By choosing the appropriate level of inhibition, both models were able to match experimental estimates of mean orientation tuning width for simple cells. For example, Heggelund and Albus [1978] report that simple cells have mean tuning width (HWHH) of 19.5°. Model tuning widths (HWHH) above 5% contrast were between 18.7° and 20.8° for both the computational and conceptual models.

The width of the tuning is largely determined by the strength of the inhibition and the tuning of the LGN input. Fig. 1.5 shows cortical tuning halfwidths, using either narrowly or broadly tuned LGN inputs, for a variety of levels of inhibition. Tuning narrows with stronger inhibition, but remains contrast-invariant above 5% contrast. Tuning to a long moving bar (width 0.62°, velocity $3.75^o/sec$) is slightly broader but shows identical sharpening with increasing levels of inhibition: filled circles show bar tuning at 50% contrast for the narrowly tuned LGN inputs.

For higher levels of inhibition and broadly tuned input, tuning at 5% contrast narrows slightly in the conceptual model. This is due to the fact that spike threshold was optimized for the default level of inhibition (see Methods), and could be corrected if spike threshold were separately optimized for each set of parameters. At 2.5% contrast and high levels of inhibition (Fig. 1.5B, thin lines), the

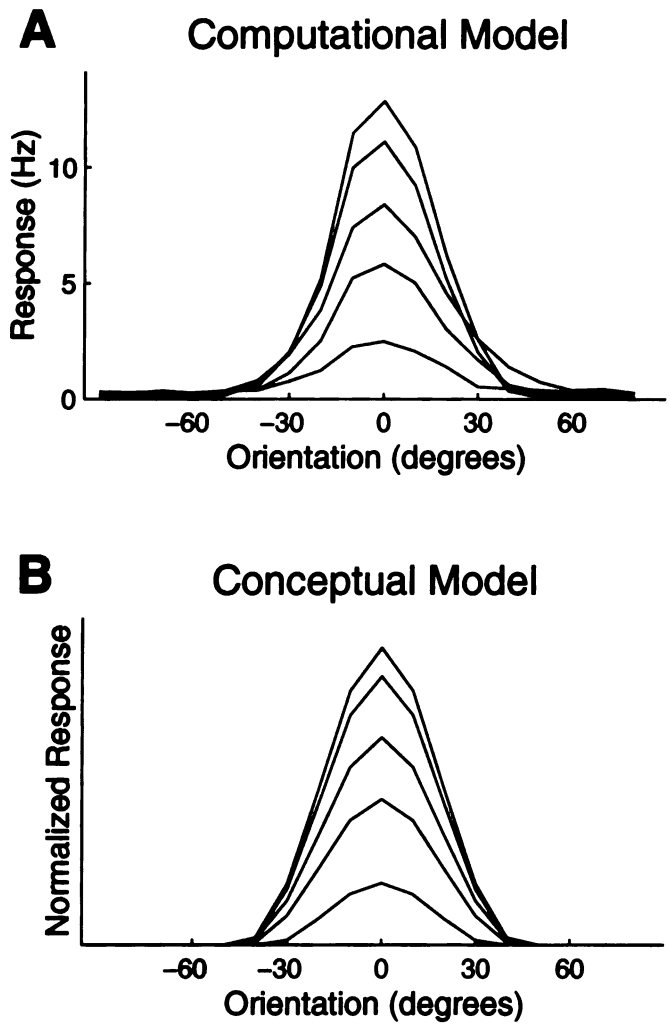

Figure 1.4:
Contrast-invariant tuning. Response vs. orientation for gratings of 2.5%, 5%, 10%, 25% and 50% contrast. **A.** Computational model. **B.** Conceptual model. Both models yield contrast-invariant tuning at 5% contrast and above.

conceptual model predicts much narrower tuning, for reasons that are more general (see below).

The conceptual and computational models yield qualitatively similar results. Simple additions to the conceptual model led to progressively closer quantitative matches to computational model behavior. A significantly improved match was obtained by adding inhibitory thresholds and using correlation-based inhibitory connectivity from cells with a range of RF properties (rather than from only the single cell with precisely opposite spatial phase). Using simulated synaptic noise (and hence changing the threshold linear function to a smoother function near spike threshold) led to an even closer match between the models, and nearly eliminates the difference in responses to 2.5% contrast gratings (see below). However, incorporation of these features required additional unconstrained parameters and we began to lose the simplicity that was the strength of the conceptual model. Therefore, the results of these investigations are not further reported.

The behavior of our correlation-based model is presented below, in three steps. First, we

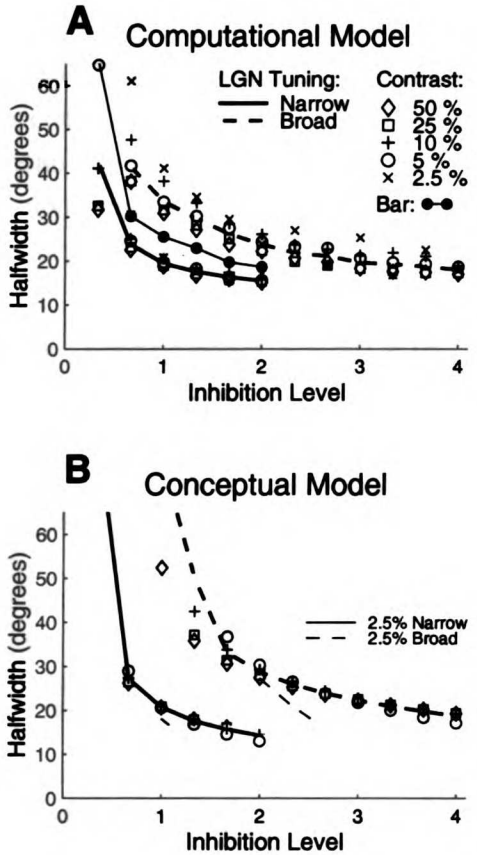

Figure 1.5:

Increasing inhibition leads to sharper tuning. Tuning halfwidth at half height (HWHH) vs. level of inhibition for gratings of 2.5%, 5%, 10%, 25% and 50% contrast. Thick solid (lower) curve: mean tuning HWHH above 5% for RFs with large subfield aspect ratios and narrow LGN tuning (matched to data from Jones and Palmer, 1987); Thick dashed (upper) curve: mean tuning HWHH for RFs with small subfield aspect ratios and broad LGN tuning (matched to data from Ferster et al. [1996]). Level of inhibition is normalized so that 1 is the level that produces physiological halfwidths for narrow LGN input (Fig. 1.4). Overlapping symbols indicate contrast-invariance. Tuning gradually sharpens with increased levels of inhibition. **A**. Computational model. **B**. Conceptual model. In conceptual model, tuning narrows slightly at 5% contrast for large levels of inhibition due to the fact that spike threshold is optimized for default parameters, *i.e.* inhibition level of 1 (see Methods). Responses to 2.5% contrast gratings at high inhibition levels for both narrow (solid) and broad (dashed) LGN tuning are shown using thin lines. At very low contrast, conceptual model predicts much narrower tuning.

analyze the reasons why anti-phase inhibition achieves contrast-invariant tuning, using the simple conceptual version of the model. Second, we incorporate correlation-based intracortical excitation into the computational model, and present results from this completed model. Finally, we explore the robustness of this computational model to variations in the key parameters controlling model behavior. A schematic representing the behavior of both models is shown in Fig. 1.6, and will be referred to throughout the text.

1.3.5 Conceptual Model: Anti-Phase Inhibition Cancels the Untuned Component of the Input

Recall that the main obstacle to achieving contrast-invariant tuning is the untuned component of the LGN input, which increases with contrast due to the rectification of LGN responses at higher contrasts. The ability of anti-phase inhibition to overcome this problem is most easily demonstrated in the context of the two-cell conceptual model. Here we introduce the term "feedforward", by which we mean input from LGN to cortical cells not mediated by cortical excitatory cells. Thus, the geniculocortical input represents feedforward excitation, while the pathway from LGN to cortical inhibitory cell to cortical excitatory cell represents feedforward inhibition.

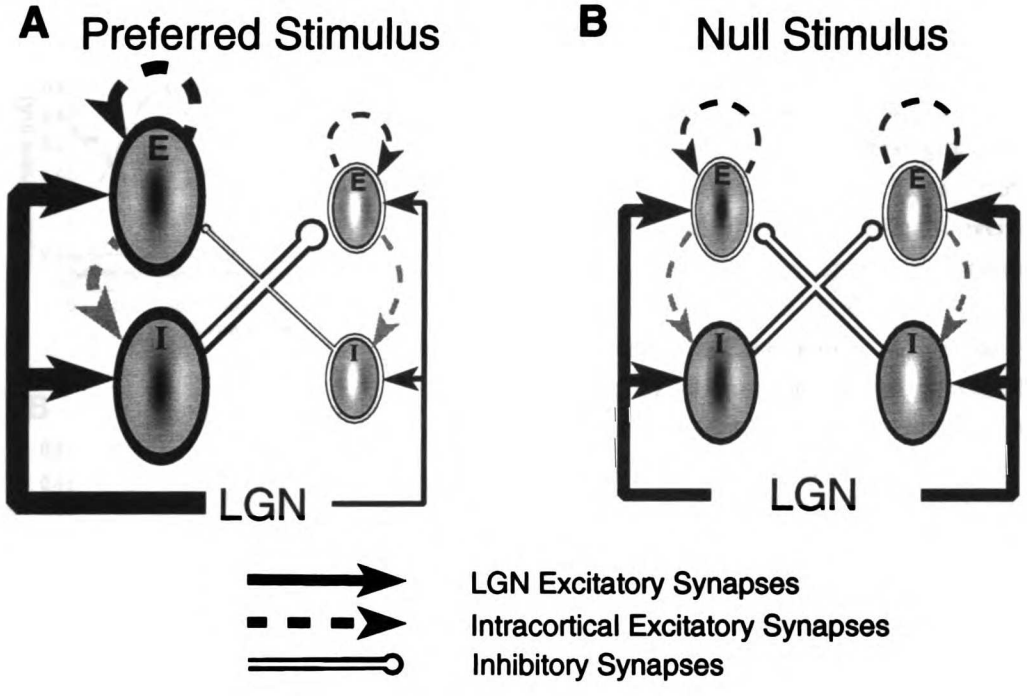

Figure 1.6:

Behavior of model using correlation-based connectivity. Schematic representing behavior of the model in response to preferred (**A**) and null (**B**) stimuli. The excitatory cell described in the Results is in the upper left; its inhibitory "anti-phase partner" is in the lower right. E: excitatory cells; I: inhibitory cells. Solid lines: excitation and depolarization; open lines: inhibition and hyperpolarization. Line thickness and size of RF icon represent magnitude of activity. Dashed lines represent correlation-based excitation, which is included in the complete computational model only (Figs. 1.8-1.11). Some simulations were performed without cortical excitatory projections onto inhibitory neurons (gray dashed lines), but this did not substantially affect network behavior (see Fig. 1.13B).

Suppose an excitatory simple cell receives total input A^e from the LGN, while its inhibitory anti-

phase partner receives LGN input A^i. Assuming for simplicity that inhibitory cell response is linear, the total feedforward input to the excitatory cell is $A^e - wA^i$, where $w > 1$ is the total strength of the inhibitory synaptic connection multiplied by the gain of the inhibitory cells. During the peak response to the preferred orientation, LGN excitation A^e is large, while the anti-phase inhibition wA^i is weak (Figs. 1.6A,1.7A-top). Thus, the cell gives a strong response. At the null orientation, cells at all spatial phases are receiving an intermediate level of feedforward excitation $A^e \approx A^i$, and the inhibition $wA^i > A^e$ is sufficient to prevent excitatory cell spiking (Figs. 1.6B,1.7A-bottom). Since A^i and A^e both rise with contrast at the same rate, the dominance of inhibition over excitation is maintained for null stimuli of all contrasts.

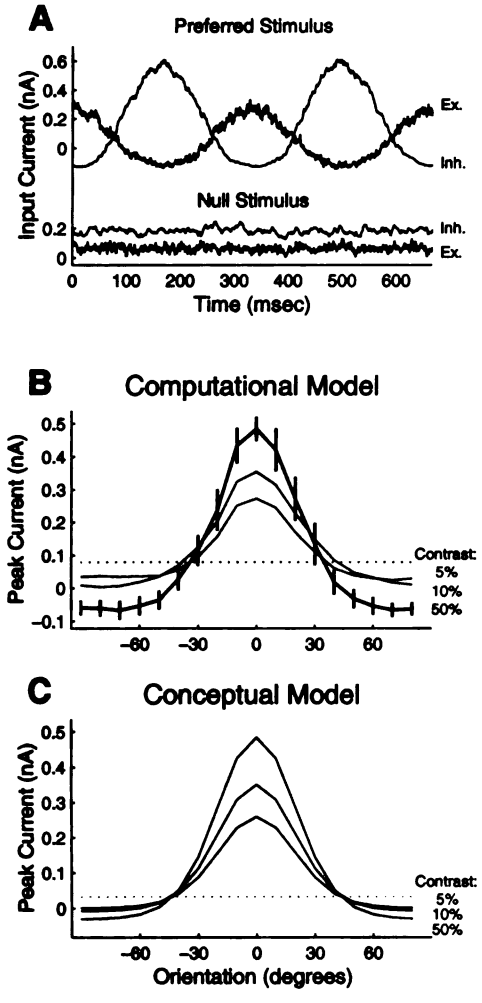

Figure 1.7:

Inputs to a cortical cell given anti-phase inhibition (inputs shown relative to background). **A.** Averaged computational model responses (40 presentations) to 50% contrast gratings. Excitatory LGN input marked 'Ex'; intracortical inhibitory input marked 'Inh'. To compare excitatory and inhibitory inputs, synaptic conductances were converted to currents obtained if the cell was voltage clamped at threshold. **B.** Peak synaptic current vs. orientation for computational model. Responses are to single presentations of 50%, 10% and 5% contrast gratings at 128°. Peak current is the first harmonic (F1) plus the mean (DC) of the stimulus-induced current (including excitation and inhibition). Error bars for 50% contrast are ± 1 std. dev. Dotted line shows approximate threshold level that would lead to contrast-invariant tuning; actual threshold in computational model is determined independently from *in vitro* data (see Methods). **C.** Peak synaptic current vs. orientation for conceptual model. Since there is no noise, true peak current is shown. Thin line shows automatically selected threshold (see Methods). For both models, mean input decreases and modulation increases with contrast. Thresholds near the crossover point of net input tuning curves result in sharp, contrast-invariant tuning.

The ability of anti-phase inhibition to achieve contrast-invariant tuning for a wide variety of

stimuli can be best understood by dividing the LGN input into two components: the phase-nonspecific component $A_{non} = (A^e + A^i)/2$ – the average of the input to the cell and to its anti-phase partner – and the remaining phase-specific component, $A_{spec} = (A^e - A^i)/2$. The total input to the cell, $A^e - wA^i$, can then be rewritten $(1-w)A_{non} + (1+w)A_{spec}$. Thus, anti-phase inhibition acts to eliminate the phase-nonspecific component of the LGN input, while amplifying the phase-specific component. For all of the commonly presented oriented stimuli (moving or flashed bars, flashed or counterphased gratings), Hubel-Wiesel RFs yield a phase-specific component tuned for orientation, and a phase-nonspecific component that is nearly or completely untuned. Thus, the effectiveness of the anti-phase model in achieving contrast-invariant tuning generalizes across stimuli.

This can by summarized by noting that the schematic circuit (Fig. 1.6) acts as a "differential phase filter": with inhibition sufficiently large, any stimulus that gives similar excitation to each of two opposite phases will cause more inhibition than excitation in excitatory cells and hence will be "filtered out". Only stimuli that predominantly excite one phase and not its opposite can "pass" through this "filter" and cause the excitatory cells to fire. The only stimuli that can accomplish this are stimuli near the preferred orientation; stimuli far from the preferred will give similar input to both phases. This argument applies to any type of oriented stimulus.

1.3.6 Conceptual Model: Dominant Anti-Phase Inhibition Provides a Contrast-Dependent Effective Threshold

Although the most important effect of anti-phase inhibition is to eliminate the phase-nonspecific component of the LGN input, this is not sufficient to achieve contrast-invariant tuning. This can be seen by setting $w = 1$, thereby causing $(1-w)A_{non} = 0$. In this case, contrast invariance can only be achieved if spike threshold is negligible, *i.e.* if any positive input leads to spiking. Otherwise, orientations that at low contrast give positive but sub-threshold phase-specific input would yield spike responses at higher contrast, because A_{spec} grows with contrast; thus, orientation tuning would broaden with contrast.

This problem is remedied by including relatively strong inhibition, ($w > 1$). Then the phase-nonspecific component $(1-w)A_{non}$ has a net inhibitory influence that increases with contrast. Since the phase-nonspecific input is untuned for orientation, it serves as a "plateau" – an input identical for stimuli of all orientations – to which the orientation-tuned, phase-specific component is added. The distance from this plateau to the cell's spike threshold can be thought of as a contrast-dependent effective threshold for the tuned input component [Bonds, 1989, Ben-Yishai et al., 1995]. With $w > 1$, this plateau is inhibitory and moves further from spike threshold with increasing contrast (Fig. 1.7B,C). By "pulling down" the tuned component, so that only a portion of it is above the spike threshold, this inhibition serves to sharpen the spiking orientation tuning relative to the tuning of the phase-specific input. If spike threshold falls near the crossover point

of the net input tuning curves for varying contrasts (Dotted lines in Fig. 1.7B,C), this inhibition sharpens the feedforward input in a contrast-invariant manner.

In the conceptual model, spike threshold for excitatory cells was automatically set at this crossover point in the input current (see Methods). Somewhat surprisingly, we have found that in the computational model, simply using a physiologically-based spiking neuron model (Troyer and Miller, 1997a, 1997b) was adequate to robustly attain contrast-invariant tuning; no parameter adjustments were required. One possible explanation is that synaptic noise "smears out" spike threshold, making it relatively easy to match "threshold" with the crossover. Also, with inhibition dominant, the orientation tuning curves cross one another where input changes rapidly as a function of orientation, so moderate changes in threshold should make little difference in tuning. Simulations with the conceptual model show that moving threshold by as much as 10% of the peak-to-peak variation in the input driven by 5% gratings changes tuning by less than 1.5^o.

At very low contrasts, the conceptual model predicts that orientation tuning will narrow. Below about 5% contrast, LGN responses do not rectify and therefore the plateau, $(1 - w)A_{non}$, does not change with contrast. Orientation tuning narrows with further decreases in contrast (Fig. 1.5B), because the tuned input component is reduced while the nonzero "effective threshold" is left unchanged. It is unclear whether one could expect to see narrower tuning in the experimental data. As mentioned above, synaptic noise eliminates sharp thresholds, and the effect may be lost in the noise. Computational model results bear this out: tuning for 2.5% contrast has HWHH similar to that at higher contrast (Fig. 5C). This conclusion is further supported by simulations in which synaptic noise was added to the conceptual model. As mentioned above, in this case the conceptual model behavior matched the broader tuning of the computational model, even at 2.5% contrast (not shown).

1.3.7 Computational Model: Adding Correlation-Based Excitation

To this point, we have not considered the effect of intracortical excitation. We have seen that correlation-based inhibition is sufficient to achieve sharp, contrast-invariant tuning. Here we show that the addition of correlation-based excitation "amplifies" these contrast-invariant responses, without altering their tuning. The conceptual model, which contains only two cortical neurons, is too simple to explore the effects of intracortical excitation in any meaningful way. Hence, the remainder of the paper will present results from the computational model only.

Intracortical excitation was incorporated using a correlation-based rule analogous to that used for intracortical inhibition: excitatory connections were determined probabilistically, such that the strongest connections are found between cells whose RFs are most strongly correlated, *i.e.* those with similar preferred orientation and *similar* spatial phase. This is illustrated schematically by the dashed lines in Fig. 1.6. That intracortical excitation comes primarily from cells of similar orientation preference and similar spatial phase is supported by the fact that EPSP's are evoked

only by stimuli of appropriate position and phase, with opposite phase to the stimuli that evoke IPSP's [Ferster, 1988, Hirsch et al., 1995a]. More direct support is provided by Freeman et al. [1997] who recorded from pairs of cat V1 simple cells isolated on a single electrode. Cell pairs had similar preferred orientations but randomly varying spatial phases. However, cross-correlations indicative of a monosynaptic excitatory connection were found only when the cells had similar absolute spatial phase (G. Ghose, private communication).

Since the dependence on correlation of intracortical inhibition and excitation differs only in sign, excitatory and inhibitory connections in our model have *precisely* the same average distribution in terms of orientation preference; they differ only in spatial phase. An example is shown in Fig. 1.8A, which illustrates the experimental V1 orientation map used to assign preferred orientations to cortical cells in the computational model. In this figure, white squares show the locations of cells making excitatory connections to the excitatory cell at the X, while black squares show the locations of cells making inhibitory connections. Excitatory and inhibitory connections to this cell have similar distributions as a function of orientation. Fig. 1.8B shows the theoretical average distribution of connections for retinotopically identical RFs, as a function of orientation difference (top) and spatial phase difference (bottom). The tightness of tuning as a function of correlation is determined by the parameter n_{pow} (see Methods). Large values of n_{pow} lead to tighter connectivity as a function of correlation, whereas smaller values of n_{pow} lead to broader connectivity. Increasing and decreasing n_{pow} had only minor effects on the behavior of the model.

For the full computational model, we reduced the synaptic strengths of the LGN input by a factor of 2, relative to the previous simulations without intracortical excitation. We thus relied on the positive feedback from same-phase intracortical excitation to amplify the response to suprathreshold stimuli. We also increased the default level of inhibition (i→e conductances) by a factor of 2, leaving the level of feedforward inhibition (LGN→i→e) roughly constant. The reduction of LGN input strength was important when comparing model behavior to experimental estimates of stimulus-induced conductance changes (discussed below). The intracortical excitation increased average firing rates to a 50% contrast stimulus of the preferred orientation by a factor of 2.1, relative to an identical circuit with intracortical excitation removed.

Intracortical amplification of an effective stimulus by feedback excitation has been used in many other models. [Douglas et al., 1989, 1995, Ben-Yishai et al., 1995, Somers et al., 1995, Suarez et al., 1995]

Our use of excitatory feedback differs in that our "amplifier" is localized in spatial phase as well as in orientation. It also differs from Ben-Yishai *et al.* and Somers *et al.* in two important respects. First, our intracortical inhibitory connections spread no further than excitatory connections – both are equally localized in orientation. Second, the resulting amplified responses have F1 tuning identical to that of the thalamocortical input alone, as observed by Ferster et al. [1996]) (shown in Fig. 1.10).

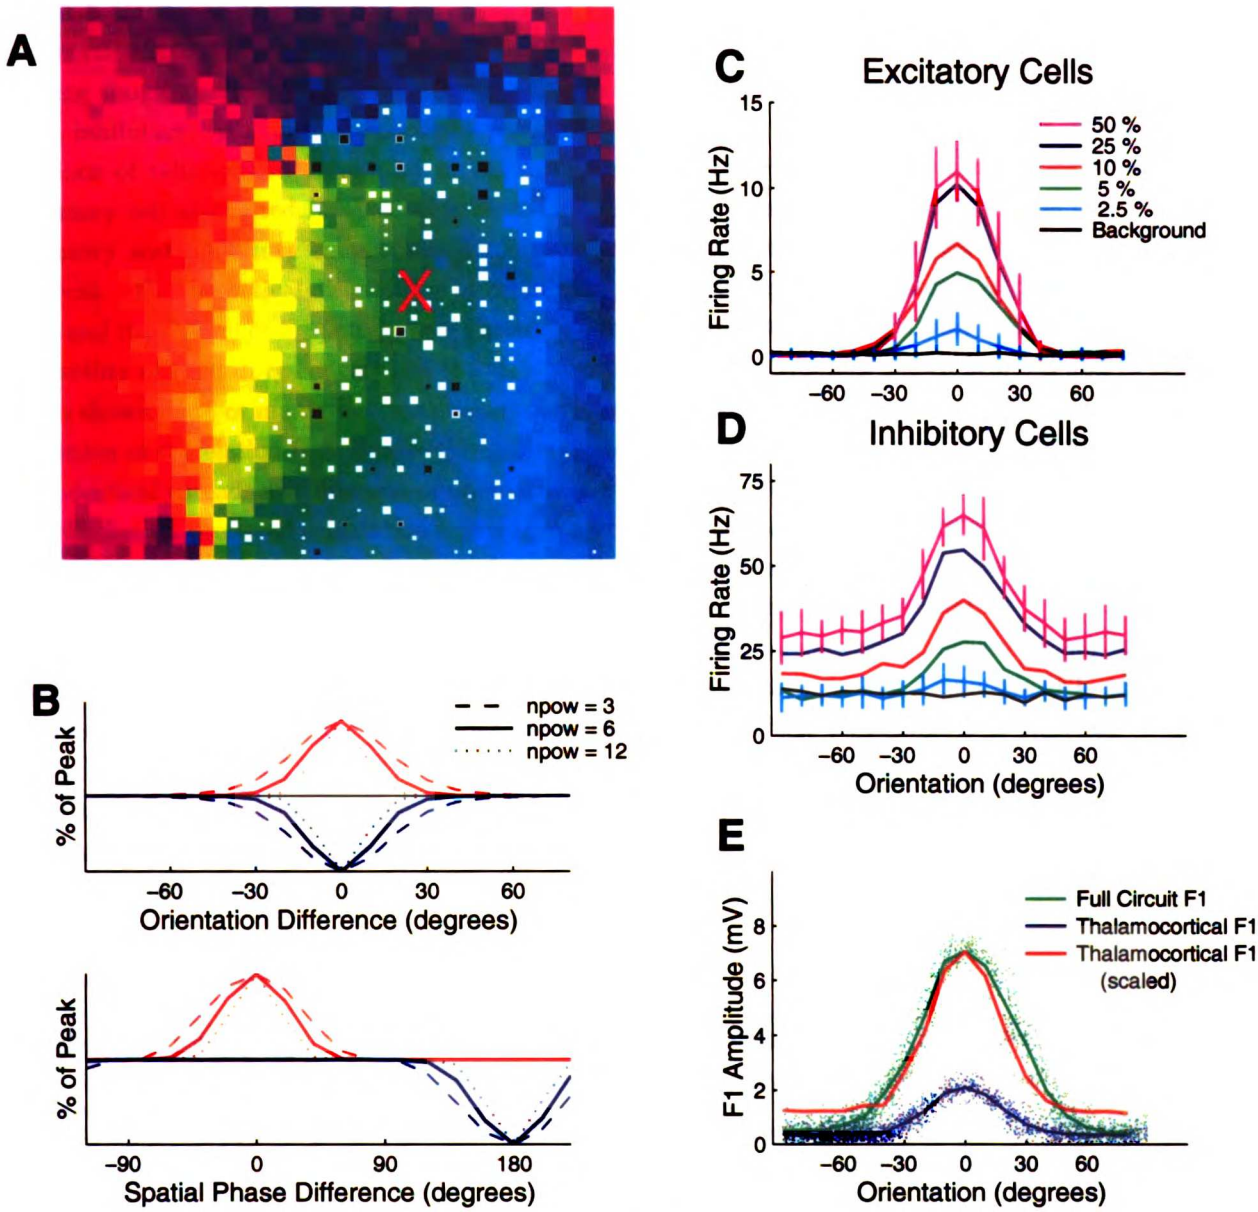

Figure 1.8:

Behavior of the full computational model. **A**. Orientation map used, and typical pattern of intracortical connections. There is one excitatory cell in every location of the cortical map (40 × 40 lattice) and one inhibitory cell in every fourth location (20 × 20 lattice). Cells were assigned preferred orientation according to illustrated 40×40 color map (red-green-blue-red representing $0^o - 60^o - 120^o - 180^o$) (map is 2/3 mm x 2/3 mm from measurement in cat V1, provided by Michael Crair and Michael Stryker); RF spatial phases were assigned randomly to each cell, while retinotopic centers progress continuously across the map (described in Methods).

Fig. 1.8 Caption, continued:
Intracortical connections were assigned probabilistically according to RF correlations (excitatory connections, yielding roughly same-phase excitation) or anti-correlations (inhibitory connections, yielding roughly anti-phase inhibition). A typical connectivity pattern is shown by the black and white squares, which illustrate locations of cells making inhibitory or excitatory intracortical synaptic connections, respectively, to the excitatory cell at the red X. Area of squares is proportional to connection strength. The distributions of excitatory and inhibitory connections across orientations are similar; on average, these distributions are identical. **B.** Theoretical distribution of connectivity as a function of the difference in preferred orientation (top) and the difference in spatial phase (bottom) between two cortical neurons with overlapping RF centers. Probability of excitatory connections shown in red; inhibitory probabilities shown inverted and in blue. All values shown as % of maximal connection probability. The parameter n_{pow} controls the width of tuning as a function of correlation (see Methods); $n_{\text{pow}} = 6$ (solid line) is the default value. Excitation and inhibition have identical spreads as a function of orientation difference, but have opposite preferences for spatial phase. Distribution vs. preferred orientation is averaged over cells of all spatial phases; distribution vs. spatial phase averaged over cells of all preferred orientations, with spatial phase measured with respect to the center of the RF for all orientations. **C-E.** All responses are to 3 Hz, 0.8 cycle/^{o} sinusoidal grating. **C,D.** Firing rates of of excitatory and inhibitory cells, vs. orientation, as function of contrast (indicated by key in C). Error bars for the 50% contrast and 2.5% contrast responses are ± 1 std. dev. **E.** Amplitude (F1) of excitatory cell voltage modulation, with and without the intracortical circuitry, vs. difference of stimulus orientation from preferred. Dots: F1 for all 1600 excitatory cells; traces: means in 10^o orientation bins, as in C,D. Blue: F1 for thalamocortical inputs alone; green: F1 with the full cortical circuit. Red trace is the thalamocortical response scaled to the peak response of the full cortical circuit; note that thalamocortical and full circuit have same tuning, as in Ferster et al. [1996].

1.3.8 Computational Model: Tuning

For the full circuit, typical currents and voltages from excitatory cells at various orientations are shown in Fig. 1.9. The behavior of a cell tuned to the stimulus is shown at the top left. At the anti-preferred temporal phase, large inhibitory currents hyperpolarize the cell to near the chloride reversal potential. As the preferred phase is approached, the inhibitory currents drop and the excitatory currents rise, depolarizing the membrane to spike threshold. Spikes in turn evoke adaptation currents, which help to shut off the cell's response and also, by acting on other cells across the network, lead to lowering of the intracortically evoked excitatory current. For a cell tuned to stimuli perpendicular to that presented (top right in Fig. 1.9), inhibition dominates at all phases and the cell is prevented from firing.

The full model, including same-phase excitation, achieves contrast invariant tuning to sinusoidal gratings (Fig. 1.8C, colored lines). Orientation tuning halfwidths at half height of the mean firing rates of excitatory cells were between $19^o - 21^o$ for contrasts ranging from 2.5% to 50%. The model was also well tuned to a moving oriented bar. Tuning width for the default parameters was 27.5, slightly broader than for gratings.

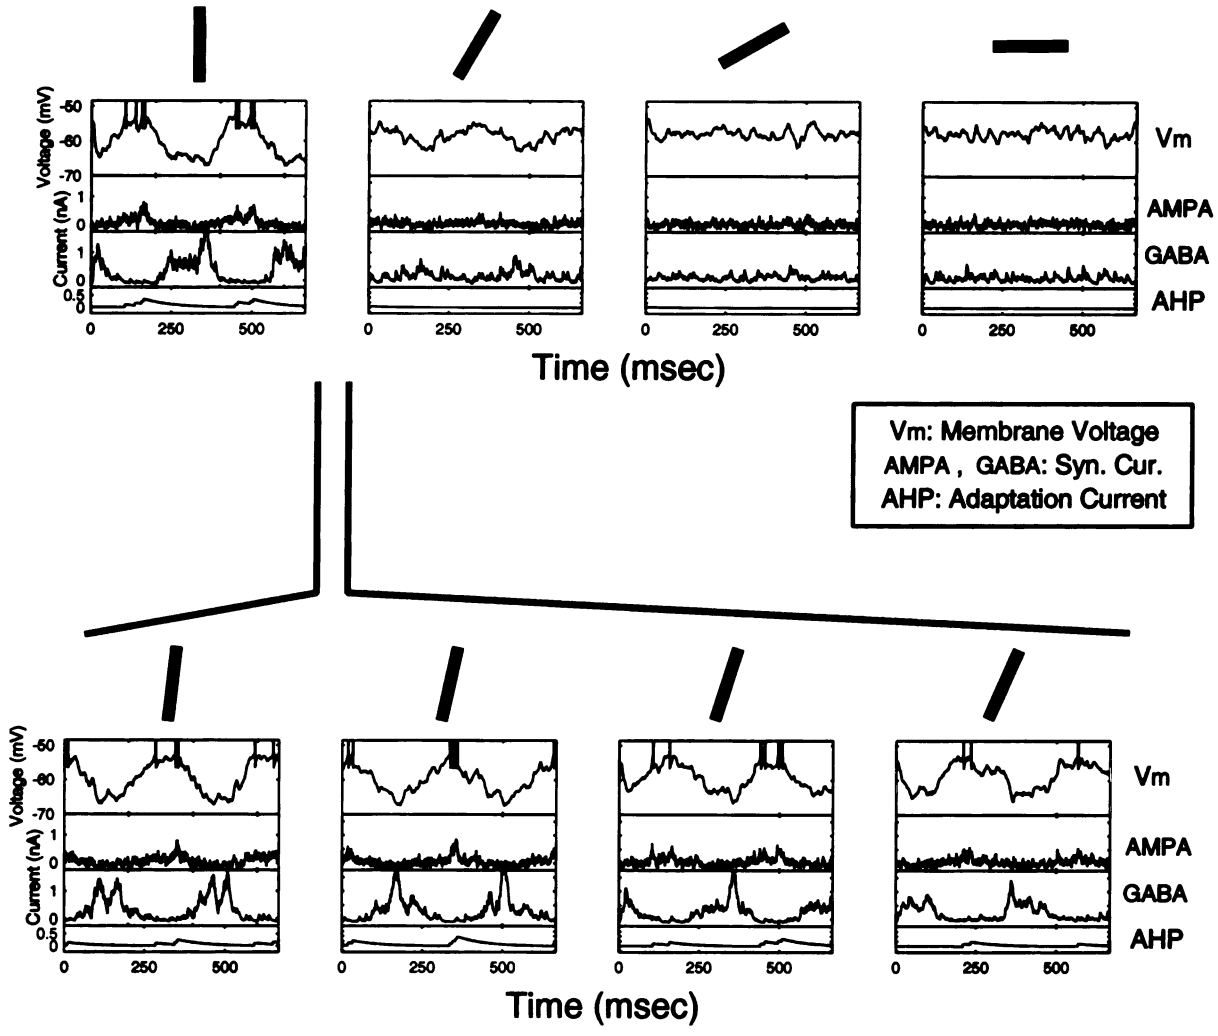

Figure 1.9:

Example traces. Voltage (V_m) and conductance in response to a 3 Hz sinusoidal grating at 50% contrast. AMPA,GABA-A: Synaptic conductances with background subtracted (converted to currents at threshold as in Fig. 1.7A). AHP: spike triggered potassium conductance. Bars show stimulus orientation relative to preferred (vertical). Top row shows orientation differences spaced at 30^o intervals; bottom row shows model behavior at orientations between 0^o and 30^o. Excitation and inhibition arrive out of phase and have similar orientation tuning. Inhibition dominates at the null.

Our model also achieves nearly identical orientation tuning of the excitation and inhibition received by a simple cell, as observed in Ferster [1986]. Fig. 1.10 demonstrates that the F1 tuning

curves of excitatory and inhibitory conductances onto cortical cells have nearly identical shapes. However, the spiking responses of excitatory and inhibitory cells differ. Like excitatory cells, inhibitory interneurons have a tuned peak in their firing rates near the preferred orientation; however, unlike the excitatory cells, at the null orientation their firing rates rise above background with increasing contrast (Fig. 1.8D). This untuned component of the inhibitory response counters the untuned component of the LGN input and is necessary to prevent excitatory spiking for null-oriented stimuli. Since the untuned response component increases with increasing contrast, inhibitory cell tuning measured as HWHH (with background subtracted) broadens with contrast from 32.3° at 5% contrast to 41.6° at 50% contrast. If, however, one measures the width of tuning after subtracting the response to null stimuli, inhibitory cells have HWHH across contrasts of 18.6° to 20.7°, similar to excitatory cell tuning.

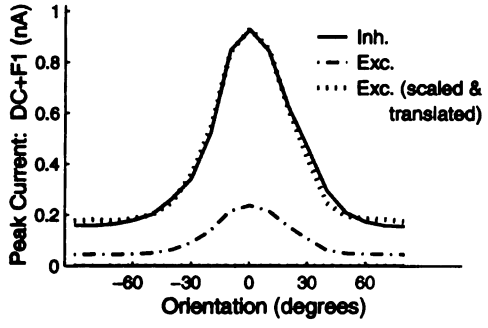

Figure 1.10:

Orientation tuning of peak excitatory current (dot-dashed line) and inhibitory current (solid line) at 50% contrast (peak current equals DC+F1 - see Fig. 1.7B). Dotted line shows excitatory current scaled and translated to match maxima and minima of inhibitory currents. Peak inhibition is larger than excitation at all orientations, but the tuned component of excitation and inhibition have nearly identical shape.

1.3.9 Orientation Tuning is Driven by LGN Input

The tuning in our model is driven by the LGN inputs, despite the strong role of intracortical inhibition. In agreement with the cortical cooling experiments of Ferster et al. [1996], the tuning of the first harmonic (F1) of the membrane potential response to moving gratings was identical for the full circuit and when all intracortical synapses were set to zero strength (Fig. 1.8E). For our default parameters, we find that the cortical circuitry amplifies the purely thalamocortical F1 by a factor of 3.4, slightly higher than the estimate (2.7) reported by Ferster et al. [1996].

A more rigorous test of the hypothesis that cortical tuning is driven by the LGN can be obtained by systematically varying the tuning of the LGN input. This can be accomplished by varying the spatial frequency of a sinusoidal grating (Fig. 1.11). As spatial frequency is increased, the orientation tuning of the F1 of the LGN inputs becomes narrower (Fig. 1.11A, top). The F1 of the full circuit closely follows the LGN F1 across spatial frequencies (Fig. 1.11A, middle), yielding a narrowing of orientation tuning of the spiking response with increasing spatial frequency (Fig. 1.11B). Note that there is an optimal spatial frequency, at which spike response to stimuli of the preferred

orientation is maximal (Fig. 1.11B, inset); thus, as spatial frequency increases, the orientation tuning width monotonically narrows, while spike response to the preferred orientation rises to a maximum and falls again. This is precisely the behavior observed in cat cortical cells [Vidyasagar and Sigüenza, 1985, Webster and De Valois, 1985, Jones et al., 1987, Hammond and Pomfrett, 1990] and is the behavior expected for a cell receiving inputs from a linear Gabor function RF. Similar behavior should occur for other stimulus manipulations that would vary the tuning of the LGN F1 but that have not yet been reported experimentally, such as use of moving or flashed ellipses of various eccentricities; the key point is that the full circuit F1 and the spiking tuning closely covary with the LGN tuning.

Inhibitory sharpening of orientation tuning is due exclusively to disynaptic inhibition driven by the LGN input, *i.e.* sharpening depends on feedforward inhibition. Feedback inhibition (resulting from cortical cell excitation of inhibitory cells – gray dashed lines in Fig. 1.6) has little effect on model behavior. In simulations run after setting these connections to zero, tuning curve halfwidths at the default parameters remain unchanged to within 1^o for 10% contrast and above (4^o change at 5% contrast), and peak spiking responses increased by less than 6.4% (see Fig. 1.13B). These small changes in output occur even though the mean inhibitory conductance in response to a 50% contrast grating at the preferred orientation is reduced by 38%.

This result can be understood by examining Fig. 1.6. With same-phase excitation and anti-phase inhibition, firing of excitatory cells of one phase increases the inhibition onto excitatory cells of the opposite phase. But excitatory cells with RFs of opposite *spatial* phase spike during opposite *temporal* phase of the stimulus. Therefore, feedback inhibition is directed onto cells that are already not spiking, and has little effect on model behavior.

1.3.10 Inhibitory Dominance and Inhibitory Blockade

The model operates in an inhibition dominated regime, as revealed by a variety of measures. For the "default" parameter settings used to measure tuning (Fig. 1.8), the total synaptic strength received by a cell (see Methods for definition) was 5 $nA\ msec$, 4.25 $nA\ msec$, and 7.5 $nA\ msec$ for the LGN inputs, cortical excitation and cortical inhibition respectively. A null oriented stimulus induces an inhibitory current that is approximately 3.9 times as large as the excitatory current for a cell clamped at threshold voltage, corresponding to an inhibitory conductance change that is 11.7 times greater than the change in excitatory conductance. Strong inhibition is consistent with the experimental result that non-specific stimulation of LGN afferents (or cortical white matter) yields a short EPSP followed by a large IPSP (e.g. Ferster and Jagadeesh, 1992). However, quantitative measurements of the balance of excitation and inhibition under physiological conditions have not been reported, and estimates based on indirect evidence are imprecise [Shadlen and Newsome, 1994].

Our model agrees with the experiments of Nelson et al. [1994], which demonstrated that intra-

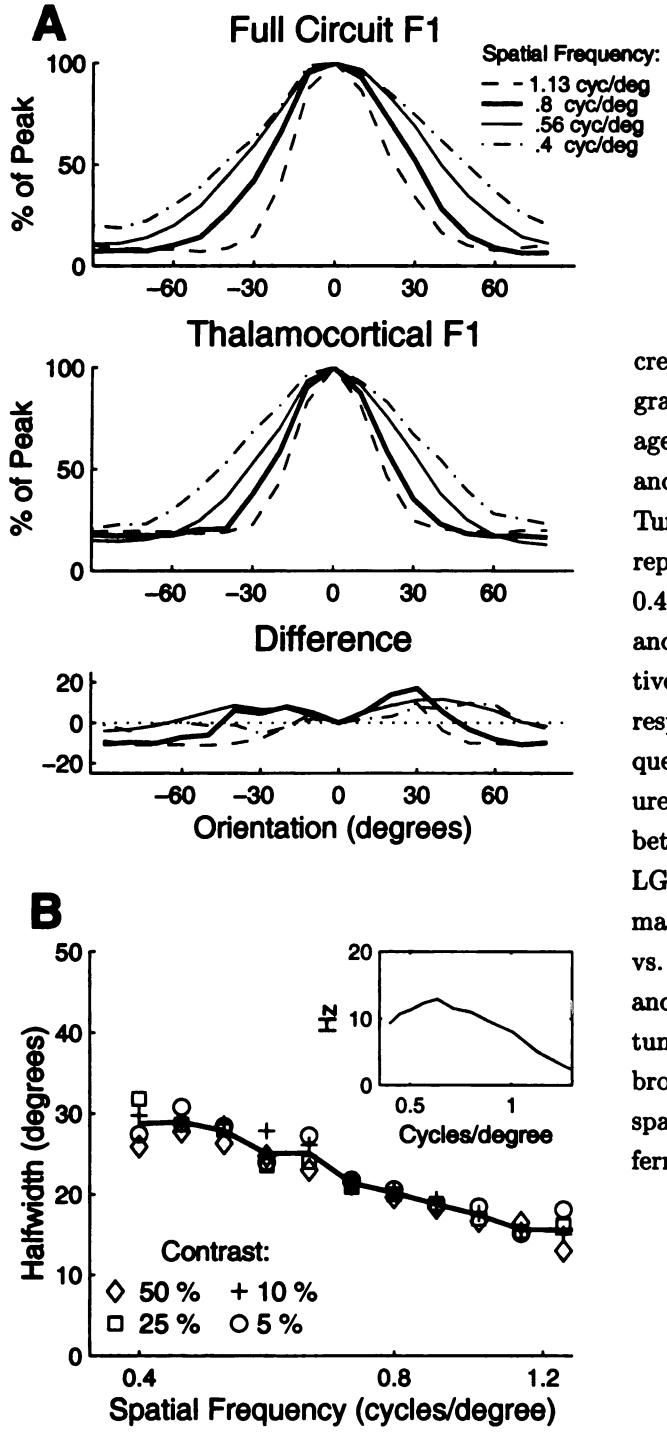

Figure 1.11:

Orientation tuning narrows with increasing spatial frequency of a sinusoidal grating (3 Hz). **A.** Tuning of the F1 voltage modulation for the full circuit (top) and with LGN excitation only (middle). Tuning curves, from widest to narrowest, represent response to spatial frequencies 0.4 (dot-dashed), 0.56 (thin), 0.8 (thick), and 1.13 (dashed) cycles/degree respectively; each curve is normalized to its peak response. Thicker line is the spatial frequency used for simulations in other figures. The bottom shows the difference between the normalized tuning curves – LGN input F1 and full circuit F1 closely match. **B.** Halfwidth at half height vs. spatial frequency for 5%, 10%, 25% and 50% contrast gratings. Orientation tuning remains contrast invariant over a broad range of spatial frequencies. Inset: spatial frequency tuning curve at the preferred orientation and 50% contrast.

cellular blockade of inhibition in a single neuron, with DC current injection sufficient to suppress excess background firing, did not disrupt cortical orientation tuning. Figure 1.12 shows extracel-

lular tuning curves for cells under inhibitory blockade with (solid line) and without (dot-dashed line) compensating DC current injection. Tuning curves without blockade are shown for comparison (dashed line). (Note that Nelson et al. [1994] used moving bar stimuli, while we use drifting gratings.) We find a small amount (1.9 Hz) of elevated spiking at the null, consistent with the report that some cells did show spiking at the null under inhibitory blockade (S.B. Nelson, private communication). This reflects the untuned mean DC component of the LGN input.

Effects of global blockade of inhibition are discussed below (Fig. 1.13A).

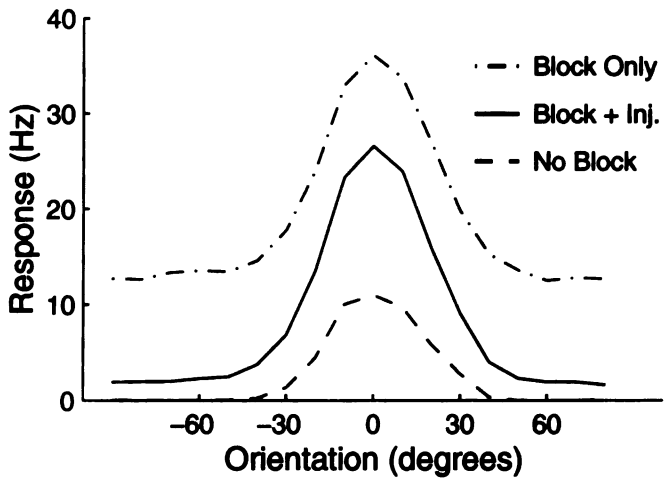

Figure 1.12:

Orientation tuning after blocking inhibition in a single cell, as in Nelson et al. [1994]. Tuning curves derived when inhibitory and adaptation currents are blocked within a single cell (see Methods). Dotted line shows tuning with inhibitory blockade only; solid line shows the tuning when negative current equal to the mean inhibitory synaptic current at background is injected into the cell. Dashed line shows tuning without blockade for comparison. Inhibitory blockade with current injection has little effect on tuning. Notice, however, a slight (1.9 Hz) rise in the response to null stimuli. Contrast equals 50%.

1.3.11 Robustness of Model Results

The effect of parameter variations on computational model behavior is shown in Fig. 1.13. Each of the subplots represents the effect of varying the strength of two variables out of the following four: the three types of synaptic connections in the model (thalamocortical excitation (LGN), intracortical excitation (e→{e,i}), and inhibition (i→e)), and spike-rate adaptation in excitatory cells. Magnitudes of a given type of connection or conductance were varied by multiplying all the corresponding unitary conductances by a constant factor. In Figures 1.13B and 1.13D, the intracortical excitation of inhibitory cells was removed (e→i = 0).

Simulations were run at four contrasts (5%, 10%, 25%, 50%) for each point in the 5×5 parameter grid. We monitored contrast-invariance of orientation tuning, average tuning width (HWHH), and

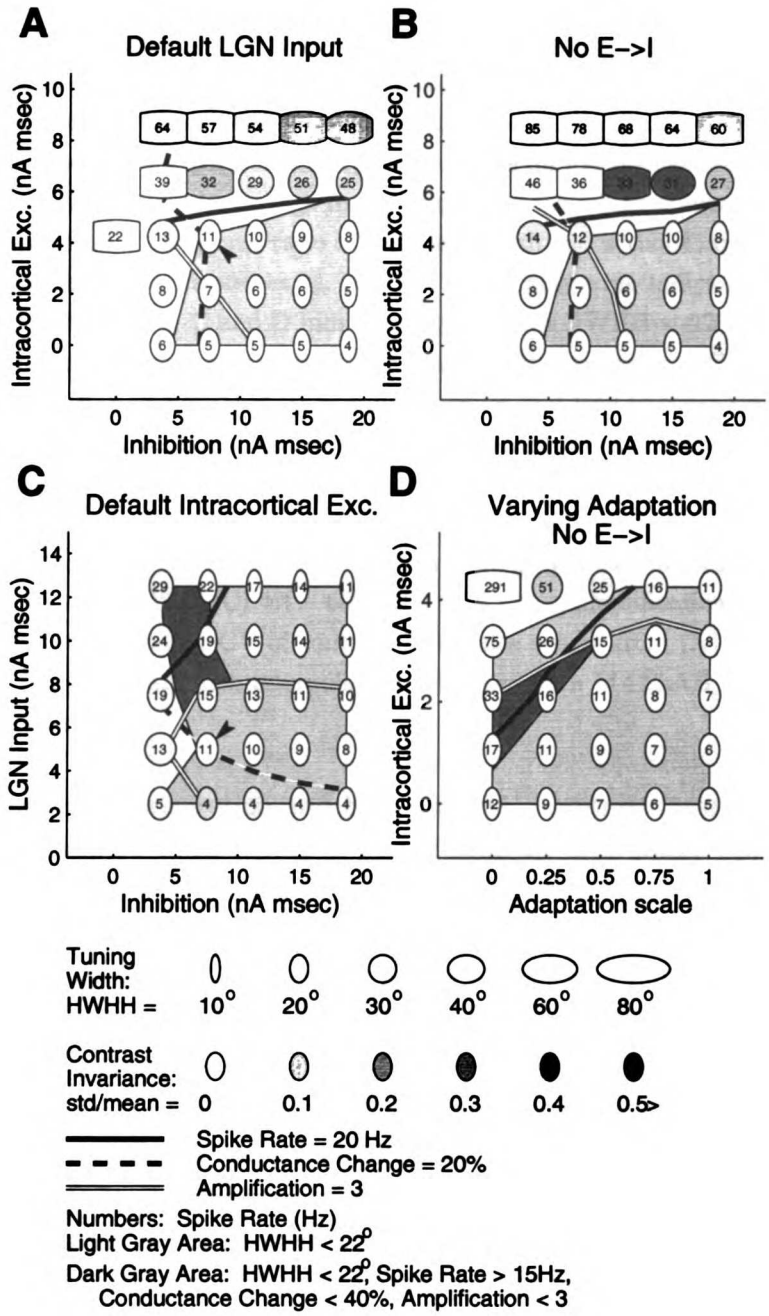

Figure 1.13:

Robustness of model behavior to changes in various model parameters. Each subplot represents the effect of varying the strength of two variables out of the following four: the three types of synaptic connections (LGN input, intracortical excitation, and inhibition), and spike-rate adaptation in excitatory cells. Mean spike tuning half-width at half-height (HWHH) for 5%, 10%, 25%, and 50% contrast gratings is represented by oval width. Oval height represents 30°.

Fig. 1.13 Caption, continued:
Mean spike rates (Hz) for preferred stimuli at 50% contrast are printed inside each oval.Darker ovals indicate a loss of contrast invariance, monitored as standard deviation divided by the mean of the HWHH over the four contrasts sampled.Points with extremely broad tuning are contrast invariant only because all contrasts give maximal HWHH. Lines show experimentally reported values for mean spike rate (bold line equals 20 Hz; Albrecht [1995]), maximal conductance change from background for a high contrast, null stimulus (dashed line equals 20%; Douglas et al. [1988]), and ratio of voltage F1 with and without input from cortical circuitry ("amplification"; white line = 3; Ferster et al. [1996]). Light gray area indicates areas of sharp tuning (HWHH $< 22^o$). Dark gray area in C and D indicates regions with HWHH $< 22^o$, amplification < 3, spike rate > 15 Hz, and conductance change $< 40\%$ (C) or $< 22\%$ (D). Arrows in A,C indicate default network parameters used in Figs. 1.8- 1.12.Note that (i) parameter values that lead to sharp tuning also yield contrast invariance; (ii) higher levels of inhibition sharpen tuning; (iii) high levels of excitation lead to instability – runaway feedback excitation – indicated by high spike rates, broad tuning and loss of contrast invariance; (iv) removing e→i connections causes little change in stable region except that amplification is reduced (also true for varying adaptation as in D, not shown). Within light gray areas, null conductance changes range from: (A) 20% –44%; (B) 21% –42%; (C) 4% – 60%; (D) 21% –23%; amplification ranges from (A) 2.7–3.8; (B) 2.7–3.5; (C) 2.4–3.9; (D) 2.2–3.2; CV of contrast invariance ranges from (A) .02–.06; (B) .03–.11; (C) .01–.15; (D) .01–.06. Gray area and line interpolations obtained with MATLAB "contour" command. See Results for detailed discussion.

firing rates at the preferred orientation and 50% contrast. Contrast-invariance is represented by color within the ovals; darker ovals represent loss of contrast invariance. Spike tuning HWHH is denoted by oval width, relative to oval height which represents 30^o. The light gray shaded regions indicate HWHH tuning of less than 22^o. Spike rate to a 50% contrast preferred stimulus is printed inside each oval. The lines show the relation of model behavior to various experimental estimates and are discussed below.

The model robustly achieves sharp, contrast-invariant tuning. Thin white ovals within the gray shaded regions indicate that, in the large regions of parameter space in which tuning is sharp, it also remains contrast invariant. Increasing inhibition leads to sharper tuning (narrower ovals) while having only moderate effects on spike rate, except with very strong levels of intracortical excitation.

This qualitative behavior is quite robust to changes in other model parameters. Increasing and decreasing the tightness of the correlation-based connectivity by changing the parameter n_{pow} (see Fig. 1.8B) by a factor of 2 had only moderate effects on model behavior. Using the more broadly tuned geniculocortical connectivity matched to the data of Ferster et al. [1996], tuning widened substantially. However, large levels of inhibition still produced sharp tuning (see also Fig. 1.5), and the level of intracortical excitation that led to feedback instability remained relatively unchanged.

The Roles of Feedforward and Feedback Inhibition

High levels of intracortical excitation (upper portions of Fig. 1.13A,B,D) leads to unstable feedback excitation, indicated by sharply increased spike rates, broadening of orientation tuning

and a loss of contrast invariance. Strong excitation is difficult to control, even for high levels of inhibition, because the model lacks the center-surround ("Mexican hat") intracortical connectivity commonly used to stabilize feedback excitation [Ben-Yishai et al., 1995, Somers et al., 1995]. A set of activated excitatory cells will excite cells of nearby orientation and similar spatial phase, while driving feedback inhibition (e→i→e) onto cells of nearby orientation but nearly opposite spatial phase. Excitation can thus spread in orientation along a series of cells linked by similar spatial phase, unchecked by feedback inhibition.

The most important parameter determining the presence of runaway excitation is the strength of intracortical e→e connections. However, feedforward inhibition from the LGN (LGN→i→e) also plays a role. As discussed earlier, the phase-nonspecific component of feedforward input is inhibitory and acts like a contrast-dependent effective threshold, increasing the depolarizing current necessary to reach threshold with increasing contrast (Fig. 1.7); intracortical excitation must overcome this inhibition in order to spread. At low contrast, the effective threshold is small and presents a weak barrier to the spread of excitation. Consequently, whenever feedback excitation resulted in a loss of contrast invariance (darker ovals), this loss was due to a widening of tuning at lower contrasts. That is, feedback instability poses the greatest problem at lower contrasts, even though peak firing rates are lowest at these contrasts.

Removing feedback inhibition by removing intracortical excitatory input onto inhibitory cells has little effect on the model behavior. This is demonstrated in Fig. 1.13B, which shows the results of using parameters identical to those of Fig. 1.13A except that e→i connections are set to 0 strength. Even though mean inhibitory conductance is substantially reduced (by 38% at the preferred orientation, 50% contrast, for default parameters), model behavior is virtually unchanged.

In agreement with experiments using large amounts of GABA-A antagonists [Sillito, 1975, Tsumoto et al., 1979], strong reduction of inhibition leads to a loss of orientation tuning (a single run with zero inhibition is shown at left of parameter grid in Fig. 1.13A). This loss of tuning is due primarily to the fact that, at very low inhibition levels, the untuned phase-nonspecific component of the LGN input remains uncancelled by feedforward inhibition (though unchecked feedback excitation also contributes). Under less severe blockade, our model predicts that, much as in intracellular GABAergeic blockade (Fig. 1.12, "block only" condition), excitatory cell tuning should resemble that of inhibitory cells: a well-tuned peak on an untuned platform whose height increases with contrast (see left portion of Fig. 1.5: orientation tuning gradually widens before being lost as inhibition decreases). Such an untuned platform at a single stimulus contrast has been observed after inhibitory blockade in cat simple cells [Sillito, 1975, Tsumoto et al., 1979]; similar results can be seen in monkey visual cortex [Sato et al., 1996].

Comparisons to Experimental Estimates

In addition to the width and invariance of orientation tuning, we monitored the conductance change to a null stimulus, the magnitude of the full circuit voltage F1, and the mean spike rate for

a 50% contrast preferred stimulus. The lines in Fig. 1.13 shows the contours through parameter space corresponding to experimental estimates of these variables. First, some results [Douglas et al., 1988, Koch et al., 1990, Douglas et al., 1991] suggest conductance change at the null is small, although there is a great deal of uncertainty about this issue (see Discussion). The dashed lines represent conductance changes that are 20% of background. Smaller changes were achieved only for sufficiently small levels of excitation and inhibition (regions to left of dashed lines). Strong levels of inhibition can sharpen tuning, but at the expense of larger conductance change at the null. Second, Ferster et al. [1996] very roughly estimated a ratio of 2.7 for the magnitude of the voltage F1 with vs. without input from cortical circuitry (we refer to this ratio as the level of "cortical amplification"). A cortical amplification level of 3 is represented by the white lines. Amplifications of this size or smaller are obtained for cortical synaptic strengths that are sufficiently small compared to thalamic input (lower left corner in Fig. 1.13A and 1.13B; upper portion of Fig. 1.13C). Third, Albrecht [1995] reported that the F1 of the spike rate of simple cells to preferred stimuli at 50% contrast is $\approx$ 38 Hz. Assuming a ratio of F1 to DC of 1.57 [Skottun et al., 1991], this results in a mean spike rate of $\approx$ 20 Hz (bold line). For default levels of LGN input (Fig. 1.13A and Fig. 1.13B), spike rates for parameters that achieve contrast-invariant tuning were relatively low.

Simply by varying synaptic strengths, we were not able to quantitatively match all three of these estimates simultaneously. The reason is as follows. Low conductance change at the null means that feedforward conductances must be small relative to the cell's background conductance. Given the responsiveness of our model neurons to these weak inputs, we find that this implies that firing rates will be small unless there is large cortical amplification of the feedforward input. However, such large cortical amplification both disagrees with the estimates of Ferster et al. [1996]; (see also Chung and Ferster [1997]) and, in our model, leads to instability and loss of tuning. By increasing thalamocortical conductances, we can achieve sharp, contrast-invariant tuning with realistic firing rates and amplification levels, but conductance changes to a null stimulus become larger. This is shown by the dark gray region in Fig. 1.13C, which shows points with spike rate > 15 Hz, tuning HWHH $< 22^o$, amplification ratio < 3, and conductance change $< 40\%$.

Varying Neuronal Gain

One obvious way out of this dilemma is that cortical cells may respond more strongly to weak thalamocortical input and moderate amplification than do our model neurons. We used extremely simple model neurons, including a simple, voltage-independent mechanism for spike rate adaptation. These were matched to experimental data on firing rate vs. constant somatic current level *in vitro* (McCormick et al. [1985]; see Methods and Troyer and Miller [1997b,a]). Our model may not accurately model responses to synaptic currents *in vivo*, where cells are subject to various neuromodulators, slow currents, and active conductances, receive fluctuating synaptic input over a spatially extended dendritic tree, and may show weaker spike rate adaptation. All of these may affect the "neuronal gain", *i.e.* the magnitude of a neuron's response to a given level of synaptic

input (*e.g.* Fox et al. [1990], Mel [1993], Nowak et al. [1997], Tang et al. [1997], Storm [1990]). A simple argument shows that neuronal gain and synaptic strength can trade off against one another: suppose a neuron's gain could be doubled, so that it fires at twice its previous rate in response to any stimulus pattern. If the strength of its synaptic conductances onto other cells were simultaneously cut in half, then each cell in the network would receive the same total input as before the change. By making these changes for all excitatory cells in the network, spike rates could be doubled without altering network behavior, including orientation tuning widths.

To investigate the changes in model behavior induced by increasing neuronal gain, we studied the effects of lowering spike-rate adaptation in model excitatory neurons (Fig. 1.13D). We also removed e→i conductances as in Fig. 1.13B; this reduces cortical amplification without altering network behavior (compare 1.13A and 1.13B). Spike rates are increased while tuning remains sharp and contrast-invariant (1.13D). Conductance changes at the null are relatively unaffected by the change in neuronal gain (changes range from 20.6% to 22.5%, excluding the unstable point at upper left). Reducing the spike-triggered AHP current in our model by $\approx$ 75% leads to a substantial region of parameter space satisfying all experimental estimates: the dark gray region shows points for which spike rates are $>$ 15 Hz, amplification ratios $<$ 3, and tuning HWHH $<$ 22°; in this region, conductance changes are $<$ 22%.

1.4 Discussion

1.4.1 Principal Findings and Predictions

We have constructed a simple model that accounts for contrast-invariant orientation tuning in layer 4 of cat visual cortex. First, we analyzed the LGN input to simple cells in response to drifting sinusoidal gratings. This input has two components: a temporally modulated, phase-specific component which is tuned for orientation, and an unmodulated, phase-nonspecific component which, assuming LGN cells are not orientation tuned, must be completely untuned. Because of LGN rectification, this untuned input increases significantly with increasing contrast, implying that no simple threshold mechanism could correct for it across contrasts. This problem is general: the phase-nonspecific LGN input component for any stimulus is poorly tuned and increases with contrast. To counteract this increase in untuned feedforward excitation, stimulation of an excitatory cell at its null orientation must evoke increasing intracortical inhibition and/or increasing withdrawal of intracortical excitation with increasing contrast. Thus, the central questions for understanding contrast-invariant orientation tuning in cat layer 4 are how the level of this contrast-dependent "pull" is computed, and which cells are its source.

Second, we have demonstrated that cortical circuitry relying on correlation-based connectivity is sufficient to robustly yield contrast-invariant orientation tuning. In our model, the contrast-dependent inhibition to layer 4 simple cells comes from inhibitory cells with similar orientation

tuning and roughly opposite spatial phase (Fig. 1.8B). The tuned input component is amplified by excitatory connections between cells with similar orientation tuning and similar spatial phase. While some details might vary, our key predictions are that the layer 4 circuitry is strongly phase-specific, inhibition-dominated, and localized in orientation, and has orientation tuning that depends on the LGN input tuning.

Beyond the circuit itself, our most striking prediction is the existence of inhibitory simple cells that are tuned for orientation, but have a contrast-dependent response to null-oriented stimuli. We also predict that (i) the DC LGN input should increase with increasing contrast at all orientations, while (ii) the dominance of inhibition should cause the net DC feedforward input (LGN plus feed-forward inhibition) to decrease with increasing contrast at all orientations. These predictions are most easily tested with null-oriented stimuli, for which feedback excitation should be negligible. Prediction (i) could be tested by blocking inhibition intracellularly in a single layer 4 cell. Note that, for (ii), the predicted decrease is in net current at threshold voltage; this need not imply voltage decrease from rest due to differences between excitatory and inhibitory reversal potentials. In addition, we predict that intracellular, or sufficiently weak extracellular, GABAergic blockade in layer 4 should reveal a tuned response component sitting on an untuned plateau response (*e.g.* Fig. 1.12).

1.4.2 Possible Sources of Contrast-Dependent Inhibition

Besides anti-phase inhibition, what are other possible sources of the contrast-dependent "pull" required to achieve contrast-invariant tuning? Since the LGN input is exclusively excitatory [Ferster and Lindström, 1983], inhibition must come from other cortical cells. One obvious possibility is inhibitory cells preferring stimuli of the orthogonal (null) orientation. A major contribution from such cells seems ruled out in cat layer 4 by the similar orientation tuning of excitation and inhibition. [Ferster, 1986]. and by the phase-specificity and spatial opponency of EPSPs and IPSPs [Ferster, 1988, Hirsch et al., 1995a].

Another source of "pull" could be withdrawal of excitation from other cortical cells. Since layer 4 simple cells have very low spontaneous rates, significant withdrawal would most likely come from complex cells, which do show contrast-dependent inhibition to a null stimulus [Sclar and Freeman, 1982]. This explanation would simply move the problem of the origin of the required inhibition onto complex cells.

One study has suggested that inhibition onto layer 4 simple cells with multiple ON/OFF subregions may come primarily from layer 4 simple cells with single-subregion RFs [Toyama et al., 1981], while multi-subregion inhibitory cells in layer 4 may primarily act on layer 3 complex cells. Thus, the population showing an untuned plateau response might be restricted to the single-subregion inhibitory cells; inhibition from these cells might block responses at the null orientation in both inhibitory and excitatory multi-subregion cells.

Carandini and Ferster [1997] have shown that a tonic hyperpolarization underlies contrast adaptation in cat simple cells. This contrast-dependent "pull" may make an important contribution toward contrast-invariant orientation tuning for steady-state stimuli. However, this mechanism would not play a role for transient stimuli such as moving or flashed bars. While contrast-invariance for these stimuli has not been explored in detail, our analysis shows that tight orientation tuning to high-contrast bars combined with robust responses to low contrast bars would require contrast-dependent cancellation of the phase-nonspecific component of the LGN input. Furthermore, contrast-dependent "pull" is required at all orientations, while contrast adaptation appears to not be induced by null-oriented stimuli [Allison and Martin, 1997].

Our analysis shows that the untuned component of the LGN input will cause simple cells to spike in response to a high-contrast, null-oriented stimulus, unless they are inhibited. Our model relies on the simplest explanation for the source of this contrast-dependent inhibition: inhibitory neurons (or a subset of them) are not strongly inhibited and *do* spike, providing the inhibition necessary to prevent responses in the remaining layer 4 neurons.

1.4.3 Experimental Results Related to Inhibitory Cell Tuning

Several studies have reported that cat layer 4 inhibitory interneurons have simple cell RF structure much like that of excitatory neurons, and are orientation tuned. [Gilbert and Wiesel, 1979, Toyama et al., 1981, Martin, 1988, Azouz et al., 1997]

However, only Azouz et al. [1997] reported details of orientation tuning in V1 inhibitory neurons. Consistent with our predictions [Krukowski et al., 1996], they found that 5 of 8 inhibitory cells identified in intracellular recording showed significant spiking response at the null; the strongest null response was found in the only layer 4 interneuron studied. However, they did not study contrast dependence of responses, and also reported that in their hands 70% of extracellularly recorded cells showed spiking at the null.

Our prediction for inhibitory cell tuning is also consistent with observations in other systems. In rabbit visual cortex, putative inhibitory neurons respond to all orientations, while other neurons respond only to a limited range of orientations [Swadlow, 1988]; orientation tuning of these cells was not otherwise reported. More generally, putative inhibitory neurons are more broadly tuned than other cells across a variety of cortical systems [Brumberg et al., 1996, Simons and Carvell, 1989, Swadlow, 1989, 1990, 1991, 1994].

1.4.4 Input-Driven Tuning

Orientation tuning in our model is driven by the LGN inputs. The model yields voltage responses that have identical F1 tuning with and without intracortical synaptic input (Figs. 1.8, 1.10; Ferster et al. [1996]). In addition, orientation tuning sharpens with increasing spatial frequency of a

sinusoidal grating, as is also observed in cat visual cortex [Vidyasagar and Sigüenza, 1985, Webster and De Valois, 1985, Jones et al., 1987, Hammond and Pomfrett, 1990]

Nonetheless, cortical circuitry plays a key role in our model. Inhibitory cells have an untuned component to their response that eliminates the untuned component of the LGN input and sharpens spike tuning. Studies in the rat whisker barrel system also indicate that the layer 4 computation is local, input-driven, and dependent on broad inhibitory tuning [Brumberg et al., 1996, 1997, Pinto et al., 1996, Simons and Carvell, 1989], suggesting that these may be general properties of layer 4 cortical circuitry.

Other models [Ben-Yishai et al., 1995, Somers et al., 1995, Adorján et al., 1997] predict that orientation tuning is determined by cortical circuitry, independent of input tuning, and thus should not change with stimulus spatial frequency. Such models do not allow simultaneous representation of multiple orientations at a single retinotopic position [Carandini and Ringach, 1997], which may play an important role in visual processing, including figure/ground separation and object recognition [Zucker, 1986].

1.4.5 Model Robustness and Experimental Constraints

Our model uses correlation-based inhibition to robustly achieve sharp, contrast-invariant, input-driven orientation tuning over a wide parameter range. The ability of a simple, two-cell conceptual model to capture the essential behavior of the more realistic computational model argues strongly for the robustness of the underlying mechanisms.

We examined the conditions under which this tuning could be achieved while satisfying several other constraints suggested by a variety of experiments: (i) small ($< 20\%$) conductance changes to a null stimulus; (ii) high ($\approx$ 20Hz) steady-state firing rates to a high-contrast preferred stimulus; (iii) only moderate ($\approx$2.7 times) cortical amplification of the LGN input. We have found that these conditions can be simultaneously satisfied if the neuronal gain of our model excitatory neurons is increased, for example by reducing the magnitude of the spike-trigged adaptation (AHP) current. Alternately, by increasing feedforward synaptic strengths, all but the conductance condition can be robustly met.

The experimental support for these constraints in layer 4 is not entirely clear. With the exception of the amplification level, our numerical targets are based on average values across cortical simple cells from all layers, but properties of layer 4 cells may differ (*e.g.* see Tolhurst and Thompson [1981], which suggests orientation tuning is broadest in layer 4). Indeed, if the layer 4 circuitry solves the problem of contrast-invariant tuning, it may be relatively easy for cells in other layers to display sharp tuning with high rates and small conductance changes. In addition, quantitative estimates of conductance change remain controversial. Different labs and techniques are producing widely varying measurements [Douglas et al., 1988, 1991, Hirsch et al., 1995b, Carandini and Ferster, 1997, Monier et al., 1997]. Furthermore, the relationship between net synaptic input and

conductance changes measured at the soma is complicated by factors such as voltage-dependent dendritic conductances (*e.g.* Yuste and Tank [1996]) and large synaptic background conductances [Bernander et al., 1991], which were not considered in previous theoretical studies [Koch et al., 1990].

Withdrawal of tonic input could also have an important effect on stimulus-induced conductance changes. Withdrawal of excitation from complex cells, discussed above, would contribute to low conductance change in response to a null stimulus. Withdrawal of tonic anti-phase inhibition might allow a large spike response to a preferred stimulus without a large change in conductance. There are some indications from experiments that this change may be small during spike responses [Carandini and Ferster, 1997, Monier et al., 1997]. Our model shows a large conductance change at the preferred orientation (mean increase 72%; peak 148%). Withdrawal of tonic inhibition would be enhanced in our model by raising the tonic background spike rate of our model inhibitory cells (currently $\approx$ 12 Hz, see Methods) to match values typically found experimentally ($\approx$ 20 Hz; Simons and Carvell [1989], Swadlow [1989, 1990, 1991, 1994], Brumberg et al. [1996]. Anti-phase i$\rightarrow$i connections, which have been observed in inhibitory simple cells in layer 4 (J. Hirsch, unpublished observations), could also contribute to withdrawal of inhibition at the preferred orientation, but at the cost of reducing inhibition at the null. This reduction in inhibition at the null could be compensated for by increasing inhibitory gain and/or i$\rightarrow$e connection strengths.

Many factors not studied here might increase the stability of feedback excitation, enabling high firing rates with small increases in null conductance (at the cost of large amplification). These factors include frequency-dependent short-term synaptic depression [Abbott et al., 1997, Priebe et al., 1997, Thomson and Deuchars, 1994, Tsodyks and Markram, 1997] and slow GABA-B mediated inhibition [Allison et al., 1996, Bush and Priebe, 1998]. In addition, weaker components of inhibitory connectivity that are phase-nonspecific or broadly tuned relative to excitation (*e.g.* Wörgötter and Koch [1991]) could help to stabilize the amplifier without dominating orientation tuning.

1.4.6 Nonlinear Response Properties With Local Circuitry

Many groups have suggested that (i) important features of simple cell responses can be approximated using a linear function of stimulus contrast [Movshon et al., 1978b, Glezer et al., 1982, Tolhurst and Dean, 1990, Albrecht and Geisler, 1991, Heeger, 1992, Carandini and Heeger, 1994, Carandini et al., 1998, 1997]; and (ii) such linear responses might be achieved by a circuit involving opponent or "push-pull" inhibition, in which inhibition *balances* excitation [Glezer et al., 1982, Tolhurst and Dean, 1990, Carandini and Heeger, 1994, Carandini et al., 1998, 1997]. Several groups further proposed that a variety of nonlinear contrast effects could be explained through addition to the linear model of a global, divisive normalization of activity levels [Albrecht and Geisler, 1991, Heeger, 1992, Carandini and Heeger, 1994, Carandini et al., 1998, 1997]; Carandini

and colleagues proposed that this normalization could arise through global, orientation-nonspecific inhibitory connectivity.

Our model uses opponent inhibition, but assumes that inhibition dominates, rather than balances, excitation. The inhibition resulting from the phase-nonspecific LGN input depends strongly on contrast but not orientation, and can produce nonlinear effects similar to the normalizing inhibition of Carandini and colleagues, although the biological substrate is quite different. For example, preliminary simulations indicate that one such nonlinearity, cross-orientation inhibition (suppression of response to a stimulus of the preferred orientation by simultaneous presentation of a null-oriented stimulus), arises naturally in our model from the anti-phase inhibition (A. Hoffman and present authors, unpublished observations).

Additional nonlinearities include effects of contrast on temporal responses. Preliminary results [Priebe et al., 1997] indicate that the adaptation current contributes to contrast-dependent temporal phase advance, and that the addition of frequency-dependent short-term synaptic depression [Markram and Tsodyks, 1996, Stratford et al., 1996, Abbott et al., 1997, Tsodyks and Markram, 1997] to our model can largely account for contrast effects on both temporal phase and temporal frequency tuning in cat V1 (related observations have been made independently by Chance et al. [1997]). Thus, we hypothesize that local rather than global circuitry can account for contrast nonlinearities as well as for contrast invariance, although this will require further study.

1.4.7 Developmental and Functional Implications

The LGN input to simple cells assumed here – a "Hubel-Wiesel" model – has previously been shown to develop from correlation-based rules of synaptic development under simple assumptions [Miller, 1994]. Our model uses cortical connectivity that should also arise from such rules, as generalized to include plasticity of inhibitory synapses. Thus our model circuit, in addition to explaining many properties of cortical functional responses, naturally suggests its own developmental origin: central features of the intracortical and thalamocortical connectivity in cortical layer 4 may arise together through similar mechanisms of correlation-based development.

The anti-phase inhibition that results from such development allows cortical circuitry local to a small number of iso-orientation columns to distinguish the form (orientation) from the intensity (contrast) of an oriented stimulus. Suppose that spatial phase is ignored, so that all cortical cells preferring the same orientation receive the same single component of LGN input, an instantaneous synaptic input rate. Then the two input variables of contrast and orientation are confounded. An intermediate input rate might correspond to the preferred orientation at low contrast, or a non-preferred orientation at higher contrast. To disambiguate these, a comparison across cells of all preferred orientations is needed, to determine for each cell whether it is receiving more or less input than cells of other preferred orientations. Recent models achieve this comparison through circuitries in which intracortical inhibition is broader in orientation than excitation [Ben-Yishai

et al., 1995, Somers et al., 1995].

Our model instead takes into account spatial phase, and predicts that the computation in layer 4 is local. If the cell and its anti-phase partners are receiving similar input (Fig. 1.6B), the preferred orientation is *not* being seen, and the cell does not respond, regardless of stimulus contrast (intensity). If the cell is receiving input and its anti-phase partners are not (Fig. 1.6A), the preferred orientation (form) is present, and the cell responds; the strength of its response reflects the stimulus contrast.

This provides a specific example of a possible, more general principle of cortical organization that should arise from correlation-based development: a cortical cell should respond to the *difference* between its preferred stimulus (call it $\mathcal{P}$) and its "anti-preferred" stimulus (call it $\bar{\mathcal{P}}$), rather than to the preferred stimulus alone. By anti-preferred stimulus, we mean the stimulus that evokes an input pattern most anti-correlated with that of the preferred. A stimulus that is *un*correlated with the preferred stimulus pattern (*i.e.* evokes an uncorrelated input pattern), while partially stimulating the cell, should also stimulate the cell's "anti-preferred partners" (cells with preferred stimulus $\bar{\mathcal{P}}$). Correlation-based inhibition prevents the cell from responding to this inappropriate stimulus, regardless of stimulus intensity. This hypothetical principle, that a cell selective for $\mathcal{P}$ actually responds to "$\mathcal{P}$ AND NOT $\bar{\mathcal{P}}$", might aptly be termed "dialectical," in recognition of the philosophical school that argues that objects exist and are known only in relation to their opposites (*e.g.* Merleau-Ponty [1962]).

Correlation-based connectivity also suggests a possible developmental explanation for columnar organization (the vertical invariance of RF properties). Why don't inhibitory and excitatory neurons in a given column take on opposite preferred orientations (or ocular dominance (OD))? We hypothesize that interconnected excitatory and inhibitory neurons share RF properties that are shared by both $\mathcal{P}$ and $\bar{\mathcal{P}}$, and differ in the RF properties that distinguish these. Thus, orientation (and OD) is invariant in a column, while spatial phase varies [Freeman et al., 1997], because the most anticorrelated stimulus pair has identical orientation (and OD) but opposite spatial phase. It will be of great interest to determine whether such a principle might apply to cortical representations of other sensory modalities.

Chapter 2

Temporal Frequency Tuning

Abstract

There are several response properties that distinguish cortical simple cells of cat primary visual cortex from the geniculate cells that provide their input. Many have argued that any explanation of these differences will necessarily involve intracortical circuitry, in particular the massive intracortical excitatory connections that have been observed throughout neocortex. In the first chapter, we demonstrated that, in fact, a simple model of the thalamocortical connections to layer 4 simple cells in conjunction with a developmentally motivated pattern for the intracortical inhibition can robustly account for the contrast-invariance of the orientation simple cell orientation tuning. Intracortical excitatory connections serve a secondary function of amplification of the tuned response. This model shows that the thalamocortical input due to drifting gratings can be divided into two components: a mean component that is predicted to be net inhibitory, and a temporal modulating component that is excitatory. We now show that this same model circuit unexpectedly accounts for the low-pass shift in the temporal frequency tuning curves of cortical cells relative to geniculate cells. This arises when NMDA-mediated excitatory conductances are included in the geniculocortical synapses, at levels observed experimentally in thalamocortical synapses. The longer time course of the NMDA conductance, as compared to AMPA conductances, selectively preserves the modulation of lower frequency gratings, while it dampens out the modulation of higher frequencies. Since cortical firing is predicted to occur only with temporal modulation strong enough to overcome the inhibitory mean component, the NMDA-induced dampening of higher frequency modulations prevents cortical response to higher-frequency inputs.

2.1 Introduction

Cells in the primary visual cortex (V1) of cats [Ikeda and Wright, 1975, Movshon et al., 1978a, Holub and Morton-Gibson, 1981, Lee et al., 1981b, Orban et al., 1981, Bisti et al., 1985, Orban et al., 1986, DeAngelis et al., 1993a] and monkeys [Foster et al., 1985, O'Keefe et al., 1998] fail to respond to fast-moving stimuli that evoke strong responses in the lateral geniculate nucleus (LGN) [Derrington and Fuchs, 1979, Troy, 1983, Sclar, 1987], the source of visual inputs to V1. This is exemplified by the temporal frequency tuning of V1 neurons, defined by studying neuronal response to a drifting sinusoidal luminance grating, of the neuron's preferred spatial frequency and orientation, as a function of the grating's temporal frequency. Cortical cells cease responding, with increasing temporal frequency, at frequencies to which LGN cells respond vigorously. Similar temporally low pass behavior is seen in other cortical areas, *e.g.* primary auditory cortex [Creutzfeldt et al., 1980]. The origin of this temporal behavior remains an outstanding puzzle for the understanding of cerebral cortical circuitry.

We recently introduced a model of the circuitry of layer 4, the input-recipient layer, of cat V1. This model addressed the orientation tuning of cortical neurons, including its invariance with changes in stimulus contrast and its variation with other stimulus parameters such as spatial frequency, while satisfying constraints from intracellular studies [Troyer et al., 1998]. Here we report that this same circuit model provides a natural and unexpected explanation of cortical temporal frequency tuning.

The proposed circuit includes two features that are essential to this explanation of temporal frequency tuning:

1. The inhibition received by a cell is "spatially opponent" to, or "push-pull" with, the excitation received, as suggested by many experiments [Palmer and Davis, 1981, Ferster, 1988, Hirsch et al., 1998]. Cells in layer 4 of cat V1 are "simple cells", cells that are excited by light or by dark in alternating, oriented subregions of visual space, known as ON-subregions (light-preferring) or OFF-subregions (dark-preferring) respectively [Hubel and Wiesel, 1962, Gilbert, 1977, Bullier and Henry, 1979]. The spatial opponency of inhibition and excitation means that these simple cells receive inhibition that is evoked by dark in ON subregions and by light in OFF subregions.

2. This inhibition is dominant, that is, it is significantly stronger than the excitation received by a cell.

The dominance of inhibition dictates that simultaneous activation of both the excitation and the inhibition received by a cell will fail to activate a cell. The spatial opponency of excitation and inhibition ensures that a drifting sinusoidal grating of the preferred orientation and spatial frequency alternately evokes excitation and inhibition, *i.e.* as light bands of the grating alternately drift over ON- and then OFF-subregions of the RF. Hence, given these two circuit features, cat layer

4 cells will cease to respond, with increasing temporal frequency, when the alternating activations of excitation and inhibition become sufficiently close together in time as to be effectively simultaneous. Our modeling task thus becomes to determine, in terms of the biophysics of the cells and synapses, at what temporal frequency excitation and inhibition become effectively simultaneous, and under what conditions this can account for the observed temporal frequency tuning of cortical cells.

2.2 Experimental Background

2.2.1 Temporal Frequency Tuning

The problem that the model is to address in this chapter is the shift in temporal frequency tuning towards lower frequencies from the thalamus to the cortex. The purpose of this section is to describe the experimental evidence for the existence of such a problem. To what extent do we know there is in fact such a shift from the LGN to primary visual cortex in the cat? How general is this transition to lower temporal frequencies in thalamocortical processing, in terms of various species and sensory modalities?

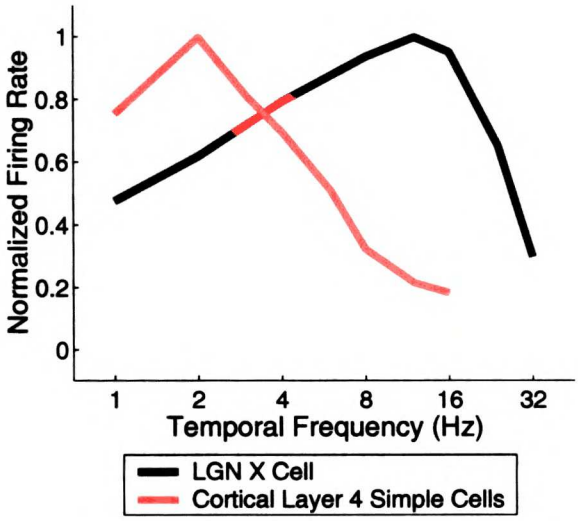

Figure 2.1:

Temporal frequency tuning in LGN versus cortex: Experimental Results. Cortical layer 4 simple cell data from A. Saul and A. Humphrey (unpublished results) and LGN X cell data from Sclar [1987]. Cortical temporal frequency tuning curves, derived from the amplitude of F1 modulation of firing rate responses at various temporal frequencies, were normalized individually to each cell's maximum response. The normalized tuning curves were then averaged over all 18 cells, and this averaged tuning curve was renormalized to its peak (red trace). LGN cell data is normalized tuning curve of amplitude of F1 modulation of firing rate responses (blue). Temporal frequency cutoffs, defined as the frequency where the response reaches half of the maximal response, are 6.1Hz for the cortical cells, and 27.4Hz for the LGN cell.

While several groups have studied the temporal properties of both cat LGN and visual cortex, there are few studies where the two regions were studied simultaneously, in the same preparation. To illustrate the degree of the problem, we show some representative temporal frequency tuning curves from two different groups in figure 2.1. A. Saul and A. Humphrey provided us with experimental results stemming from a number of different studies, from 18 cortical cells that were identified as layer IV simple cells. We compare the population tuning curve for these layer IV cells, with a tuning curve from a representative LGN X cells from Sclar [1987] in figure 2.1. From this data, it appears that there is a shift of 2-3 octaves in the temporal frequency tuning curves from LGN X cells to layer IV V1 cells. The temporal frequency cutoffs, defined as the frequency where the response reaches half of the maximal response, are 6.1Hz for the cortical cell population tuning curve, and 27.4Hz for the LGN cell. What follows is a general survey of the experimental literature of temporal frequency tuning, in order to determine if this strong a shift in tuning from the thalamus to the cortex is a reasonable assumption.

Studies of temporal frequency tuning in cat LGN

Several groups have measured temporal frequency tuning curves of cat LGN cells. One of the earlier detailed studies was done by Derrington and Fuchs [1979] who measured temporal frequency tuning curves for 13 X cells and 17 Y cells. Apart from showing some representive curves for each type, they quantified the results with the distributions of peak frequencies. The two distributions overlapped: X cells had a mean ($\pm$ standard deviation) peak frequency of 7.9 Hz ($\pm$ 2.6), and Y cells had 10.2 Hz ($\pm$ 1.9).

Troy [1983] found that X and Y cells did not have distinguishable temporal frequency tuning curves at low contrast. He measured contrast sensitivity, using gratings that were at threshold levels of contrast, and recorded from 27 X cells and 35 Y cells. Y cells were somewhat more sensitive at all frequencies, but the shapes of the two contrast sensitivity curves, averaged over all the cells of each type separately, were almost identical. For both X cells and Y cells, the optimum frequency was about 5 Hz.

These results suggest that the temporal frequency tuning curves for X and Y cells are very similar at low levels of contrast but that Y cells are somewhat more sensitive to higher temporal frequencies than X cells at higher contrast. This is consistent with the study performed by Sclar [1987]. He measured temporal frequency tuning curves for 27 X and 51 Y cells in the cat LGN at four different levels of contrast to see if the shifts in temporal frequency tuning with contrast were similar to those that had been observed in the retina, as had been reported in a series of papers by Shapley and Victor. Both X and Y cells showed a sharper increase in response, with increasing contrast, at higher temporal frequencies than at lower frequencies. However, this effect was stronger in Y cells than in X cells, resulting in a larger difference in their temporal frequency tuning curves at high contrast than at low. At 10% contrast, X cells had a mean peak frequency

of 8.1 Hz (± 6.0), and for Y cells it was 8.4 Hz (± 3.5). At 80% contrast, X mean peak frequency was 12.0 Hz (± 8.23) and for Y cells is was 16.8 Hz (± 7.6).

A more recent study similarly found little difference in the temporal frequency tuning of X and Y cells, but they did find a distinction in the temporal tuning of LGN cells divided along a different axis: non-lagged versus lagged cells [Saul and Humphrey, 1990]. Non-lagged cells were identified as having a strong transient excitatory response to the onset of a flashed spot, as opposed to the lagged cells which responded to the onset of a spot with a transient decrease in firing rate, followed by a strong sustained response. The lagged cells also often responded with a transient excitatory response to the offset of the stimulus, before returning to their background firing rate. The lagged cells, therefore, had a delayed response both to the onset and the offset of the flashed spot, and the two categories were separated quantitatively with the latencies of the increase and decrease in firing rates to the onset and offset of a flashed spot. In cunjunction with the standard definitions of X- and Y- cells, they divided their data into the four categories of lagged X (X_L), lagged Y (Y_L), non-lagged X (X_N), and non-lagged Y-cells (Y_N).

	Optimal Frequency Spots	Resolution Spots	Optimal Frequency Gratings	Resolution Gratings
X_L	1.5 ± 0.2 (24)	9 ± 1 (24)	3.1 ± 0.9 (9)	17 ± 4 (9)
X_N	3.4 ± 0.3 (63)	22 ± 2 (63)	4.7 ± 0.8 (18)	25 ± 4 (18)
Y_L	4.3 ± 0.6 (4)	13 ± 1 (4)	4.7 ± 1.1 (4)	18 ± 3 (4)
Y_N	4.1 ± 0.5 (21)	22 ± 1 (21)	6.0 ± 0.7 (10)	26 ± 5 (10)

Table 2.1:

Saul and Humphrey [1990] LGN Temporal Frequency Tuning Results
Optimal temporal frequencies and temporal frequency resolution, defined as the frequency when the response reaches 10% of the maximal response, for the four categories of cells defined in the text. All values are means ± standard errors. All units are Hz. Number of cells for each measurement are in parentheses.

They then measured temporal frequency tuning curves using both standing spots that modulated in luminance sinusoidally, and with drifting sinusoidal gratings. With both measures, they use the amplitude of the first harmonic (F1) of the response to determine optimal temporal frequencies, and temporal frequency resolutions, defined as the frequency when the response reaches 10% of the maximal response (table 2.1). The distributions of optimal frequencies overlapped over all four categories; there was some tendency for the non-lagged cells to have higher optimal frequencies than the lagged cells. The strongest distinction was in temporal resolution. The lagged cells stopped responding at much lower frequencies than the non-lagged cells. All cells responded better at higher frequencies to drifting gratings than to the modulated spots, such that the temporal res-

olutions for all of the cells are higher when measured with the gratings. The increase in resolution is bigger for lagged cells than for non-lagged, such that the difference in resolution between lagged and non-lagged is not as large when using gratings than when using spots, but it is still significant.

More recent studies have reported LGN temporal frequency tuning curves tuned to significantly higher frequencies than those reported in Saul and Humphrey [1990]. Hamamoto et al. [1994] recorded temporal frequency tuning curves for 38 LGN X-cells, using drifting sinusoidal gratings of moderate contrast (30%). They reported mean bandwidths (width of tuning curve at half maximal response, measured in octaves) of 4.2 ± 0.4. The mean temporal resolution, defined as the frequnecy where the reponse reaches 10% of the peak was 36.9 ± 3.3 Hz. Dan et al. [1996] found temporal filters with significantly higher peak frequencies than those observed in Saul and Humphrey [1990]. They measured the temporal filtering properties of LGN X-cells using both full-field white noise stimuli and drifting sinusoidal gratings. With both methods, they found increases in responses with increasing frequency up to 15Hz, with filter functions roughly proportional to the square of the frequency.

It is possible that a number of these differences in results from different groups can be explained from differences in sampling due to different experimental techinques. Alan Saul has pointed out that the lagged cells that are tuned to lower temporal frequencies, although quite numerous, tend to be small cells that are overlooked with standard low-impedance electrodes, and that many studies might, therefore, be excluding them inadvertently from reported population tuning curves (personal communication). Dan et al. [1996] mention this issue in their discussion section as well.

Studies of temporal frequency tuning in cat V1

One of the earliest extensive studies of the temporal frequency tuning of cat cortical cells in area 17 was done by Ikeda and Wright [1975]. They first divided the cortical cells they encountered into 4 categories, using their response to stationary light and dark bars presented in the center of the receptive field for the categorization. A cell was called "sustained" if its response remained at least 3 spikes/sec above the spontaneous firing rate throughout a 5 second bar presentation, and "transient" if it did not. They also determined if the cell was simple or complex, both by whether or not it responded similarly to light and dark stationary bars, and by standard receptive field mapping techniques to determine if it had segregated on and off subfields. With this method they found all 4 subtypes in area 17, simple sustained (SS), simple transient (ST), complex sustained (CS) and complex transient (CT).

To characterize temporal frequency tuning, they used both drifting sinusoidal gratings, and stationary contrast reversal flickering gratings, using square-wave contrast modulation. In general, they found that sustained cells, both simple and complex, preferred low frequencies, either prefering the lowest frequency they tested (0.25 Hz), that is, a pure low pass filter, or having very shallow fall offs at frequencies lower than the preferred frequency. Transient cells, both simple and complex,

were all band-pass, with higher preferred frequencies than for sustained cells. This categorization in the temporal domain is not terribly compelling. They originally divided cells into sustained and transient pools based on whether or not they responded well to stationary bars maintained for 5 seconds, which can be thought of as one bar of a 0.2 Hz flickered grating, so the fact that the two pools had different low frequency cutoffs was going to be true by definition. They don't comment on the statistical significance of the separation of the means for peak temporal frequency and high frequency cutoff, and they do have overlapping ranges (table 2.2). However, the sustained and transient cells also had distinctive spatial frequency tuning curves, essentially the opposite traits from the temporal tuning. Sustained cells were bandpass in spatial frequency, and transient cells were low pass in spatial frequency (down to 0.1 cycles/deg). Between the spatial and temporal separations of the two populations, Ikeda and Wright argued that they had found the cortical cells in area 17 that corresponded to the sustained (X cell) and transient (Y cell) retinal and geniculate pathways.

	SS ($n = 6$)	CS ($n = 12$)	ST ($n = 6$)	CT ($n = 12$)
Peak Temporal Frequency				
Mean	1.88	0.97	3.37	3.0
Range	0.1-5.0	0.1-4.0	1.0-8.0	1.0-10
High Frequency Cutoff				
Mean	5.8	4.9	9.6	6.7
Range	3.5-12.0	3.0-9.0	6.0-12.0	2.0-12.0
Low Frequency Cutoff				
Mean	0.13	0.17	1.37	1.10
Range	0-0.5	0.-0.7	0.25-2.5	0.25-4.0

Table 2.2:

Ikeda and Wright [1975] Temporal Frequency Tuning Results

Peak and high and low temporal frequency cutoffs for the 4 categories of cells. All units are Hz. Cutoff frequencies are defined as the frequency when the response reaches half of the maximal response.

Movshon et al. [1978a] studied the spatial and temporal frequency tuning of cells in both areas 17 and 18 of the cat. Their results essentially confirmed the results of Ikeda and Wright [1975] in that they found two categories of cells that divided along the same lines of their spatial and temporal tuning properties as the sustained and transient cells of Ikeda and Wright [1975]. However, they did not find both types of cells in area 17. Sustained cells were found in area 17 and transient cells were found in area 18. The study focuses on spatial frequency tuning; the temporal tuning

properties presented are not nearly as extensive. They report that area 17 cells responded best to gratings slower than 2-4 Hz, and progressively less well to faster gratings, and that there was a shallow fall off, if at all, for low frequencies down to 0.5 Hz. While some area 18 neurons had low-pass tuning curves similar to the area 17 neurons, most of them were band pass, with preferred frequencies from 2-8 Hz. They report that both areas had similar high frequency cutoffs, although they do not report what the values were.

The anatomical separation of the two physiological types of cortical cells reported in Movshon et al. [1978a] to areas 17 and 18 parallels the difference in the projections the two receive from the LGN. There is physiological evidence that Area 17 receives mostly, or perhaps purely, X cell input [Ferster, 1990b,a], although anatomically it receives both X and Y cell input [Humphrey et al., 1985a,b]. Both anatomically [Humphrey et al., 1985b] and physiologically [Ferster, 1990b,a], area 18 receives input only from Y cells. However, a recent study argued that regions in area 17 preferring low spatial frequencies are distributed as sparse patches, which in turn was argued to support a patchy distribution of Y cells to area 17 [Shoham et al., 1997, Hübener et al., 1997]. This raises the possibility that the presence or absence of "Y-like" properties in area 17 might simply reflect sampling differences in different studies.

Holub and Morton-Gibson [1981] also measured the spatial and temporal frequency tuning curves of cells near the border of areas 17 and 18 in cat visual cortex. They had a different way of categorizing cells. For each grating, they measured both the mean firing rate and the amplitude of the first harmonic fit of the PSTH. If the temporal (and spatial) tuning curves, with peak response normalized to 1, using each of these two measures, both matched at two or more frequencies higher than the peak frequency, then the cell was called "periodic." If the tuning curves both matched below the peak frequency, but not above, then it was called "intermediate." Cells whose two tuning curves did not agree except at the lowest frequencies were called "aperiodic." They point out that these categories were related to the Hubel and Wiesel [1962] standard definitions of simple and complex, but not identical. Roughly 80% of periodic cells were simple, and 80% of aperiodic cells were complex, and the intermediate cells were of either type or hard to classify at all. They also comment that they believe there was a smooth distribution of "periodicity" in the cells they encountered.

They do not comment on finding any temporal frequency tuning curves that were low-pass, like the sustained cells described above. Instead of reporting high and low frequency cutoff frequencies, as in Ikeda and Wright [1975], they report bandwidths, which are defined as the ratio of the high to low cutoff frequencies, in octaves, a measure that implies they always found a non zero low frequency cutoff. Overall, their results are similar to the properties of the transient cells of Ikeda and Wright [1975] (whose mean ST cell would have a bandwidth of 2.8 and mean CT cell a bandwidth of 2.6) or the area 18 cells of Movshon et al. [1978a], although with slightly higher peak frequencies (table 2.3). They found no trend in peak frequency with periodicity, but they report a statistically

	Periodic	Intermediate	Aperiodic	Aggregate
Peak Frequency	$5.93 \pm 3.95(36)$	$3.93 \pm 2.27(26)$	$5.73 \pm 3.97(26)$	$5.28 \pm 3.62(88)$
Bandwidth	$2.54 \pm 1.17(25)$	$2.98 \pm 0.93(21)$	$3.64 \pm 1.41(19)$	$3.01 \pm 1.24(65)$

Table 2.3:

Holub and Morton-Gibson [1981] Temporal Frequency Tuning Results

Peak temporal frequencies (Hz) and bandwidths (octaves) for the three types of categories of cells described in the text. Values are means ± standard deviation. Numbers in the parentheses are the number of cells that went into the measurement. Bandwidth is defined as the ratio of the high to low cutoff frequencies, where the cutoff is the frequency at which the response reaches half of the maximal response.

significant trend for more periodic cells to have a narrower bandwidth.

Another study demonstrated that temporal frequency tuning of cat cortical cells was different when gratings were presented in the preferred versus the non-preferred directions [Saul and Humphrey, 1992]. While they found a wide range of tuning curves in both directions, overall their population results are similar to the values reported by Holub and Morton-Gibson [1981] (table 2.4). Furthermore, their population results in the preferred direction were very similar between Area 17 and Area 18, and between simple and complex cells.

Since in our modeling work we only study the responses of cells to stimuli in one direction, we will not discuss the variety of cells that they found. To summarize their findings, they found a number of cells that had suppressed responses to stimuli in the non-preferred direction only at low frequencies; at high frequencies the cells had very similar responses, and therefore lost their direction selectivity. For some cells, the "preferred" direction actually switched from low to high temporal frequencies. This is seen in their group data by comparing the distribution of optimal frequencies to temporal resolutions. For many cells, the optimal frequencies were higher in the non-preferred direction than in the preferred, while the temporal resolutions were not significantly different. This can be observed, to some extent, in the population data as well, if considering shifts in the two directions in terms of octaves, particularly in Area 18 (table 2.4).

One group has reported responses in cat cortex, LGN, and retina to a few different temporal frequencies in the same set of experiments [Lee et al., 1981a,b]. Their sampling of temporal frequencies was sparce, recording only at 1, 4 and 10 Hz. However, they conclude that "the temporal frequencies evoking maximal responses were clearly lower in simple cortical cells than in retinal ganglion cells and LGN neurons." All 25 retinal ganglion cells and 18 LGN cells, both sustained and transient (X and Y), responded best at 10 Hz. But, of the 31 simple cells that they recorded from, only 3 responded best at 10 Hz, 18 at 4 Hz and 10 at 1 Hz.

	Area 17	Area 18	Simple	Complex
Peak Frequency (Pref.)	$2.9 \pm 0.3(54)$	$3.2 \pm 0.6(14)$	$2.8 \pm 0.3(53)$	$3.4 \pm 0.5(18)$
Peak Frequency (Nonpref.)	$4.5 \pm 0.4(42)$	$7.3 \pm 0.6(9)$	$4.9 \pm 0.5(37)$	$5.3 \pm 0.8(14)$
Resolution (Pref.)	$13.5 \pm 1.0(57)$	$14.8 \pm 1.3(14)$	$12.9 \pm 0.9(53)$	$16.2 \pm 2.1(18)$
Resolution (Nonpref.)	$15.5 \pm 1.2(41)$	$21.0 \pm 2.3(9)$	$15.5 \pm 1.2(36)$	$19.0 \pm 2.3(14)$
Bandwidth (Pref.)	$3.3 \pm 0.2(50)$	$3.4 \pm 0.3(12)$	$3.3 \pm 0.2(46)$	$3.4 \pm 0.3(16)$
Bandwidth (Nonpref.)	$3.1 \pm 0.1(35)$	$2.7 \pm 0.1(8)$	$3.0 \pm 0.1(32)$	$3.2 \pm 0.2(11)$

Table 2.4:

Saul and Humphrey [1992] Cortical Temporal Frequency Tuning: Preferred Versus Non-Preferred Directions

Peak temporal frequencies (Hz), temporal resolution (Hz), and bandwidths (octaves) using gratings in the preferred and non-preferred directions. Values are means ± standard errors. Numbers in the parentheses are the number of cells that went into the measurement. Resolution is the frequency where response reaches 10% of the maximal response. Bandwidth is defined as the ratio of the high to low cutoff frequencies, where the cutoff is the frequency at which the response reaches half of the maximal response.

There have been some more recent studies that have recorded temporal frequency tuning curves from adult cat cortical cells. Often, these studies are focused on issues other than the actual tuning curves, some of which will be discussed below, and therefore only give qualitative descriptions of the actual tuning. Baker [1990] recorded from between 31 and 38 cortical cells, and report a range of preferred frequencies from 0.7-4.0 Hz. He did not find many low-pass cells, which he points out is different from the results reported by Movshon et al. [1978a]. McLean and Palmer [1994] recorded multidimensional spectral responses, including temporal frequency, from 50 simple cells in area 17 and 18. Details of their results are described more below, but as far as the absolute temporal frequency tuning curves are concerned, they say simply that their results are consistent with previous results, in the sense that, in general, area 17 simple cells respond well to low temporal frequencies, and area 18 simple cells are more bandpass in their temporal tuning.

More recently, it has been demonstrated that the spatial and temporal frequency tuning curves of simple cells can be predicted from their spatiotemporal receptive field structure [DeAngelis et al., 1993a,b]. Spatiotemporal receptive fields were generated, using standard reverse correlation techniques in response to noise stimuli, for 91 simple cells from adult cats, and temporal frequency tuning curves were recorded for 24 of those cells. The population characteristics of the 91 predicted temporal frequency tuning curves show preferences for low frequencies, both in the distributions of peak frequencies (mean of 2.6Hz) and in the high frequency cutoff, again defined as the frequency

of half-maximal response (mean of 7.2 Hz). For most of the cells where a comparison to the actual tuning curves was possible, the predicted curves were accurate at all frequencies. For some cells there were deviations at frequencies lower than the preferred frequency; in these cases, the predicted curves overestimated the responses at low frequencies. These deviations did not effect the accuracy of the predictions of either the peak frequency or the high frequency cutoff.

The same group has also done a similar study of the spatiotemporal receptive field structure of LGN cells [Cai et al., 1997]. Receptive fields were generated for 67 LGN cells, 56 of which were identified as X cells and 11 as Y cells. They did not find statistically significant differences between X and Y cells with respect to most spatial and temporal parameters, so they pooled both types for the population results reported in the paper. The LGN study is performed purely in the time domain and there is, therefore, no analysis of the predicted temporal frequency response from these experimentally determined receptive fields. However, they do make comparisons of the receptive field structure of the LGN cells and the simple cells from their earlier study, using the two main time domain parameters that they calculated from the receptive fields: the peak response time (T_{peak}) and the duration of the temporal response profile (D). Both of these quantities were significantly shorter for the LGN cells compared to the simple cells (table 2.5). Simple cells peak later and have more prolonged responses compared to their LGN inputs, consistent with the fact that they prefer lower temporal frequencies than their LGN inputs, when converting these response profiles to the frequency domain.

Development of temporal frequency tuning

There have been a small number of studies of the development of temporal frequency tuning in cat visual cortex. The studies described in the previous section [DeAngelis et al., 1993a,b] were done not only with adult cats, but also with 4-week and 8-week old kittens, in order to study the development of spatiotemporal receptive fields, and their match to frequency tuning curves. The distributions of both optimal frequency and high frequency cutoff shifted to higher frequencies with increasing age. The mean values of optimal temporal frequency were 1.42 Hz at 4 weeks, 1.75 at 8 weeks, and 2.6 Hz for adults as mentioned above. High frequency cutoffs were 4.1 Hz at 4 weeks, 4.7 Hz at 8 weeks and 7.2 Hz for adults.

The LGN study by the same group, described above, was also performed at the same ages, so a direct comparison of the differences in the temporal structure of LGN and cortical receptive fields could be performed for all three age groups [Cai et al., 1997]. The fact that simple cells respond later and have longer lasting responses held true at all ages (table 2.5). Furthermore, some of the developmental shift in the temporal properties of simple cells towards faster responses is paralleled by a similar shift in their LGN inputs. However, the LGN does not seem to be speeding up as quickly, with age, as the cortex. The difference in these two parameters between the LGN and the cortex shrinks with age. As the authors observe in the discussion of Cai et al. [1997], "although

temporal selectivity improves with age in both areas, there does appear to be an early component of development that is mainly intracortical. Both T_{peak} and duration decline substantially for simple cells between 4 wk and 8 wk postnatal, whereas the corresponding values for LGN neurons remain essentially unchanged during this time period." It should also be noted that a developmental study of the temporal frequency tuning of LGN cells in Old World monkeys reported a significant shift in the high temporal frequency cutoff towards higher frequencies with increasing age [Hawken et al., 1997]. At the time of birth, no cells had cutoffs above 10Hz, while by the age of eight months many cells had cutoffs of at least 20Hz. Since there is no companion study of the development of primary visual cortical temporal frequency tuning in primates, it is not possible to tell if this shift in the LGN could completely account for a potential developmental shift in the cortex.

	T_{peak} LGN	T_{peak} Cortex	D LGN	D Cortex
4-week kitten	55.4 (36)	114 (79)	108.0 (36)	242.4 (79)
8-week kitten	58.2 (41)	90.7 (63)	115.9 (41)	211 (63)
Adult	35.9 (67)	68.2 (91)	79.6 (67)	139.6 (91)

Table 2.5:
Development of Temporal Receptive Field Properties for LGN and Cortex
Mean peak response time (T_{peak}) and duration of the temporal response profile (D) for cat LGN and cortical simple cells for three different age groups as reported in Cai et al. [1997], DeAngelis et al. [1993a]. (For details of the definitions of these two parameters, please refer to Cai et al. [1997], DeAngelis et al. [1993a]). All values are in msec. Numbers in the parentheses are the number of cells for each category.

There has also been a study looking at the development of temporal frequency tuning under different types of visual experience [Gary-Bobo et al., 1995]. Kittens were raised for 6 weeks either in a normally lit room (NR), in the dark (DR), or in the dark for 6 weeks followed by a 6 hour period of exposure to light only to be returned to the dark before electrophysiology was performed 12 hours later (DR+6). They recorded from a total of 346 cells in area 17 in all the animals, and reported the distribution of peak optimal temporal frequencies for 36 cells from DR kittens, 29 from NR kittens and 69 from the DR+6 kittens. The distribution of peak temporal frequencies for cells from the DR kittens was shifted towards lower frequencies compared to the NR kittens. Furthermore, the DR+6 distribution closely resembles the NR distribution. Combining these results with the results of DeAngelis et al. [1993a] indicate that something is preserving the lower frequency tuning of 4 week old kittens at least until 6 weeks when they are raised in the dark. While this could simply

be due to a freezing of the temporal tuning of the LGN inputs in kittens during dark rearing, it is interesting to note that the developmental study of LGN tuning in primates described above noted that binocular deprivation did not prevent the normal development of temporal tuning [Hawken et al., 1997].

Temporal frequency tuning versus contrast

There is a small amount of data in the literature on the shifts of temporal frequency tuning with contrast, indicating that visual cortical temporal selectivity shifts towards higher frequencies with increasing contrast, consistent with what is observed both in the LGN, as described above [Sclar, 1987], and in the retina [Shapley and Victor, 1978]. Albrecht [1995] shows temporal frequency tuning curves at multiple contrasts for two cat cells, and one monkey cell. The cat cells shift to higher frequencies, both in terms of peak temporal frequency and high frequency cutoff, with increasing contrast (table 2.6). For more details about what is known experimentally with regards to shifts in temporal frequency tuning versus contrast, refer to Kayser et al. [1999].

Contrast	Center Freq. Cell #1	Cutoff Freq. Cell #1	Center Freq. Cell #2	Cutoff Freq. Cell #2
5%	3.9	7.9	3.2	7.8
10%	4.5	10.5	2.8	9.8
15%	6.2	11.9	3.4	10.6
30%	8.2	13.9	4.1	14.1
60%	8.1	15.2		
70%			5.4	15.4

Table 2.6:
Temporal Frequency Tuning versus Contrast from Albrecht [1995] Cutoff frequency refers to the high frequency at which the response reaches 50% of the maximum. Cell #1 was bandpass, and also had a low frequency cutoff ranging from 2.2 Hz to 0.5 Hz with different contrasts, while cell #2 was essentially low pass, with response exceeding half of its maximal response at all frequencies tested lower than the peak. All values are in Hz.

Temporal frequency versus velocity tuning

One question that has been discussed throughout these studies of the temporal properties of visual cortex is the potential confusion between a cell's sensitivity to the temporal frequency versus the

velocity of a visual stimulus. If one tests the response of a cell to gratings of different temporal frequencies, but only at a single spatial frequency, then it is impossible to tell if the tuning is determined by how rapidly the bars of the grating repeat, or how quickly a single bar of the grating passes over the receptive field; that is, if the cell is tuned to temporal frequency or velocity. These are significantly different properties of the stimulus. Temporal frequency tuned cells are answering a question about a property of the stimulus that is local in space; what is the time course of the change in luminance at one point in space? Velocity is a non-local spatial property of the stimulus; how quickly does the stimulus move from there to here?

There is a simple relationship between the temporal frequency, spatial frequency and velocity of a drifting grating. The velocity in degrees per second is equal to the temporal frequency in cycles per second divided by the spatial frequency in cycles per degree. Therefore, repeating the temporal frequency tuning curves at non-optimal spatial frequencies can resolve the confusion. For a grating with higher spatial frequency and narrower bars than the preferred, the velocity of a single bar will be slower at a fixed temporal frequency, and the reverse is true for lower spatial frequencies. Therefore, if the temporal frequency tuning curves at different spatial frequencies are simply scaled versions of each other, or if the spatiotemporal filter is separable as the product of a spatial filter and a temporal filter, then the cells are in fact tuned to temporal frequency. The other possibility is that the spatiotemporal filter is tilted, such that gratings with the same velocity evoke similar responses, independent of spatial and temporal frequencies.

Tolhurst and Movshon [1975] originally addressed this question. They recorded from 15 simple and 11 complex cells in cat striate cortex, and measured contrast-sensitivity curves. They report that the shapes of the temporal frequency tuning curves were very similar for a wide range of spatial frequencies, and vice versa, for all of the cells they looked at. Ikeda and Wright [1975] address this as well for four of the cells from their study. Three sustained cells, two simple and one complex, were reported to be tuned to temporal frequency and one complex transient cell was better described as velocity tuned.

Holub and Morton-Gibson [1981] also extensively studied this question. They recorded temporal-frequency response functions at two different spatial frequencies, and found insignificant differences in the shapes and peaks of the two tuning curves. Differences in shapes were assessed both with bandwidths and with the high frequency roll-off which is defined as $20log(R_o/R_p)$ where R_o is the response one octave higher than the optimal frequency, and R_p is the preak response. Furthermore, they demonstrated that spatial-frequency response functions that held velocity constant rather than temporal frequency could be predicted from the individual spatial and temporal frequency tuning curves, calculated at the peak temporal and spatial frequencies respectively.

Bisti et al. [1985] found similar results in area 17, but their results in area 18 were quite different. In area 17, they report that they extended the results of Tolhurst and Movshon [1975], by using high contrast drifting gratings, rather than threshold contrasts used to record contrast-sensitivity curves.

With contrasts of either 10 or 25%, they found that the shapes and positions of spatial frequency tuning curves did not vary significantly over a wide range of temporal frequencies. In area 18 they found that the optimal spatial frequency decreased with increasing temporal frequencies, although the overall bandwidth of the tuning curves stayed roughly the same. This shift, however, was only evident when using drifting gratings. With standing square-wave gratings with contrast that varied sinusoidally, the spatial frequency tuning did not shift with increasing temporal frequency, as was observed in area 17.

While the focus of this literature review so far has been on the use of sinusoidal gratings as visual stimuli, an entirely parallel, and perhaps more extensive, set of studies have been performed by various groups using light and dark bars or edges as the main stimulus to probe temporal properties of LGN and cortical neurons. Single bars are not repeating stimuli, and, therefore, by definition are probing for the tuning of cells to the velocity of a stimulus, as opposed to temporal frequency tuning. Furthermore, since bars are transient stimuli, it is unclear what the most appropriate response property is for studying the velocity tuning in response to single bars: total number of spikes, peak firing rate, or mean firing rate over a particular time window. These issues are discussed in detail in the review by Orban [1991]. In general, the trends in temporal frequency tuning described above are similar to the trends observed for velocity tuning in response to bars. There is a significant shift in temporal tuning towards lower velocities from the LGN to the cortex [Orban et al., 1985]. The same study also reported that this shift in velocity tuning is present in all cortical layers; in particular, it is not any weaker in layer 4, where the cells that receive direct LGN input are found. More recently, Baker [1990] directly compared the response properties of cortical cells in response to moving bars to their response to drifting sinusoidal gratings. He demonstrated that the optimal bar velocity for cat cortical cells can be inferred from the ratio of their preferred temporal frequency to their preferred spatial frequency, as determined from their responses to drifting gratings. It seems safe to conclude that the velocity responses to bars are not at odds with the temporal frequency results that have been described in detail.

Most recently, McLean and Palmer [1994] made thorough spatial and temporal frequency searches of a total of 50 simple cells in both areas 17 and 18. For 30 of these cells, a 3-D spectral response profile was generated, a two-dimensional spatial frequency search (including orientation and spatial frequency) and temporal frequency. For the other 20 cells, a more finely sampled 2-D search in spatial and temporal frequency was performed at the cells preferred orientation. For all of the area 17 cells and all but one of the area 18 cells, spatial and temporal frequencies were essentially separable. For the 3-D spectral response profiles, the centroid of the 2-D response profile in each spatial frequency/orientation plane did not shift with varying temporal frequency. For the cells that had more finely sampled 2-D search of spatial and temporal frequencies, the responding region had axis aligned with the coordinate axis and had a separable profile, indicating that the peak spatial frquency did not vary with temporal frequency, and vice versa. The one exceptional

area 18 neuron had a peak spatial frequency that increased with increasing temporal frequency, and vice versa, that was indiciative of a velocity tuned cell, as opposed to being tuned to temporal frequency.

Temporal Frequency Tuning in Cat Suprasylvian Cortex

It is interesting to note that the general trends of temporal frequency tuning observed in areas 17 and 18 are similar to what has been observed in other cat visual cortical regions that receive input from other thalamic cells, in addition to intracortical input. There have been a few studies of the temporal tuning properties of neurons in the posteromdeial lateral suprasylvian cortex of the cat (PMLS) [Zumbroich and Blakemore, 1987, Morrone et al., 1986]. The PMLS receives direct input from the C-laminae of the LGN and the pulvinar/lateralis posterior complex, as well as intracortical input from areas 17,18 and 19 [Sherk, 1986]. It is suspected that the spatial properties of cells in the PMLS are shaped by Y-cell and W-cell LGN input, as opposed to X-cells [Zumbroich and Blakemore, 1987]. Overall, the temporal frequency tuning properties do not seem to be significantly different from areas 17 and 18. Zumbroich and Blakemore [1987] recorded temporal frequency tuning curves for 9 cells in PMLS, and 5 cells in the lateral banks of the lateral suprasylvian (PLLS). Most of the optimal temporal frequencies fell between 2.5 and 10Hz. Four of the PMLS cells were "low-pass" meaning they did not reach half of the maximal response by the lowest frequency tested (.175Hz), but all the PMLS cells were still responding at .175 Hz. The other nine cells had bandwidths (the width from the low to high frequency cutoffs when half maximal response was attained) ranging from .63 octaves to more than 4 octaves with a mean of 2.8 octaves for PMLS and 2.5 octaves for PLLS. Morrone et al. [1986] report very similar values. 26 cells in PMLS had optimal frequencies ranging from 2-10 Hz, and 19 cells had bandwidths ranging ranging from 1-4 octaves, with the distribution peaking around 2 octaves.

There has also been a study of the development of temporal tuning in the PMLS [Zumbroich et al., 1988]. They report the mean optimal temporal frequency and the mean high temporal frequency cutoff or temporal acuity, defined as the frequency when the response reaches two standard errors above the spontaneous activity, as a function of age for the kittens from 10 days old to 8 weeks old and adults. They found a rapid increase in both the optimal and cutoff frequencies. Optimal frequencies increased from $\approx$1Hz to $\approx$4Hz, and cutoffs from $\approx$4Hz to $\approx$14Hz. Most of these shifts had occurred by three weeks of age.

Comparisons to rat and mouse

In a recent study of rat V1, Girman et al. [1999] report that cells respond best to high temporal frequencies, although from the numbers they present, the tuning does not look significantly different from cat. They looked at both the peak of the temporal frequency tuning curve, and the high frequency cutoff, which they refer to as the highest frequency to give any response. The distribution

of peak frequencies is evenly distributed from 0.43Hz to 6.88Hz, and the distribution of cutoff frequencies peaks at 27.5Hz. It is important to note that peak frequency was highest in layer IV, where it was about 7.5 Hz. The difference from cat V1 responses is more dramatic when considering the response to the velocity of the grating stimulus, as they point out in the discussion. The distribution of cutoff velocities peaked at 500 degrees/sec, and 25% of the cells they looked at responded to velocities of 700 degrees/sec, which were found most often in the vicinity of layer IV. However, considering the fact that these cells preferred very low spatial frequencies – most cells had optimal responses at less than 0.08 cycles/degree – it is possible that most of the difference from cat V1 in velocity tuning can be explained based on the difference in spatial properties rather than temporal properties.

Mangini and Pearlman [1980] looked at the velocity tuning, but not temporal frequency tuning, of cells in mouse V1. They characterized three types of cells, oriented, non-oriented and large field, non-oriented cells. Layer IV cells were mostly non-oriented cells. All of the oriented cells, and most of the non-oriented cells, had a peak velocity in the range from 10-200 degrees/sec. They found one non-oriented cell that peaked at 300 degrees/sec and one at 1000 degrees/sec. The large field, non-oriented cells, which were found almost exclusively in layer V, preferred signficantly higher velocities than the other two types, ranging from 300-2000 degrees/sec.

Comparison to temporal frequency tuning in macaque monkeys and other primates

There is experimental evidence that there is a similar shift in temporal frequency tuning between the LGN and the cortex in macaque monkeys. However, because the anatomy of V1 is more complicated than area 17 in the cat, and there is some evidence for laminar differences in the temporal frequency tuning of macaque V1 cells, it is possible that the total shift in temporal frequency tuning occurs in multiple stages in the macaque, with only some of the shift occuring in the input layers of V1 at the first stage of convergence of inputs from the LGN. Further shifts could then occur with mechanisms similar to those that we suggest in this study.

Hawken et al. [1996] directly compared the temporal frequency characteristics of monkey V1 and LGN cells. They found that both the peak temporal frequency and the high frequency cutoff were significantly higher in LGN cells than in V1 cells (table 2.7). They found no significant different in the tuning characteristics of simple and complex cells. Furthermore, they determined the laminar position for many of the cortical cells that they recorded from. While there was no significant shift in the peak frequency between cells in the layers that receive direct input from the LGN and those in the other layers, there was a difference in the high frequency cutoff. The 20Hz overall difference in high frequency cutoff between the LGN and V1 is divided into two stages, a 10Hz shift from LGN to the input layers of V1, and an additional 10Hz shift between the input and output layers of cortex.

	Peak Frequency	High Frequency Cutoff	Number of Cells
LGN	16 ± 8	41 ± 9	37
V1 (Total)	10 ± 8	25 ± 11	75
V1 (simple)	9.6 ± 7.5	23 ± 11	
V1 (complex)	11.2 ± 7.5	26 ± 12	
V1 (input layers)	12.1 ± 7.5	30 ± 9.5	19
V1 (output layers)	8.8 ± 6.2	21.2 ± 8.5	36

Table 2.7:

Hawken et al. [1996] Temporal Frequency Tuning Results for Macaque V1 and LGN

Peak temporal frequencies and high frequency cutoffs for the LGN and V1 cells. Values are means ± standard deviation, in Hz. High temporal frequency cutoff is the frequency where the response reaches half of the maximal response. Input layers are layers 4a, 4cα and 4cβ. Output layers are layers 2, 3, 4b, 5 and 6.

Foster et al. [1985] found that macaque V1 cells were tuned to even lower temporal frequencies. They found an average optimal temporal frequency of 3.7Hz and a median high frequency cutoff (the frequency of the half-maximal response) around 8Hz. In the discussion section of Hawken et al. [1996] the authors speculate that a possible factor in the discrepancy between their results and those of Foster et al. [1985] was the differences in the anesthetic agents used in the two protocols. It is worth noting that Hawken et al. [1996] used ketamine, an NMDA antagonist, during initial surgery, while Foster et al. [1985] did not. It is unclear if ketamine has lingering effects on NMDA that would last until the time of recording, well after the initial surgery, but it is another possible distinction between their protocols.

A recent study of physiological properties of LGN and visual cortex in the New World owl monkey showed a similar shift in temporal frequency tuning from the LGN to the cortex [O'Keefe et al., 1998]. While both LGN and V1 cells responded over a lower range of frequencies in owl monkey compared to macaque, the overall drop in peak temporal frequency was similar. Other than peak frequencies, O'Keefe et al. [1998] unfortunately report different properties of the tuning curves for LGN and cortical cells, which makes direct comparison difficult. For LGN cells they report the high frequency cutoff which is the frequency at which response reaches half of the maximum, and something called the transience which is the ratio of the low-frequency response (at 1/10 of the peak frequency) to the peak response. For cortical cells they report the temporal resolution, which is the frequency at which the response is almost gone (1 spike/sec) and the bandwidth, which is the ratio of the high to low frequency cutoffs in octaves. However, the fact that the median temporal resolution for the cortical cells, a much stricter high frequency limit, is

lower than the high frequency cutoff for the LGN cells (table 2.8) attests to the significant low pass shift from the LGN to the cortex. There are also laminar differences in temporal frequency tuning, similar to what is seen in macaques. The tuning of layer 4 cells is intermediate between the tuning of the LGN and the supergranular layers, which suggests that there might be multiple levels of low pass shifting in primate V1.

	Peak Frequency	High Frequency Cutoff	Temporal Resolution	Bandwidth (Octaves)
LGN Parvo	6.4 (20)	22.6 (21)		
LGN Magno	4.0 (22)	16.1 (23)		
V1 (total)	3.0 (185)		10.3 (195)	3.8 (185)
V1 (simple)	3.4 (87)		10.6 (92)	3.7 (92)
V1 (complex)	2.7 (98)		10.3 (103)	3.7 (104)
V1 (layer 2/3)	2.3 (83)		7.9 (89)	3.8 (83)
V1 (layer 4)	4.2 (50)		14.8 (51)	3.8 (50)
V1 (layer 5/6)	4.6 (16)		14.1 (18)	4.1 (16)

Table 2.8:

O'Keefe et al. [1998] Temporal Frequency Tuning Results for Owl Monkey V1 and LGN

All values are medians. Values for peak frequencies, high frequency cutoffs and temporal resolution are in Hz. High temporal frequency cutoff is the frequency where the response reaches half of the maximal response. High frequency resolution is the frequency where the response reaches 1 spike/sec. Bandwidth is the ratio of the high frequency cutoff to the low frequency cutoff in octaves.

Comparison to temporal tuning in auditory cortex and medial geniculate body (MGB)

A similar low pass shift in the temporal tuning of cortical neurons as compared to their thalamic inputs has been observed in the auditory system. The temporal properties of primary auditory cortex (A1) and of the auditory region of the thalamus that supplies input to A1, the medial geniculate body (MGB), has been probed with a number of different types of time-varying stimuli. A common method is to use a carrier signal, such as a pure tone or a noise, whose amplitude modulates in time. As with responses to modulating visual stimuli, the responses to amplitude modulating auditory signals is characterized either by the first fourier component of the response at the modulation frequency, which is considered a measure of the degree to which the cell is able to phase-lock to the stimulus, or by the mean firing rate. Another measure is the average

number of action potentials locked to each cycle of modulation, which is refered to as the degree of entrainment. Alternative methods include measuring the responses to frequency modulating stimuli, and to click trains with different rate of repitions of the clicks.

There have been a number of studies that have looked at the temporal tuning of populations of MGB and A1 cells with a variety of methods, and in various species (see Schreiner [1999] for a review). While there are a wide range of optimal frequencies reported in both populations, there is a general trend for cortical cells to prefer lower frequencies. Distributions of optimal frequencies in the MGB range from a few hertz to over 100Hz, with means of these distributions occuring around 35Hz. For cortical cells, very few neurons have optimal frequencies above 50Hz. The means of different reported distributions ranged from 6.8Hz to 20.6Hz, with a grand mean of 13.6Hz.

More direct evidence for shifts in temporal tuning come from a simultaneous study of the tuning of thalamic and cortical cells [Creutzfeldt et al., 1980]. They measured the response of MGB neurons and A1 neurons in unanesthetized guinea pigs to amplitude modulated tones with different frequencies of modulation, and recorded the peak amplitude of the response at steady-state, measured 0.5 to 1 second after the onset of the stimulus. When plotting this peak amplitude for populations of cells as a ratio to the onset response to a 10Hz stimulus, as a function of modulation frequency, they find that the set of ten thalamic cells reach half of their maximal response at 50 Hz, while twelve cortical cells reach half maximal response at 25Hz. Furthermore, they were able to find pairs of thalamic and cortical cells that had cross-correlations that were suggestive of monosynaptic connections. In one example of such a pair, the thalamic cell was able to follow a modulating stimulus up to 100Hz, although the amplitude started to decrease after 60Hz. The postsynaptic cortical cell, however, was not able to follow the stimulus beyond 20Hz. For a set of four cortical cells that were recorded with a monosynaptically connected thalamic cell, the peak amplitude fell to half maximal response at 15Hz, even more sharply than the total set of 12 cortical cells.

The developmental trends of temporal frequency tuning shifting towards higher frequencies with age that were described above for the visual cortex have also been demonstrated in the auditory cortex. Eggermont [1993] found that optimal frequencies in response to both click trains and amplitude modulated stimuli in anesthetized cats increased from ≈4Hz in kittens younger than 30 days to ≈10Hz in adult cats. The developmental changes stabilized by 4 to 6 months of age, and the time constant of this change was 19.6 days.

Comparison to somatosensory cortex and thalamus

There have been studies in different species of the ability of thalamic and cortical cells to follow repetitive stimuli sensory stimuli at different repition frequesncies. The discussion in Simons [1985] points out that somatosensory cortical neurons can not follow repetitive stimuli greater than 40-80Hz while the thalamocortical fibers to the cortex are entrainable to at least 200Hz. He refers to

the following articles, that I will not describe in detail here: [Ferrington and Rowe, 1985, Hellweg et al., 1977, Mountcastle et al., 1969, Simons, 1978].

Summary

There is a wide range of evidence for shifts in temporal tuning towards lower frequencies from the thalamus to the cortex, in various species and in different sensory modalities. While there is significant variability in different measures of the degree of these shifts, the existence of at least some degree of shift in tuning seems fairly universal. On exception, perhaps, may be a comparison of the tuning of cortical cells with the tuning of the lagged cells reported by Saul and Humphrey [1990], however, they acknowledge the important role of non-lagged input to the cortex, and they do not argue that cortical temporal tuning is derived purely from the lagged cell input. There is also a developmental trend in temporal frequency tuning, such that the youngest kittens are tuned to the lowest frequenies.

2.2.2 NMDA Receptor Contributions to Thalamocortical and Intracortical Excitatory Transmission

Terminology

To make reading more transparent, we shall use "NMDA" to mean "NMDA-receptor-mediated" in expressions such as "NMDA component", "NMDA conductance", "NMDA connection", "NMDA contribution", "NMDA current", and "NMDA responses". We shall similarly use "AMPA" to mean "AMPA-receptor-mediated" in similar contexts.

Pharmacological studies in visual cortex *in vivo*

While much of the focus on NMDA receptors has been on their role in the developing central nervous system and in mediating synaptic plasticity, there is also substantial experimental evidence that NMDA receptors play a role in normal excitatory transmission in the adult animal. In fact, some have argued that many of the studies that have demonstrated a role for NMDA receptors in cortical plasticity in visual cortex have been confounded by their role in synaptic transmission and thus in determining the activity level of cortical cells. Recently, however, Roberts et al. [1998] were able to reduce NMDA receptor efficacy enough to have an effect on visual cortical plasticity without having an effect on the response of cells to visual stimuli, by using antisense DNA complementary to the mRNA that codes for the NR1 components of the NMDA receptor, rather than a more standard pharmacological agent.

Many studies have used local iontophoretic infusion of APV, a selective NMDA-receptor blocker, to study the effect on visual responses in primary visual cortex. Fox et al. [1989] found a strong reduction of visual responses in the superficial layers for younger animals. In addition, spontaneous

activity was reduced in all layers. They later elaborated on this result, to describe more precisely the effect on visual responses [Fox et al., 1990]. APV was found to reduce the gain of the contrast-response curves of cells in the superficial layers, without affecting the threshold contrast required to see the smallest response. In layers IV, V and VI, application of APV reduced spontaneous activity with little or no effect on gain. Furthermore, application of low levels of NMDA increased the gain in all layers except for layer IV, while application of quisqualate (an AMPA-receptor agonist) increased the level of firing without affecting gain.

Hagihara et al. [1988] found slightly different results. With iontophoretic application of APV, they found suppression of visual response of some cells in all layers. Overall, 34% of the cells they encountered were APV-sensitive. The proportion was the largest in layer V, where 60% of the cells were suppressed, and lowest in layers IV and VI.

Iontophoretic application of drugs has very local effects, and it is possible that in some layers it does not sufficiently cover the entire dendritic field of cells to suppress visual responsivity. It is also possible that a cell depends on NMDA receptors indirectly, because its responses depend on excitation from other cells whose activation in turn depends on NMDA receptors. To study the effects of globally blocking NMDA receptors throughout a small region of cortex, Miller et al. [1989] infused 50mM of the NMDA receptor antagonist APV into the visual cortex of adult cats for one day. In the region surrounding the cannula they found blockade to iontophoretic application of NMDA, and not to application of agonists specific for non-NMDA receptors, as well as profound suppression of response to visual stimuli. The degree of suppression of the visual response correlated well with the degree of NMDA receptor blockade. They argued that this was strong evidence for the NMDA receptor playing a crucial role in normal cortical activity in the adult.

Kasamatsu et al. [1998] found similar results. They infused a lower dosage of APV (10mM) into adult cats and measured visual reponses. Under two different conditions, either infusing continuously throughout the recording session for a total of 33-48 hours, or infusing for 10 hours and recording subsequently after waiting either 34 or 51 hours, they found a high percentage of unresponsive cells. (Among the cells that did respond, the normal pattern of binocularity had been disrupted by the drug application.) Bear et al. [1990], however, reported little or no effect on the strength of visual responses of visual cortical cells after 1 week of 50mM APV infusion in two adult cats. It is possible that there is some form of cortical compensation for NMDA receptor blockade over this longer time, so this result is not necessarily inconsistent with the other two reports on APV infusion and need not indicate a lack of dependence of normal cortical responses on NMDA receptors.

Studies in other cortical regions and other species

Some *in vivo* studies have also looked for NMDA contributions to the excitatory input from the thalamus to the cortex. Salt et al. [1995] recorded intracellularly in cells in cat motor cortex, and

found that the EPSPs evoked from thalamic inputs were insensitive to iontophoretic application of NMDA receptor antagonists, but were blocked by application of CNQX, an AMPA-receptor antagonist. EPSPs evoked by intracortical stimulation, on the other hand, were sensitive to both types of antagonists.

There has also been an intracellular study done in an *in vitro* preparation of turtle visual cortex that preserves the geniculocortical pathways [Larson-Prior et al., 1991]. They also found that the EPSPs evoked by stimulation of the geniculocortical fibers were insensitive to APV. Intracortical stimulation evoked EPSPs with three components separated by their latency to the peak response. The slowest of the three was sensitive to APV.

Armstrong-James et al. [1993] looked at the spike responses of neurons in layers I-IV of rat barrel cortex to whisker stimulation during iontophoretic application of APV and DNQX. They found that most of the spike responses were sensitive to NMDA receptor antagonists. However, the shortest latency spikes were insensitive to APV, and they concluded that the thalamocortical synaptic responses were not mediated by NMDA receptors.

In vitro evidence for NMDA conductances in intracortical synapses

Many results from Alex Thomson's lab [Thomson and Deuchars, 1994, Thomson and West, 1993, Thomson et al., 1993], discussed below, have suggested a significant NMDA component in intracortical connections between pyramidal cells. More recently, Markram et al. [1997] recorded from pairs of adjacent layer V pyramidal cells in slices of young rat somatosensory cortex to measure the physiological properties of synaptic connections between cortical cells that were due to a single presynaptic cell's axonal connections to another cortical cell. They measured EPSPs at various baseline potentials, both with and without APV, and looked at the fraction of the EPSP amplitude and integrated voltage that was suppressed by APV, as a function of the baseline voltage. The fraction of the EPSP that was sensitive to APV had a voltage dependence that was typical for NMDA conductances. By this technique, they concluded that at -80mV the EPSP had a negligible NMDA component, at -60mV the NMDA component made up 17% of the integrated potential, at -50mV this went up to 26% and at -30 mV it was 55%.

The same lab performed a similar set of experiments in pairs of layer 4 cells [Feldmeyer et al., 1999]. They found that at -60mV, with 1mM $Mg^{+}+$ present, NMDA receptors contributed $39.3 \pm 12.5\%$ of the integrated EPSP integral. In $Mg^{+}+$ free solution, the NMDA/AMPA receptor ratio of the EPSC was 0.86 ± 0.64.

Another group has recorded from layer 4 cells in a tangential slice preparation that isolates layer 4, and preserves the connections within layer 4 [Fleidervish et al., 1998]. They report that these intra-layer 4 connections are largely mediated by NMDA, and that there are two different types of NMDA, that are distinguished by the degree of voltage dependence.

In vitro evidence for NMDA conductances in thalamocortical slice

Crair and Malenka [1995] looked for the presence of NMDA conductances in thalamocortical synapses during development. Using a thalamocortical slice preparation of rat somatosensory cortex, they stimulated in the thalamus and recorded EPSPs intracellularly in layer IV cells. These were determined to be monosynaptic by the fact that their latency was independent of stimulation strength and frequency. They then recorded EPSCs under voltage clamp, holding the cell both at -90mV, where NMDA conductances are under voltage block, and at +40mV in the presence of CNQX, to block non-NMDA conductances. They found the presence of NMDA conductances in slices from ages 3 to 14 postnatal days, but the ratio of the peak of the NMDA current to the peak AMPA current went down with age. For slices aged P3-7, the ratio was 1.62 $\pm$ 0.9, while for slices aged P8-14, the ratio was 0.30 $\pm$ 0.2. Furthermore, the time course of the decay of the NMDA currents increased with age, from 248 $\pm$ 62msec for P3-7 cells, to 113 $\pm$ 26 msec for P8-14.

Gil and Amitai [1996] provided strong evidence for the presence of NMDA conductances in adult thalamocortical connections. They used a thalamocortical slice preparation of somatosensory cortex from adult mice. The first set of experiments were performed in a Mg^{2+}-free solution to enhance NMDA responses. In these conditions, electrical stimulation of either the thalamus or intracortically in layer 6 evoked large synchronized population events, revealed both in field potential recordings and as large depolarizing events recorded intracellularly in layer 5 cells. Application of CNQX, to block non-NMDA glutamate receptors, eliminated these events if triggered by thalamic stimulation, but not those triggered by intracortical stimulation. That is, NMDA connections alone were sufficient to propagate these large population events, following the cortical stimulation. With CNQX application, monosynaptic EPSPs, albeit reduced in size, could still be observed in layer 5 intracellular recordings following thalamic stimulation, providing strong evidence for the presence of NMDA receptors in these direct thalamocortical synapses. Furthermore, focal application of the Gaba-A antagonist, bicuculline, in layer 4, in the presence of CNQX reestablished the population events triggered by thalamic stimulation. This result was a good indication that some excitatory synapses onto inhibitory interneurons also have a significant NMDA component.

In the second set of experiments, a 2mM Mg^{2+} solution was used to study NMDA conductances under more physiological conditions. With bath application of CNQX, intracellular recording in layer 5 cells again revealed short latency EPSPs suggestive of monosynaptic connections following either thalamic or intracortical stimulation. These EPSPs had voltage sensitivites suggestive of NMDA conductances, and were eliminated by application of APV.

Developmental changes in NMDA conductances

Aside from the study by Crair and Malenka [1995] described above, there have been a few other studies that have looked at developmental changes in NMDA conductances. Carmignoto and Vicini [1992] measured the time course of NMDA EPSCs in layer IV spiny stellate and pyramidal cells

from rat visual cortical slices of various ages, evoked by stimulation of the borders between the white matter and layer VI. While Crair and Malenka [1995] found that the NMDA EPSC decay was well fit by a single exponential whose time constant increased with development, Carmignoto and Vicini [1992] found that, for almost all cells at all ages, the decay was best fit by double exponentials with a fast and a slow component. (For some of the cells from the youngest slices, the cells were well described by a single exponential, essentially corresponding to a pure slow component.) The time constants of the two components were relatively stable over the course of development (approximately 200 msec for the slow component and 63 msec for the fast component). What distinguished the cells of different ages was the relative strength of the slow component. For slices 9 to 14 days old, before eye opening, the mean contribution of the slow component was 92 ± 12% (n=18). For adult slices, the contribution of the slow component was never greater than 35%, and had a mean of approximately 12%. Furthermore, they found that dark-rearing rats prolonged the period in which NMDA EPSCs were dominated by the slower decay component.

Differences in NMDA conductances onto excitatory and inhibitory cells

Ling and Benardo [1995] argue that inhibitory interneurons in the cortex, unlike those in the hippocampus, do not have a significant NMDA component in the excitation they receive. They recorded intracellularly in layer 5 pyramidal cells in slices from mature rat somatosensory cortex. Electrically evoked EPSCs and IPSCs were isolated by holding the cell at the inhibitory (-75mV) and excitatory (0mV) synaptic reversal potentials respectively. Increasing stimulation intensity of shocks applied in cortical layer 6, evoked larger and larger EPSC amplitudes, while the IPSC amplitudes saturated at stimulus intensities of less than double the threshold intensity. The IPSC amplitude at saturation did not depend on manipulations that increased (0 Mg^{2+}) or reduced (application of CPP, an NMDA-receptor antagonist) the efficacy of NMDA conductances; however, the stimulus intensity required to reach saturation did change. Both the maximal IPSC and EPSC amplitudes were significantly reduced by application of CNQX, but the residual EPSC amplitude was eliminated by subsequent application of CPP, while the residual IPSC was unaffected by either CPP or removal of Mg^{2+}. This implied that the CNQX-resistant component of the IPSC was due to monosynaptic connections from inhibitory interneurons that were close enough to the stimulating electrode to be directly activated. Moving the stimulating electrode to the deep white matter eliminated the CNQX-resistant component of the IPSC, even without Mg^{2+}. Furthermore, intracellular recording from fast-spiking, putative inhibitory layer V/VI cells showed that electrically evoked EPSPs were mostly eliminated by CNQX.

A more direct method for measuring the properties of cortical synapses is with paired intracellular recordings that isolate the postsynaptic potentials that are evoked by the axon from a single presynaptic cell. Alex Thomson's laboratory has been performing such experiments for a number of years, and has demonstrated some differences in the properties of EPSPs onto excitatory

pyramidal and inhibitory non-pyramidal cells [Thomson and Deuchars, 1994, Thomson and West, 1993, Thomson et al., 1993]. Although pharmacological studies were not performed, they did measure the voltage dependence of these single-axon EPSPs. Pyramidal-to-pyramidal EPSPs displayed the non-linear voltage dependence characteristic of having a significant NMDA component, while pyramidal to non-pyarmidal EPSPs had a linear voltage dependence.

More recently, Angulo et al. [1999] used pharmacology to directly study the relative contribution of NMDA and AMPA receptors in synaptic conductances between presynaptic pyramidal and postsynaptic fast-spiking interneurons, using paired intracellular recordings in slices of rat cortex. Their findings indicated a relatively small NMDA component in these synapses. Holding the membrane potential near rest, at -72 mV, application of AMPA-receptor antagonists eliminated almost all (98%) of the evoked EPSC. At higher holding potentials, a small NMDA component of the EPSC was revealed. From these measurements, they calculated a ratio of the magnitude of the NMDA conductance at 40mV to the magnitude of the AMPA conductance at -60mV of 19/

Perouansky and Yaari [1993] studied NMDA EPSC's in excitatory and inhibitory cells in rat hippocampal slices. The overall size was similar in the two sets of cells, but the kinetic properties were different. While the rise times were similar, both decay times, fit by double exponentials, were faster in the inhibitory cells (34 and 212 msec, vs. 66 and 354 msec in excitatory cells). However, the component of the decay mediated by the faster rate was higher in excitatory cells than in inhibitory cells. They also found a category of inhibitory cells that had no fast component of NMDA EPSC decay at all.

Summary

A number of studies have demonstrated the presence of NMDA-mediated excitatory conductances, both in intracortical and thalamocortical connections. There is some evidence that the presence of NMDA conductances is more prominent onto excitatory cortical cells than onto inhibitory cortical cells. Furthermore, it has been demonstrated that the decay time constant of NMDA conductances decrease, and the relative proportion of NMDA to AMPA conductances decrease over the course of development.

2.3 Methods

The model used in this and the following chapters is in almost all essential details identical to the "computational" model of the first chapter. Therefore, we will only describe here the differences from that model. There are a few essential additions that will be described in detail, and minor parameter changes will simply be listed in the "Default Parameter Sets" section below.

2.3.1 LGN inputs

The model of LGN inputs was almost identical to the model from the first chapter, only substituting a different set of experimental data [Sclar, 1987] for the contrast response curves used previously [Cheng et al., 1995]. Because X cells behave quite linearly, we were able to plausibly approximate LGN firing rates in response to sinusoidal moving gratings as the rectified sum of two factors: (1) a sinusoidal modulation with the same temporal frequency as the stimulus, and with amplitude determined as described below; and (2) background firing rates of 15 Hz and 10 Hz for ON and OFF cells, respectively [Levine and Troy, 1986, Kaplan et al., 1987]. By rectified sum, we mean that negative values of the sum were set equal to zero.

The size of the sinusoidal modulation can be described in terms of the first harmonic (F1) of the LGN firing rates either before or after rectification (the F1 is the amplitude of the response component at the same temporal frequency as the stimulus). For most of the simulations presented in this paper, the pre-rectification F1 – the amplitude of the sinusoidal modulation described in the previous paragraph – in response to a given stimulus of the preferred spatial frequency was chosen separately for ON and OFF cells (because of their different background firing rates) so that the post-rectification F1 – the F1 of the geniculate response after rectification – matched the F1 values determined by Sclar [1987] (figure 1) for LGN responses to stimuli of different contrasts and temporal frequencies. Sclar [1987] report the amplitude of modulation of the response of a typical LGN X-cell at multiple contrasts (80,40,20 and 10%) and multiple temporal frequencies (0.5, 1, 2, 4, 8, 16, 32 Hz). For intermediate levels of contrast or temporal frequency, matlab's 2-D cubic interpolation (interp2) was used to find the best fit modulation amplitude. The pre-rectification F1s were further modified for stimuli of nonoptimal spatial frequencies by use of the following center-surround LGN spatial filter [Peichl and Wassle, 1979, Linsenmeier et al., 1982]:

$$17e^{-k^2\sigma^2_{\text{center}}/4} - 16e^{-k^2\sigma^2_{\text{surround}}/4}; \quad \sigma_{\text{center}} = 15', \sigma_{\text{surround}} = 1^o. \tag{2.1}$$

Here, k is the spatial frequency of the grating; its wavelength λ is given by $\lambda = 2\pi/|k|$. All gratings were shown at 0.8 deg/cycle, so the modulation amplitude was reduced by the amount predicted by this filter (0.86) relative to its value at the preferred spatial frequency of the LGN (0.54 cyc/deg, the k for which the filter value is maximized). For the figures where the LGN input temporal frequency tuning curves are compared to the temporal frequency tuning curves generated from the mean firing rates of the excitatory cortical cells in the model, we calculated the LGN mean rates as the time average of the rectified sum of the pre-rectified sinusoid and the background firing rate for ON and OFF cells separately, and averaged the ON and OFF results together.

2.3.2 Thalamocortical Receptive Fields

The thalamocotical receptive fields used throughout this and the following chapters are the more broadly tuned, second set of Gabor parameters described in the first chapter. The Gabor function

is determined with three parameters: (1) the spatial frequency of the sinusoid, which was chosen to be 0.8 cycles/degree, corresponding to a half-cycle width of 0.625^o, (2) the number of subregions, defined as the ratio of the width of the Gaussian envelope (at 5% of peak) to the half-cycle width of the sinusoid, and (3) the aspect ratio of a single subfield, defined as the ratio of the Gaussian envelope length to the sinusoid half-cycle width. The "broadly-tuned" Gabor parameters from the first chapter were defined as the mean values for simple cell physiological RFs reported in Jones and Palmer [1987] (2.65 subregions and an aspect ratio of 4.54) compressed by a factor of 0.7, yielding 1.85 subregions and an aspect ratio of 3.18. They were chosen originally to approximate the tuning widths ($\approx 35^o$ half-width at half height) of the modulation of the membrane potential due to thalamocortical synaptic inputs alone that were suggested by the cooling experiments of Ferster et al. [1996]. These parameters are now more firmly supported as the appropriate choice by quantitative measurements of the tuning width of the modulation of the membrane potential, not under cooled conditions but due to the full cortical circuit, by Carandini and Ferster [2000]. They report a distribution of half-widths at half height with a mean of 38^o, which closely match the values that theres parameters achieve in the model. This will be discussed in more detail in the following chapter.

2.3.3 Model of NMDA conductances

While the other time varying conductances are modeled as a simple difference of single exponentials, as described in the previous chapter, we modeled the decay of the NMDA conductance as a double exponential with a fast and a slow time constant:

$$g_{\mathrm{NMDA}}(t) = \sum_{t_j<t} \bar{g}_{\mathrm{NMDA}}(V_{\mathrm{shadow}}) \left(f_{\mathrm{fast}} e^{-(t-t_j)/\tau^{\mathrm{fall}}_{\mathrm{NMDA,fast}}} + (1 - f_{\mathrm{fast}}) e^{-(t-t_j)/\tau^{\mathrm{fall}}_{\mathrm{NMDA,slow}}} - e^{-(t-t_j)/\tau^{\mathrm{rise}}_{\mathrm{NMDA}}} \right) \tag{2.2}$$

where the sum is over presynaptic spike times t_j, and f_{fast} represents the contribution of the faster exponential to the total decay term. V_{shadow} is essentially the membrane potential without the pasted on spikes. That is, while a cell is in its refractory period, the actual voltage of the cell is held at V_{reset}. V_{shadow} is never reset and evolves continuously as if unaware of the spike. We chose to use V_{shadow} as the voltage that the NMDA conductances are aware of, to avoid discontinuities in synaptic currents during spiking events. Furthermore, since NMDA conductances are on the dendrites, it is appropriate to link them to a voltage that reflects the net synaptic input, and is not limited by spike threshold. Parameters were taken from the data for adult rats in a developmental study of NMDA conductances in the rat visual cortex [Carmignoto and Vicini, 1992]: $\tau^{\mathrm{fall}}_{\mathrm{NMDA,fast}} = 63\ msec$, $\tau^{\mathrm{fall}}_{\mathrm{NMDA,slow}} = 200\ msec$, $f_{\mathrm{fast}} = 88\%$. We chose $\tau^{\mathrm{rise}}_{\mathrm{NMDA}} = 5.5 msec$ to set the 10-90% rise time of the NMDA EPSC to be equal to $7.8 msec$ as has been observed experimentally [Lester et al., 1990]. For the model of temporal frequency tuning of younger animals, we held the time constants fixed, and only adjusted the contribution of the faster time constant, as suggested by [Carmignoto and Vicini,

1992] ($f_{\text{fast,young}} = 10\%$). We did not adjust $\tau^{\text{rise}}_{\text{NMDA}}$ for this case, as it only increased the 10-90% rise time slightly ($9.3msec$). The voltage dependence of $\bar{g}_{\text{NMDA}}$ followed the model described in Jahr and Stevens [1990]:

$$\bar{g}_{\text{NMDA}}(V) = 1/\frac{1 + (a_1(V) + a_2(V,C))(a_1(V)B_1 + a_2(V,C)B_2)}{[Aa_1(V)((b_1(V) + B_1) + Aa_2(V,C)(b_2(V) + B_2)]} \tag{2.3}$$

where the lower case parameters have voltage dependence, V in units of mV, and a_2 also depends on magnesium concentration, C in units of uM. Parameters are as follows: $a_1(V) = e^{(-0.016V-2.91)}msec^{-1}$, $a_1(V,C) = Ce^{(-0.045V-6.97)}uM^{-1}msec^{-1}$, $b_1(V) = e^{(0.009V+1.22)}msec^{-1}$, $b_2(V) = e^{(0.017V+0.96)}msec^{-1}$, $A = e^{(-2.847)}msec^{-1}$, $B_1 = e^{(-0.693)}msec^{-1}$, $B_2 = e^{(-3.101)}msec^{-1}$, $C = 100uM$. With this model of voltage dependence, the NMDA channels are still 35.5% open at the value for threshold used in the model, $-52.5mV$ (Figure 2.2). In response to a high contrast optimal grating, V_{shadow} reaches a peak of approximately -46 mV, where the NMDA channels are 43.8% open.

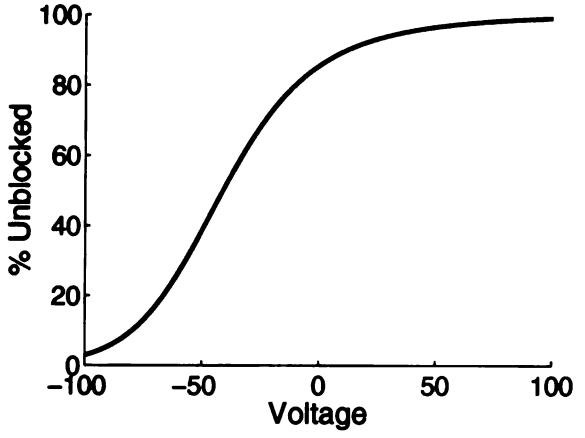

Figure 2.2:
NMDA Voltage Dependence using the model from Jahr and Stevens [1990]. Percentage of Model NMDA channels open as a function of membrane potential.

Figure 2.3 compares two measures of the relative strength of NMDA and AMPA conductances. One, the NMDA/AMPA amplitude ratio (horizontal axis of Fig. 2.3) was used by Crair and Malenka [1995]. This is the ratio of the amplitude of NMDA EPSC's with the cell held at +40mV to the amplitude of AMPA EPSC's with the cell held at -90mV. The other, the %NMDA integrated current (vertical axis of Fig. 2.3), is the percent of the temporally-integrated current (*i.e.*, of the total charge transfer) through excitatory conductances that is mediated by NMDA conductances, when the postsynaptic cell is clamped at the spike-threshold voltage. As can be seen, even for relatively modest NMDA/AMPA amplitude ratios, the integrated current is dominated by NMDA. This is because of the fact that NMDA has a much longer time constant; a small amplitude NMDA EPSC will still contribute a significant amount of integrated current as compared to a short, high amplitude AMPA EPSC.

Since the main parameters we use to set synaptic strengths is the total integrated current, as discussed in the previous chapter, we will use the %NMDA integrated current as our main

description of the relative strengths of NMDA and AMPA conductances in the various types of excitatory synapses (thalamocortical to excitatory cells, thalamocortical to inhibitory cells, and intracortical). Crair and Malenka [1995] report an NMDA/AMPA amplitude of 0.30 for the oldest thalamocortical slices that they studied (8 to 14 postnatal days) which corresponds to 91.2% of the integrated current mediated by NMDA. We have used 90% NMDA as our default for full strength of NMDA in thalamocortical synapses. For their youngest slices (3 to 7 postnatal days) they measured an NMDA/AMPA amplitude ratio of 1.62 which corresponds to a 98.2% NMDA component of the integrated current, which we will also use in our test of the development of temporal frequency tuning.

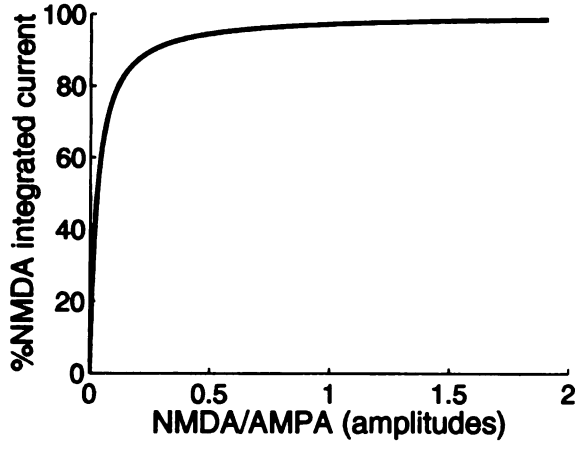

Figure 2.3:
Comparison of two measures of relative strengths of NMDA and AMPA. X-axis is the ratio of the amplitude of NMDA EPSC with the cell held at +40mV to the amplitude of AMPA EPSC with the cell held at -90mV (the measure used by Crair and Malenka [1995]. Y-axis is the percent of the temporally-integrated current (*i.e.*, of the total charge transfer) through excitatory conductances that is mediated by NMDA conductances, when the postsynaptic cell is clamped at the spike-threshold voltage.

2.3.4 Synaptic Depression

In recent years there has been a great deal of attention focused on experimental measurements and quantitative models of the short term plasticity of cortical synapses. It has been observed that synaptic strengths weaken with repeated use, and then recover to full strength after a period of quiescence. The time course of this process has been described with a model that involves only two parameters: f $(0 \leq f \leq 1)$ the factor by which the strength of a synapse gets scaled immediately after a spike, and τ, the time constant of an exponential decay back to full strength [Tsodyks and Markram, 1997, Abbott et al., 1997].

In practical computational terms, this form of depression is implemented by storing, for each presynaptic cell, the most recent time of spiking and the most recent synaptic strength scaling factor, $A \leq 1$. Synaptic scaling factors then only have to updated at the moment of each presynaptic spike,

using the iterative equation:

$$A_i = 1 - (1 - A_{i-1}f)e^{\frac{-(t_i - t_{i-1})}{\tau}} \tag{2.4}$$

where A_i is the current scaling factor, A_{i-1} is the scaling factor at the time of the most recent presynaptic spike, t_i is the current spike time and t_{i-1} is the time of the most recent spike. We used parameters determined from paired-pulse experiments that were performed on putative thalamocortical synapses onto cells in layer 4 ($\tau_{TC} = 99msec, f_{TC} = 0.563$; [Stratford et al., 1996]) and intracortical synapses within layer 4 ($\tau_{IC} = 57msec, f_{IC} = 0.875$; [Tarczy-Hornoch, 1996]; see Kayser et al. [1999] for details on parameters found in the literature using different experimental protocols.) For all simulations using synaptic depression but no cortical amplifier, depression was implemented for the thalamocortical synapses onto both excitatory and inhibitory cells using the parameters τ_{TC} and f_{TC}. For all simulations using synaptic depression and a cortical amplifier, depression was implemented for the thalamocortical synapses as above and in the intracortical E⇒E synapses using the parameters τ_{IC} and f_{IC}. Synaptic depression was not included in the intracortical I⇒E synapses, nor in the E⇒I or I⇒I connections when these connections were present.

2.3.5 Fixing Parameters to Achieve Appropriate Background Firing Rates

As in Troyer et al. [1998], we included non-specific excitatory input to all of the cortical cells in the model, excitatory and inhibitory, to bring the cells close to threshold so that they have appropriate background firing rates, that is, approximately 0.5Hz for excitatory cells and 20-30Hz for inhibitory cells. The cells also receive input from the LGN inputs, which have background firing rates set at 15Hz for OFF cells and 10Hz for ON cells. Thus, either increasing the strength of the non-specific input or the strength of the thalamocortical input will raise the background firing of the cortical cells. Furthermore, the degree of noise in (*i.e.* the variance of) the non-specific input affects background firing rates. If the mean level of non-specific input is kept constant, but the degree of noise is increased, there will be more frequent threshold crossing, and therefore higher firing rates.

There are, therefore, three parameters that were adjusted to achieve appropriate background firing rates for both the excitatory and inhibitory cells in the model. One parameter is the total strength of the thalamocortical synapses. The other two parameters specify the non-specific input, which is modeled as a presynaptic Poisson spike train with a purely AMPA mediated excitatory synapse. This input is parameterized by the amplitude of the conductance for a single presynaptic spike (Ks where s is a unitary conductance of amplitude .89 nS and K is a unitless scaling factor) and the rate of the Poisson spike train (r in Hz). A model AMPA conductance of amplitude 1 nS produces a time-integrated conductance of $t = 1.5$ nS msec. The mean conductance M produced by this input is the product of the rate times the time-integrated conductance produced by a single presynaptic spike ($M = Kstr$). Increasing the strength of an individual synaptic conductance, but keeping the mean conductance constant by reducing the rate of the Poisson spike train, increases the variance in this conductance, *i.e.* the noise of the non-specific input. Increasing the noise increases

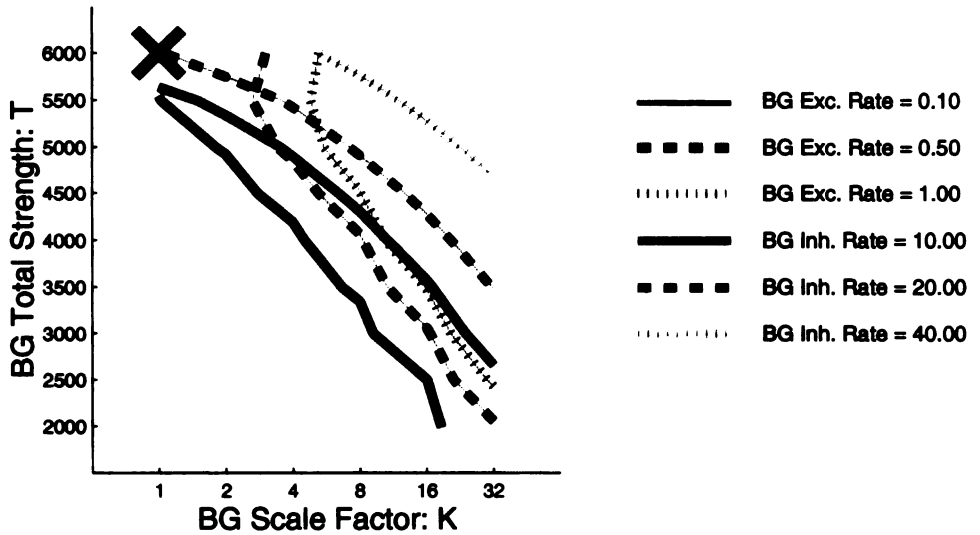

Figure 2.4:

Contour plots of background firing rates for excitatory and inhibitory cortical cells, varying the two primary parameters defining the non-specific Poisson input onto all cortical cells, with no synaptic depression. T is proportional to the total mean conductance produced by the non-specific input, and K is proportional to the amplitude of the conductance of a single presynaptic spike. $T = Kr$ where r is the rate of the Poisson spike train. Thalamocortical NMDA onto E cells = 90%. TC NMDA onto I cells = 0%. Intracortical NMDA = 95%. Each contour illustrates where in parameter space the background rate of excitatory cells (blue) and inhibitory cells (red) have a particular value, as shown in the key. Large black X indicates the default parameters for no depression used in this and the following chapters: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec. Other parameters are: Total inhibition = 5 nA msec; intracortical excitation = 3 nA msec.

the postsynaptic spike rate: for example, if the total input is packed into a very low rate Poisson spike train with massive conductances at each spike, then the postsynaptic cell is very likely to cross threshold with every presynaptic spike. Therefore, we have parameterized this input with a measure proportional to the total strength, $T = Kr$, and the parameter K, which is a measure of the noise.

These parameters were adjusted separately without and with synaptic depression (figures 2.4 and 2.5, respectively). In general, increasing either T or K increases both the excitatory and inhibitory background firing rates. However, eventually, as inhibitory background rates get high, this trend reverses for the excitatory cells. As either T or K increases, the increase in background inhibition received by the excitatory cells dominates over the increase in either the mean or the noise level of the non-specific input. The total thalamocortical strength needed to be increased

significantly when including synaptic depression, since the thalamocortical synapses at background are in a somewhat depressed state, due to the background firing rate of the LGN cells. Picking default parameter sets for these three parameters was an imprecise and somewhat arbitrary process. There was not a unique solution varying all three parameters simultaneously to fit only the two background firing rates. The contour plots for varying levels of thalamocortical strengths look similar to those of figures 2.4 and 2.5, but are translated slightly as the thalamocortical strengths are varied. Furthermore, the background rates were also dependent on other parameters that we vary in our modeling study, so background firing rates are not completely fixed through our other parameter searches. Background rates are clearly dependent to some degree on total intracortical excitatory and inhibitory strengths. But, they are also dependent on the proportion of NMDA in the thalamocortical synapses onto excitatory and inhibitory cells, even when keeping the total thalamocortical excitatory strength constant; higher proportions of NMDA lead to less noise in the thalamocortical input, and somewhat less background firing, in general. Our default parameter sets are: without depression, $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec; with depression, $T = 4000$, $K = 8$, thalamocortical excitation = 20 nA msec. The background inhibitory rates were higher with the default parameters with depression ($\approx$ 30 Hz) than with the default parameters without depression ($\approx$ 20 Hz).

2.3.6 Adaptation Conductance

The only other parameter we have changed from the parameters of chapter one is the amplitude of the adaptation conductance, $\bar{g}_{\text{adapt}}$, which was reduced by a factor of 5 from a value of 3 nS to 0.6 nS. The time constants for the adaptation conductance were not changed. The strength of adaptation was not a well constrained parameter in the model of the first chapter, and in the robustness section we showed that reducing the adaptation could allow for more reasonable firing rates without demanding excitatory strengths that were unreasonable strong. We have tested that with this decreased level of adaptation, we still get reasonable levels of gain, as measured from plots of firing rates versus injected current, and from plots of firing rate versus instantaneous membrane potential, as will be discussed below.

2.3.7 Default Parameter Sets

As described in the first chapter, the main parameters that controlled the behavior of the model were the total synaptic strengths for the thalamocortical synapses (onto either excitatory and inhibitory cortical cells cells), the intracortical excitatory synapses and the intracortical inhibitory synapses. The thalamocortical synaptic strength was set at two values, with and without synaptic depression, along with the two parameters defining the non-specific background input onto all cortical cells, to achieve appropriate background firing rates, as described above. The intracortical excitatory and inhibitory strengths were set to match a number of different experimental constraints, as described

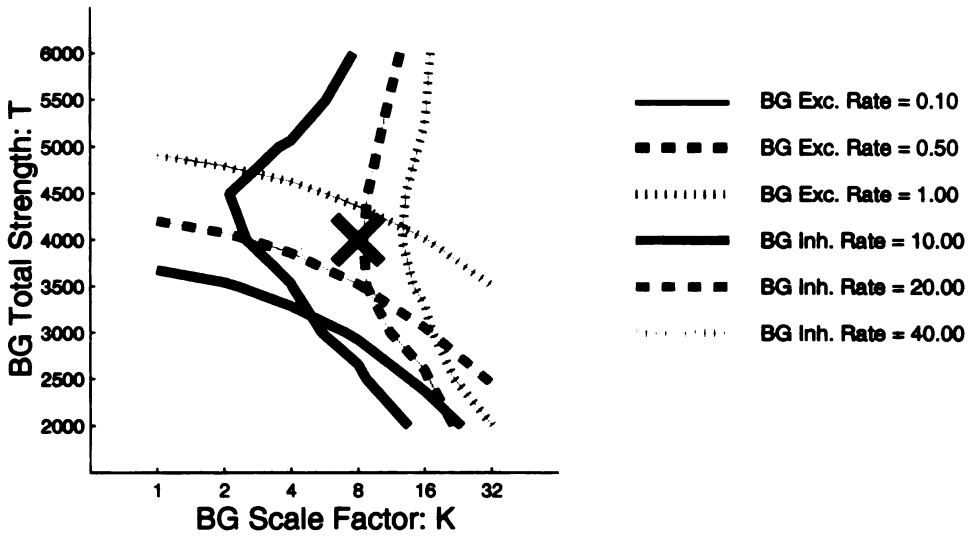

Figure 2.5:

Contour plots of background firing rates for excitatory and inhibitory cortical cells, varying the two primary parameters defining the non-specific Poisson input onto all cortical cells, with synaptic depression. T is proportional to the total mean conductance produced by the non-specific input, and K is proportional to the amplitude of the conductance of a single presynaptic spike. $T = Kr$ where r is the rate of the Poisson spike train. Thalamocortical NMDA onto E cells = 90%. TC NMDA onto I cells = 0%. Intercortical NMDA = 95%. Each contour illustrates where in parameter space the background rate of excitatory cells (blue) and inhibitory cells (red) have a particular value, as shown in the key. Large black X indicates the default parameters with depression used in this and the following chapters: $T = 4000$, $K = 8$, thalamocortical excitation = 20 nA msec. Other parameters are: Total inhibition = 7.5 nA msec; intracortical excitation = 7 nA msec.

in detail in the next chapter. That work led us to default parameter values, on which we will focus our study in the present chapter. We will use four different default parameter sets, for the cases with and without synaptic depression, and with and without the cortical "amplifier" (which simply means with or without intracortical excitation):

1. No synaptic depression and no amplifier: thalamocortical excitation = 5 nA msec, intracortical inhibition = 3.5 nA msec, intracortical excitation = 0 nA msec.

2. No synaptic depression, with amplifier: thalamocortical excitation = 5 nA msec, intracortical inhibition = 3.5 nA msec, intracortical excitation = 4 nA msec.

3. With synaptic depression and no amplifier: thalamocortical excitation = 20 nA msec, intracortical inhibition = 5 nA msec, intracortical excitation = 0 nA msec.

4. With synaptic depression, with amplifier: thalamocortical excitation = 20 nA msec, intracortical inhibition = 5 nA msec, intracortical excitation = 7 nA msec.

The thalamocortical strengths needed to be increased significantly when including synaptic depression, since the synapses are ordinarily in an intermediate state of depression, due to the background firing rate of the LGN cells. Likewise, for the two cases with the amplifier turned on, the intracortical excitatory strength is larger with depression than without, since the intracortical excitatory synapses experience synaptic depression as well. The background and stimulus-induced firing rates of the cortical excitatory cells are, in general, not as large as those of the LGN cells, and therefore they will not be strongly depressed, so the increase in total synaptic strength need not be as large to counter-balance the net effect of the depression.

We will use the parameter set with no depression and no amplifier in the beginning sections of this chapter, to demonstrate the basic behavior of the model when introducing thalamocortical NMDA. The other parameter sets were used subsequently, to demonstrate the effects of depression and the amplifier. As mentioned above, changing proportions of NMDA in any of the categories of excitatory synapses (thalamocortical⇒E, thalamocortical⇒I, or intracortical excitatory) does not change the total synaptic strength, just the percentage of the total integrated threshold current that is mediated by NMDA versus AMPA.

2.3.8 Simulations

Simulations were implemented as described in the first chapter, with a few subtle modifications. Stimulus presentations were all for one second, at a number of different temporal frequencies, so the number of cycles presented varied accordingly. By default, a "blank screen" was run for one second preceeding each stimulus presentation, both to allow the network to achieve stability before the stimulus was presented, and to determine background statistics, such as background firing rates, for every set of parameters that were run. Background and stimulus-induced output statistics were determined from the second half-second of the "blank screen" and stimulus presentations, respectively, again, to allow the network to first achieve steady-state, which, when including NMDA, took on the order of a few hundred milliseconds.

Orientation tuning curves were calculated by binning responses from all cells in the network by their preferred orientation in 5^o bins (as opposed to 10^o bins in chapter one – the bin size did not make a significant difference). All population results, such as temporal frequency tuning curves, were determined by averaging over the pool of cells in the preferred orientation bin (35 excitatory cells and 10 inhibitory cells). The cutoff temporal frequency was determined by taking the mean stimulus-induced firing rates (subtracting off background firing rates) for the cells in the preferred orientation bin, at a series of temporal frequencies (1, 2, 4, 6, 8, 12, 16, 24 and 32Hz) and using matlab's 1-D cubic spline interpolation routine, "interp1," to find the closest frequency, higher than the preferred frequency, at which the response reaches half of the maximal response.

2.3.9 A Simple Analytic Model of Simple Cell Response

To test the intuitions of the effect of NMDA conductances acting as a low-pass filter in the thalamocortical connections, we have implemented a simple analytic model of simple cells responses, to compare to the results of the detailed, network model. We consider the response of an excitatory simple cell to a drifting grating stimulus of its preferred orientation and spatial frequency with temporal frequency f. We let $L(t)$ be the mean rate of LGN input spiking to the cell, as a function of time, in response to this stimulus. $L(t)$ is periodic with period $2\pi/f$. We consider the cell to receive antiphase inhibition from an inhibitory simple cell receiving LGN input $L_{\text{inh}}(t)$ that is identical except 180^o out of phase: $L_{\text{inh}}(t) = L(t + \pi/f)$.

We ignore reversal conductance effects, and model synapses as injecting currents. We consider three anatomical types of synapses: *el*, LGN→exc. cell; *il*, LGN→inh. cell; *ie*, exc. cell →inh. cell. The excitatory synapses are of three types, N_s (NMDA with a slow decay component), N_f (NMDA with a fast decay component), or A (AMPA), while we restrict consideration of inhibitory synapses to a single type, G (GABA-A). We let a stand for anatomical type and g stand for conductance type. Following a presynaptic spike at time t_0, a synapse of type (a, g) injects a current $w_a^g \frac{1}{\tau_a^g} \exp\left[-(t - t_0)/\tau_a^g\right]$ for $t \geq t_0$. Here, w_a^g is the weight and τ_a^g the decay time of the synaptic type. These currents on cell type c, $c \in \{e, i\}$ (exc. or inh.), are converted to voltage by convolution with $\frac{R_c}{\tau_c} \exp\left(-t/\tau_c\right)$ where τ_c is the cell's membrane time constant and R_c its input resistance. Replacing the actual spike train with the mean rate of spiking, the inhibitory cell's stimulus-induced voltage response V_i, subtracting off the background voltage response, is

$$V_i(t) = \frac{R_i}{\tau_i} \exp(-t/\tau_i) \star \sum_{g \in A,N} \frac{w_{il}^g}{\tau_{il}^g} \exp(-t/\tau_{il}^g) \star L(t + \pi/f) \tag{2.5}$$

Here, $\star$ indicates convolution.

We model each cell's spiking response $r_c(t)$ as $r_c(t) = [k_c(V_c(t) - V_c^{th}) + r_c^0]^+$ where V_c is the cell's stimulus-induced voltage response, r_c^0 is the cell's firing rate in the absence of a stimulus, k_c is the gain of the cell, V_c^{th} is a threshold voltage, and $[]^+$ indicates rectification: $[x]^+ = x$, $x \geq 0$; $= 0$, otherwise. The parameters, k_i,k_e,V_e^{th}, and V_i^{th} were fit to the instantaneous voltage and firing rate responses of cells in the full model in response to an optimal high contrast grating, with no amplifier, no thalamocortical NMDA and no synaptic depression (figure 2.6). The fit parameters were k_e =14 spikes/sec/mV, k_i =20 spikes/sec/mV, V_e^{th} =1 mV, V_i^{th} =0 mV. (Note that this value for k_e is larger than, but of a similar magnitude as, the gain predicted by the linear threshold model of Carandini and Ferster [2000] – 7.2 spikes/sec/mV – who comment that their model purposely picked a threshold on the low end of the range of thresholds observed for individual spikes, and compensates by lowering the gain parameter.) For the inhibitory cell, we simplify by assuming that r_i^0 is large enough that the spiking rate never rectifies. Furthermore, since the fit value for threshold was equal to 0, we can write the stimulus induced firing rate, with background firing rate subtracted off, as $r_i(t) = k_i V_i(t)$.

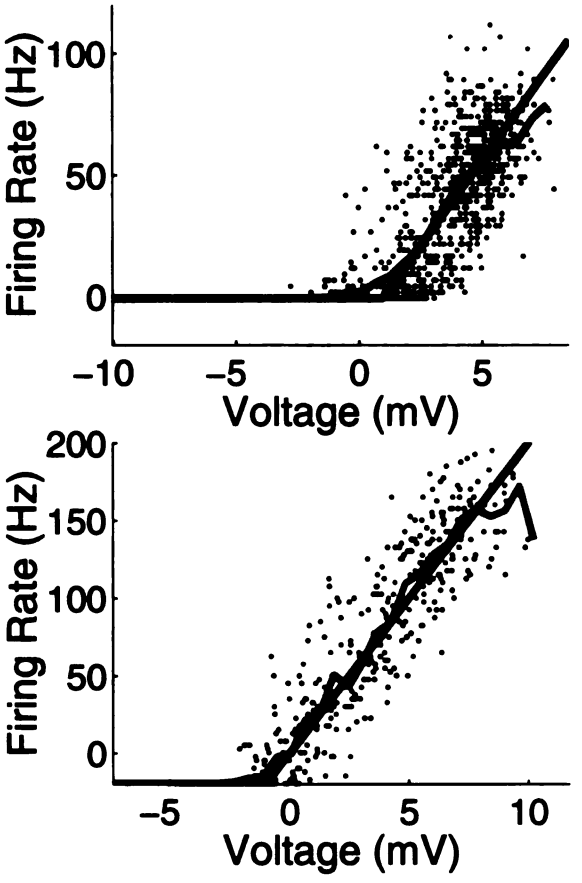

Figure 2.6:

Fitting a linear threshold spike rate model to the results of the full model. The dots represent instantaneous stimulus-induced (subtracting off background responses) voltages and firing rates of all the excitatory cells (top) and inhibitory cells (bottom) in the preferred orientation bin during a single presentation of a 2Hz, preferred orientation, high contrast grating, with no amplifier, no synaptic depression and no thalamocortical NMDA. Red traces follow the means of the firing rates over evenly spaced voltage bins. Black traces are the fits of the linear threshold model. The fit parameters, defined in the text, were, k_e =14 spikes/sec/mV, k_i =20 spikes/sec/mV, V_e^{th} =1 mV, V_i^{th} =0 mV. Note that the spike rates go below zero because they represent the rate above or below the background firing rate. The rectification of the linear threshold model occurs at the negative of the background firing rate (-0.7 Hz for excitatory cells and -19.6 Hz for inhibitory cells).

The excitatory cell's stimulus-induced voltage response is then

$$V_e(t) = \frac{R_e}{\tau_e}\exp(-t/\tau_e) \star \left[\sum_{g\in N,A} \frac{w_{el}^g}{\tau_{el}^g}\exp(-t/\tau_{el}^g) \star L(t) - \frac{w_{ei}^G}{\tau_{ei}^G}\exp(-t/\tau_{ei}^G) \star r_i(t) \right] \quad (2.6)$$

$$= \frac{R_e}{\tau_e}\exp(-t/\tau_e) \star \left[\sum_{g\in N,A} \frac{w_{el}^g}{\tau_{el}^g}\exp(-t/\tau_{el}^g) \star L(t) - \right.$$

$$\left. \frac{w_{ei}^G}{\tau_{ei}^G}\exp(-t/\tau_{ei}^G) \star \left(\frac{k_i R_i}{\tau_i}\exp(-t/\tau_i) \star \sum_{g'\in A,N} \frac{w_{il}^{g'}}{\tau_{il}^{g'}}\exp(-t/\tau_{il}^{g'}) \star L(t+\pi/f) \right) \right] \quad (2.7)$$

We assume that this voltage response is dominated by the mean, DC, and the first harmonic, F1 (the component at the temporal frequency of the grating stimulus). This assumption can be justified quite generally for simple cells with LGN inputs defined by a Gabor function with two or more subregions (T. Troyer, A. Krukowski and K.D. Miller, unpublished) and also is well satisfied

in simulations. $L(t+\pi/f)$ has DC equal to, and F1 equal in magnitude but opposite in sign to, those of $L(t)$; call the latter $\mathrm{DC_L}$ and $\mathrm{F1_L}$ respectively. Each convolution with an exponential $\frac{1}{\tau}\exp(-t/\tau)$ multiplies the Fourier transform of the voltage at frequency p by $1/(1+i2\pi p\tau)$, which multiplies the amplitude of the response at that frequency by $[1+(2\pi p\tau)^2]^{-\frac{1}{2}}$. Thus, the amplitudes of the DC and the F1 of the excitatory cell's voltage response are:

$$\mathrm{DC}_e = R_e\left[\left(\sum_{g\in N,A} w_{el}^{g} - w_{ei}^{G}k_iR_i\sum_{g'\in N,A} w_{il}^{g'}\right)\mathrm{DC_L}\right] \tag{2.8}$$

$$\mathrm{F1}_e = \frac{R_e}{[1+(2\pi f\tau_e)^2]^{\frac{1}{2}}}\left[\sum_{g\in N,A}\frac{w_{el}^{g}}{[1+(2\pi f\tau_{el}^{g})^2]^{\frac{1}{2}}} + \frac{w_{ei}^{G}}{[1+(2\pi f\tau_{ei}^{G})^2]^{\frac{1}{2}}}\frac{k_iR_i}{[1+(2\pi f\tau_i)^2]^{\frac{1}{2}}}\sum_{g'\in N,A}\frac{w_{il}^{g'}}{\left[1+(2\pi f\tau_{il}^{g'})^2\right]^{\frac{1}{2}}}\right]\mathrm{F1_L} \tag{2.9}$$

To compare temporal frequency tuning curves from the analytic model to those generated with mean stimulus-induced excitatory cell firing rates from the full model, we calculated mean firing rates by generating sinusoidal voltage traces with modulation amplitude equal to $\mathrm{F1}_e$ and mean equal to DC_e, passing them through the linear threshold model with gain $k_e = 14$ spikes/sec/mV and threshold V_e^{th} =1 mV, and taking the temporal average.

2.4 Results

2.4.1 Review of Previous Results: Orientation Tuning

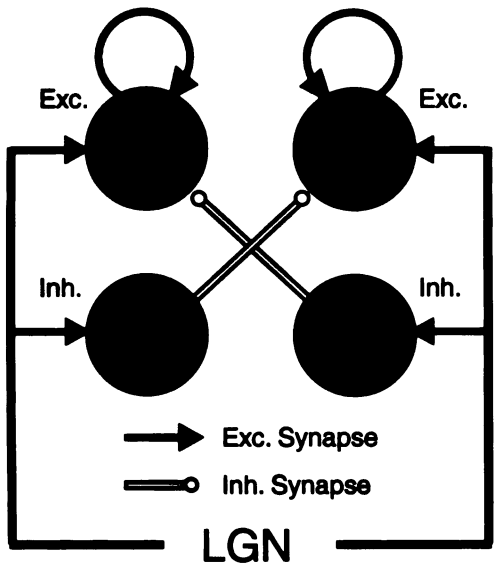

Figure 2.7:
The cartoon circuit that achieves contrast-invariant orientation tuning.

We first review the behavior of the anti-phase inhibition circuit in response to stimuli of different orientations, as described in detail in the first chapter. A cartoon version of the cortical circuit is represented in figure 2.7. Shown are four pools of cortical cells. The circles represent the input receptive fields from the LGN. Red ovals are ON subfields, that is, regions of input from On-Centered LGN cells, and green ovals are OFF subfields, regions of input from Off-Centered LGN cells. The top two pools contain excitatory cells, and the bottom two pools contain inhibitory cells. All of the cells have identical retinotopic position; their receptive fields are completely overlapping in visual space, although the circles have been separated in the figure for ease of illustration. All four pools of cortical cells have similar orientation tuning, but the left column has the opposite absolute spatial phase from the right column. That is, what distinguishes the left cells from the right is the relative position in the visual world of the ON and OFF subfields. While only two spatial phases are illustrated here, there is a similar circuit between cells of any given spatial phase and its opposite.

In general, the circuit acts as a differential spatial phase filter. Any stimulus that activates simple cells of one spatial phase and not those at the opposite spatial phase will generate a response. But a stimulus that simultaneously provides similar input to cells at two opposite spatial phases will be filtered out, due to the dominance of the anti-phase inhibition. In particular, a drifting sinusoidal grating at the preferred orientation, at any given moment in time, will activate all the cells (both excitatory and inhibitory) at one spatial phase, and not those at the opposite phase. The activated

inhibitory cells inhibit the excitatory cells of the opposite phase, which are not receiving any input anyway. The activated excitatory cells receive no inhibition, since their anti-phase inhibitory-cell partners are not activated, and so are able to respond robustly (Fig. 2.8). As the grating moves through a cycle, the activated spatial phase moves from one spatial phase to its opposite and back again, so that any particular cell's excitatory and inhibitory currents each modulate strongly but peak at opposite times in the cycle.

On the other hand, a grating at the orientation perpendicular to the preferred orientation (the null orientation) activates cells at both spatial phases simultaneously. Thus, excitation and inhibition show little temporal modulation. Since the anti-phase inhibition is dominant in this model, the excitatory cells are not able to respond (Fig. 2.8).

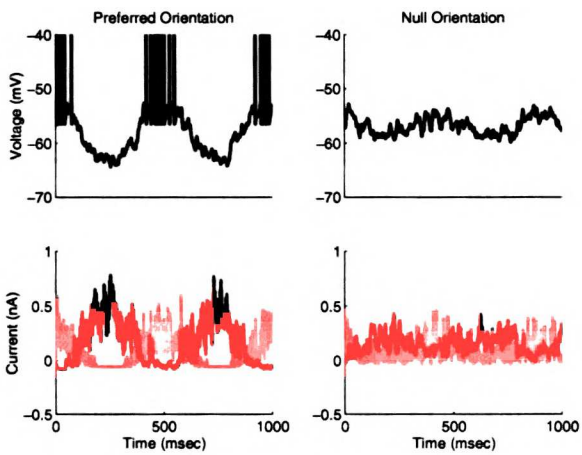

Figure 2.8:

Current and voltage traces for a single cell chosen randomly from the pool of cells at the preferred orientation (left) and at the orthogonal orientation (right) in response to a high contrast (80%), 2Hz grating. Excitatory current traces (AMPA) are shown in red. Inhibitory current traces (GabaA) are shown in blue. Default parameter set with no synaptic depression and no amplifier (inhibition = 3.5 nA msec; intracortical excitation = 0 nA msec; thalamocortical excitation = 5 nA msec.)

This behavior can be summarized by saying that the circuit acts to suppress phase non-specific input, but to pass through, and in fact amplify, phase-specific input. In the case of drifting gratings, the phase non-specific component of the input is the non-modulating, or DC, component of the input. The phase-specific component is the modulating, or F1, component.

2.4.2 Temporal Frequency Tuning: The Problem

On what time scale does this circuit determine whether or not the two phases in the cartoon circuit above are activated "simultaneously"? Clearly, for the sample 2Hz grating used in figure 2.8, the excitation and the inhibition are well separated from each other. If the grating were to move from an inhibitory to an excitatory subfield of the simple cell before the inhibition had a chance to decay away, then the excitation and inhibition would no longer be well separated, even for a grating at the preferred orientation, and the cell would not have an opportunity to respond. In the case of

the standard model with fast excitatory (AMPA) and inhibitory (Gaba-A) conductances, and no intracortical excitation, this does not occur for frequencies within the range for which the LGN cells provide significant input.

Because of the dominance of inhibition in the circuit model, the mean response is negative (more precisely, the mean conductance has a subthreshold reversal potential; this need not yield a negative voltage response). Spikes are driven by modulations of voltage about this mean: excitation and inhibition alternate each half-cycle, so the excitatory phase can drive the cell suprathreshold and yield spikes. The temporal frequency response will cutoff at frequencies for which this modulation is sufficiently attenuated that it cannot bring the cell to threshold. That is, excitation and inhibition become effectively simultaneous at temporal frequencies for which the voltage response is temporally *demodulated*, so that it becomes effectively reduced to the mean response.

Ignoring reversal potential effects, cortical voltage responses can be understood as arising from three sequential effects. LGN X-cell firing rates roughly follow the time course of the grating stimulus, for temporal frequencies to which the cells respond. This time course is convolved with the time course of synaptic conductances, which one can model as an exponential decay with time constant τ_c, to produce intracellular currents. These in turn are convolved with an exponential with the cell's membrane time constant, τ_m, to produce voltage response. The larger of these two time constants, $\tau_{\max} = \max\{\tau_c, \tau_m\}$ will dominate the filtering due to the two convolutions. As a result, the tuning curve of the amplitude of voltage modulation vs. temporal frequency f (in cycles/sec) should be roughly proportional to the product of two factors: the tuning curve of LGN firing rate vs. f; and a convolution-induced factor $1/\sqrt{1+(2\pi f \tau_{\max})^2)}$.

In particular, at the frequency $f_{\frac{1}{10}} = (3/2\pi\tau_{\max})$, voltage modulations are attenuated by 90% relative to those expected based on simple proportionality to LGN firing rates. We can consider $f_{\frac{1}{10}}$ as a ballpark estimate of this cutoff frequency. Time constants of cortical excitatory cells range from 15-24 msec in the absence of a stimulus Hirsch et al. [1998], and are lowered by stimulus-induced synaptic conductances, say to 10-15 msec. Thus, if filtering is dominated by the membrane time constant, then $f_{\frac{1}{10}} \approx 32 - 48$ Hz. This is 2-5 times the temporal frequency at which cat V1 cells typically cease responding, so this effect is not likely to play a major role in the temporally lowpass behavior of cat V1 cells.

2.4.3 Simple Solution: NMDA in Thalamocortical Synapses

If significant portions of synaptic currents are carried by slow synaptic conductances, the voltage responses at higher temporal frequencies will be demodulated. We now demonstrate that this can be induced by NMDA conductances in the thalamocortical connections of the model described previously (figure 2.9, NMDA conductances included at sites 1 and 2). For the first portion of the results, we focus only on the thalamocortical excitatory connections, and leave out the intracortical excitatory connections, which will be referred to as the "amplifier" throughout. The role of NMDA

conductances in the amplifier (figure 2.9, site 3) will be discussed below. The presence of an NMDA-mediated component of thalamocortical EPSPs has been demonstrated in the thalamocortical slice preparation [Crair and Malenka, 1995, Gil and Amitai, 1996]. We matched the relative levels of NMDA and AMPA currents in the thalamocortical connections to the levels found in the most mature thalamocortical slices studied by Crair and Malenka [1995] (90% NMDA component of integrated threshold current; we take this level of NMDA as our standard throughout this chapter; see Methods for details). Futhermore, pharmacological studies have demonstrated that blockade of normal NMDA mediated excitatory transmition can completely block [Miller et al., 1989] or interfere with [Fox et al., 1989] the response of cortical cells to visual stimuli *in vivo*.

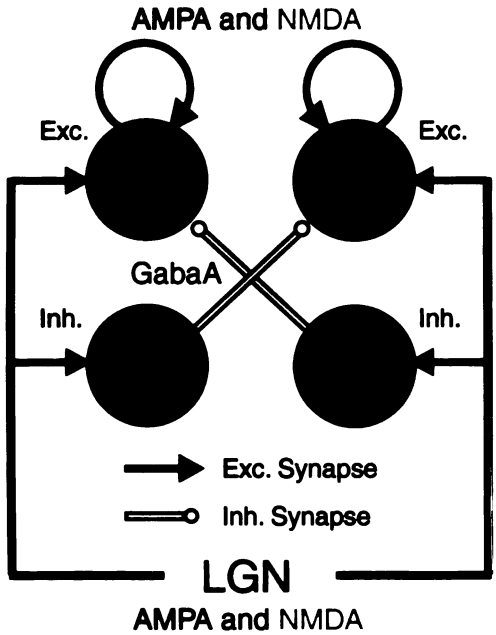

Figure 2.9:
The cartoon of the circuit showing where NMDA is being included in the circuit. The effect of including NMDA in three different sets of connections will be discussed:

1. The thalamocortical connections onto excitatory cortical cells.
2. The thalamocortical connections onto inhibitory cortical cells.
3. The intracortical connections between excitatory cells.

We model a network of conductance based integrate and fire cells, as described in detail in Troyer et al. [1998]. The effect of NMDA in the thalamocortical connections can be seen from looking at traces of a single cell in the model, chosen at random from the bin of cells that are tuned to the orientation of the grating presented (Fig. 2.10). Without NMDA (Fig. 2.10, left) the excitation and inhibition are sufficiently well separated at all frequencies that the cell is able to respond whenever the excitation is high and the inhibition is low. The long decay time constant of NMDA demodulates the input at higher temporal frequencies (Fig. 2.10, right). At 16Hz the modulations of both excitation and inhibition are greatly reduced, and the dominance of the mean inhibition over the mean excitation prevents the cell from firing. With NMDA, the temporal frequency tuning of the mean firing rate of all the cortical cells tuned to the orientation of the grating was significantly low pass shifted relative to the tuning of the LGN inputs, whereas without NMDA the tuning is similar to the LGN input (Fig. 2.11).

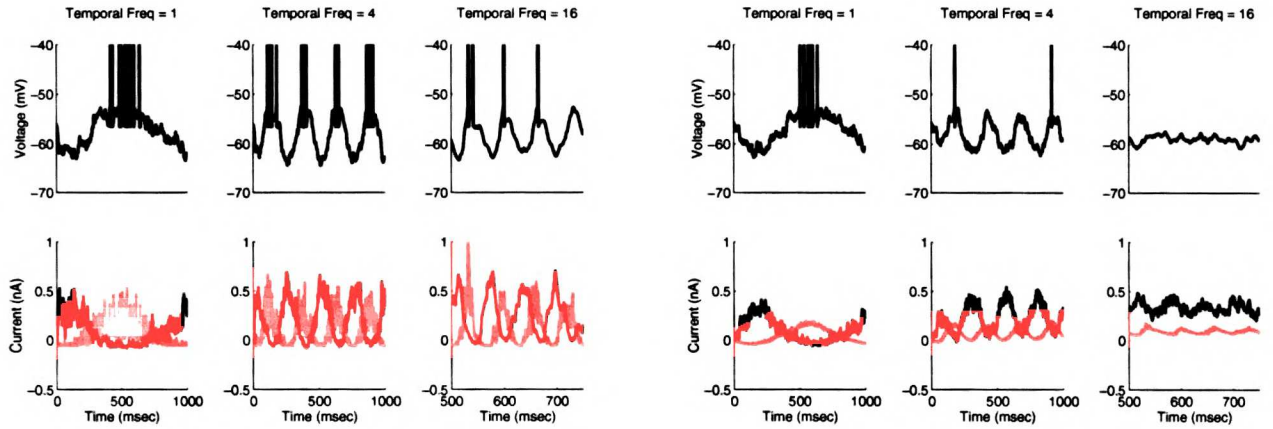

Figure 2.10:

An example of excitatory and inhibitory input to a single cell at three different temporal frequencies. Excitatory current traces (AMPA plus NMDA) are shown in red. Inhibitory current traces (GabaA) are shown in blue. Voltage responses are above in black. Left: No NMDA. Right: With NMDA. Thalamocortical NMDA onto excitatory and inhibitory cells is included; the relative amplitudes of NMDA to AMPA conductances are slightly weaker than those found in thalamocortical synapses at the oldest ages studied in Crair and Malenka [1995], which is roughly equivalent to 90% of integrated threshold current mediated by NMDA (see Methods). Default parameter set with no synaptic depression and no amplifier, as in figure 2.8.

2.4.4 Robustness with Different Levels of Thalamocortical NMDA

In the results shown above, thalamocortical NMDA was included in equal strengths onto both excitatory and inhibitory cortical cells. It is not in fact necessary to include NMDA onto both sets of cortical cells, but is sufficient to include NMDA onto either excitatory or inhibitory cells. Of particular interest is the case of NMDA onto excitatory cells only, because there is some experimental evidence that inhibitory cortical cells receive less NMDA input than excitatory cells [Ling and Benardo, 1995, Angulo et al., 1999].

Figure 2.12 shows traces for a single cortical cell, in the same format as figure 2.10, but now for the case of thalamocortical NMDA only onto the excitatory cells. The modulation of the inhibition is now preserved at all frequencies, since the inhibitory cells that provide the inhibition onto excitatory cells are receiving LGN input through AMPA synapses alone. But the excitation is demodulated at 16Hz as seen in figure 2.10. The degree of overlap between excitation and inhibition at 16Hz due to the demodulation of the excitation alone is sufficient to suppress the response of cortical cells.

Figure 2.13 shows traces with thalamocortical NMDA only onto inhibitory cells. This is the

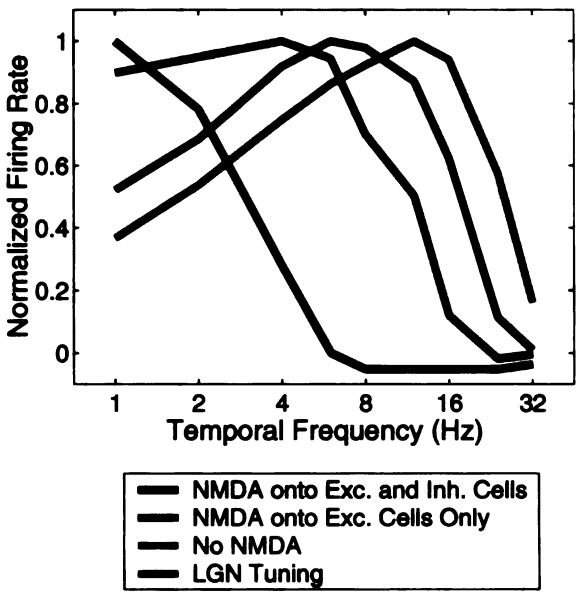

Figure 2.11:

Temporal frequency tuning of the population of cells found at the preferred orientation with and without thalamocortical NMDA, as compared to the tuning of LGN input, taken from Sclar [1987]. We show two cases: 90% NMDA on all thalamocortical synapses, as in Fig. 2.10, or 90% NMDA only on thalamocortical synapses to excitatory cells and 0% NMDA onto inhibitory cells (there are experimental suggestions of much less NMDA onto inhibitory than excitatory cells). Either case produces a strong low-pass shift in tuning. Default parameter set with no synaptic depression and no amplifier. Model curves: at each temporal frequency, mean firing rates of all cells at the preferred orientation were averaged together; the resulting population tuning curve was normalized to have a peak of 1. LGN data curve: the temporal frequency tuning curve of mean firing rate, calculated from the F1 amplitudes reported in Sclar [1987] as described in methods, was normalized to have a peak of 1.

reverse situation of figure 2.12. Now, the inhibitory currents are demodulated at high frequencies, while the excitatory currents are well modulated. But, again, the degree of overlap is sufficient to suppress response at 16Hz. In fact, we shall see that including NMDA onto inhibitory cells alone is more efficient at suppressing high frequency response than including NMDA onto excitatory cells alone. However, as mentioned above, the experimental evidence favors the case of NMDA onto excitatory cells alone.

We have looked at the population temporal frequency tuning as a function of different levels of feedforward NMDA. For all of the following figures, the temporal frequency tuning curves are the mean firing rates, averaged over the pool of excitatory cells in the preferred orientation bin, that is, cells with Gabor filter orientation within 2.5 degrees of the orientation of the grating presented. As was argued from the traces for a single cell in the model, including NMDA in either

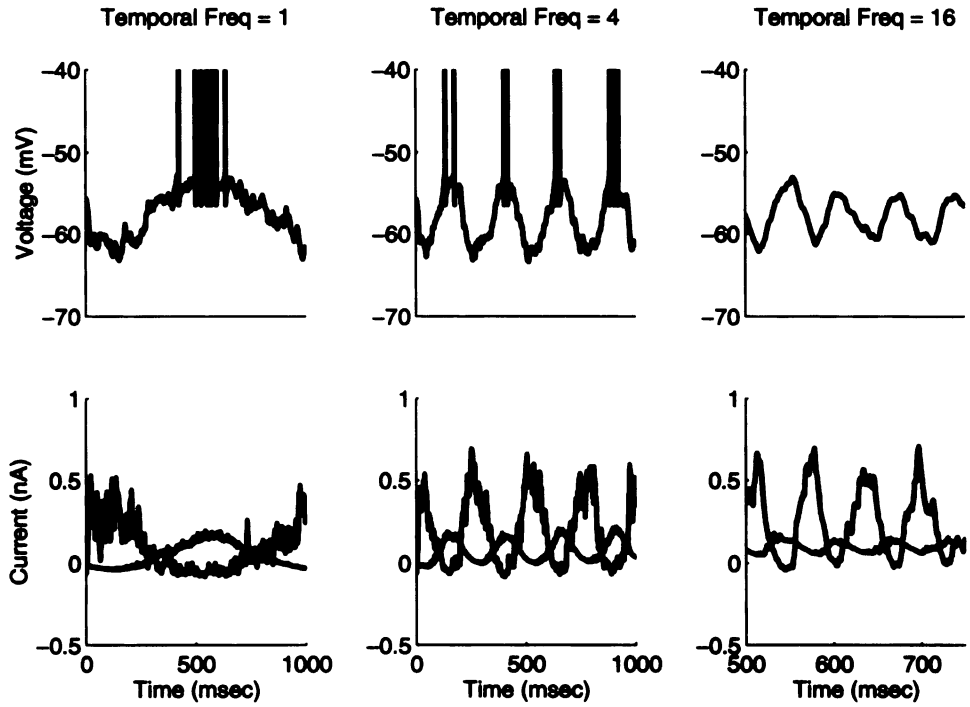

Figure 2.12:

Excitatory and inhibitory current inputs to and voltage response of a single cell in the model with 90% NMDA (fraction of integrated threshold current) included in the thalamocortical connections onto excitatory cells only. Traces are as in figure 2.10. Default parameter set with no synaptic depression and no amplifier.

the thalamocortical connections onto excitatory cells (figure 2.14 top left) or onto inhibitory cells (figure 2.14 top right) does not significantly change the firing rate at low temporal frequencies, but it suppresses the response at higher frequencies.

A quantitative measure often used to describe temporal frequency tuning curves experimentally is the high frequency cutoff, defined as the temporal frequency at which the response is reduced to 50% of the peak. It is a convenient measure for describing the tuning curves of both low pass and band pass cells since neither the peak temporal frequency nor the bandwidth are well defined for cells that respond well to the lowest frequencies commonly used experimentally (see Experimental Literature Review section above for more details). In the figures to follow, therefore, we plot the high frequency cutoff as a function of the fraction of NMDA used. Above about 40% NMDA (fraction of integrated thalamocortical circuit at threshold), the log of the high frequency cutoff decreases roughly linearly with percent NMDA in the thalamocortical synapses either to excitatory or to inhibitory cells (figure 2.14, lower left), although with different slopes in the two cases. NMDA is more effective at achieving a low pass shift in the thalamocortical synapses onto inhibitory cells.

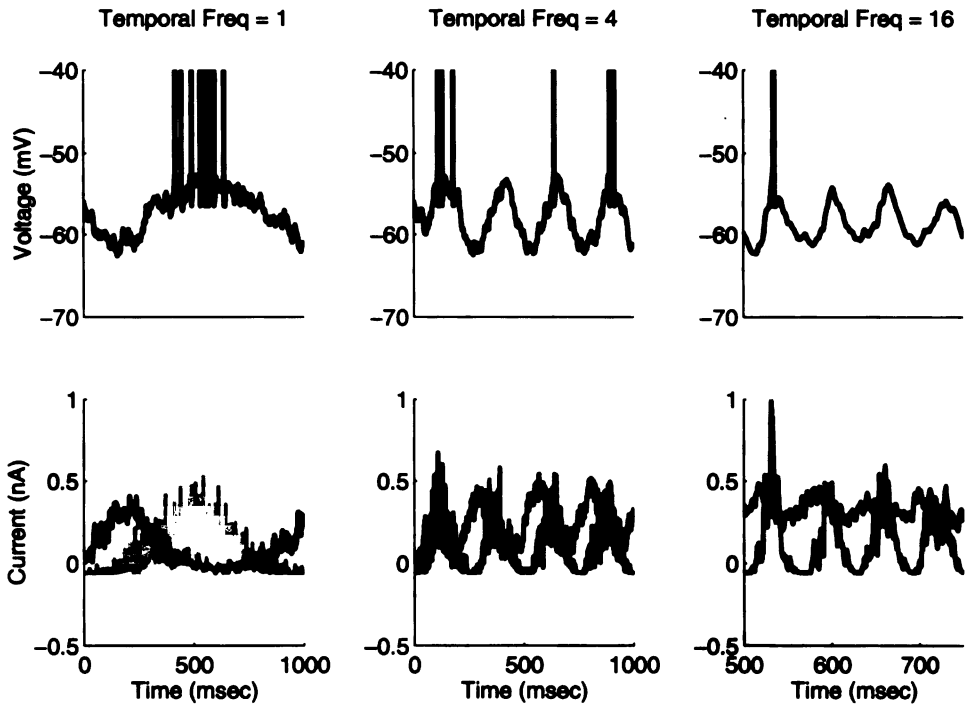

Figure 2.13:

Excitatory and inhibitory current inputs to and voltage response of a single cell in the model with 90% NMDA (fraction of integrated threshold current) included in the thalamocortical connections onto inhibitory cells only. Traces are as in figure 2.10. Default parameter set with no synaptic depression and no amplifier.

An alternative measure of the relative amount of NMDA is the NMDA/AMPA amplitude ratio, used in Crair and Malenka [1995]: the ratio of the amplitude of the NMDA EPSC at +40mV to the amplitude of the AMPA EPSC at -90mV. The cutoff frequency is close to its saturating level with an NMDA/AMPA amplitude ratio of only 10%, well below the 30% level measured by Crair and Malenka [1995].

We have also looked at the synergistic effect of increasing levels of NMDA onto excitatory cells, given an intermediate or a large amount of NMDA onto the inhibitory cells (figure 2.15). Overall, NMDA onto inhibitory cells dominates over NMDA onto excitatory cells. With the maximum amount of NMDA (90% integrated threshold current or 28% NMDA/AMPA amplitude ratio) onto inhibitory cells, including NMDA onto excitatory cells does not change the temporal frequency tuning curves very much (figure 2.15, upper right), and barely lowers the temporal frequency cutoff (figure 2.15, lower left, red trace).

However, in the converse situation of having 90% NMDA onto excitatory cells, including NMDA onto inhibitory cells shifts the temporal frequency tuning curve towards lower frequencies (fig-

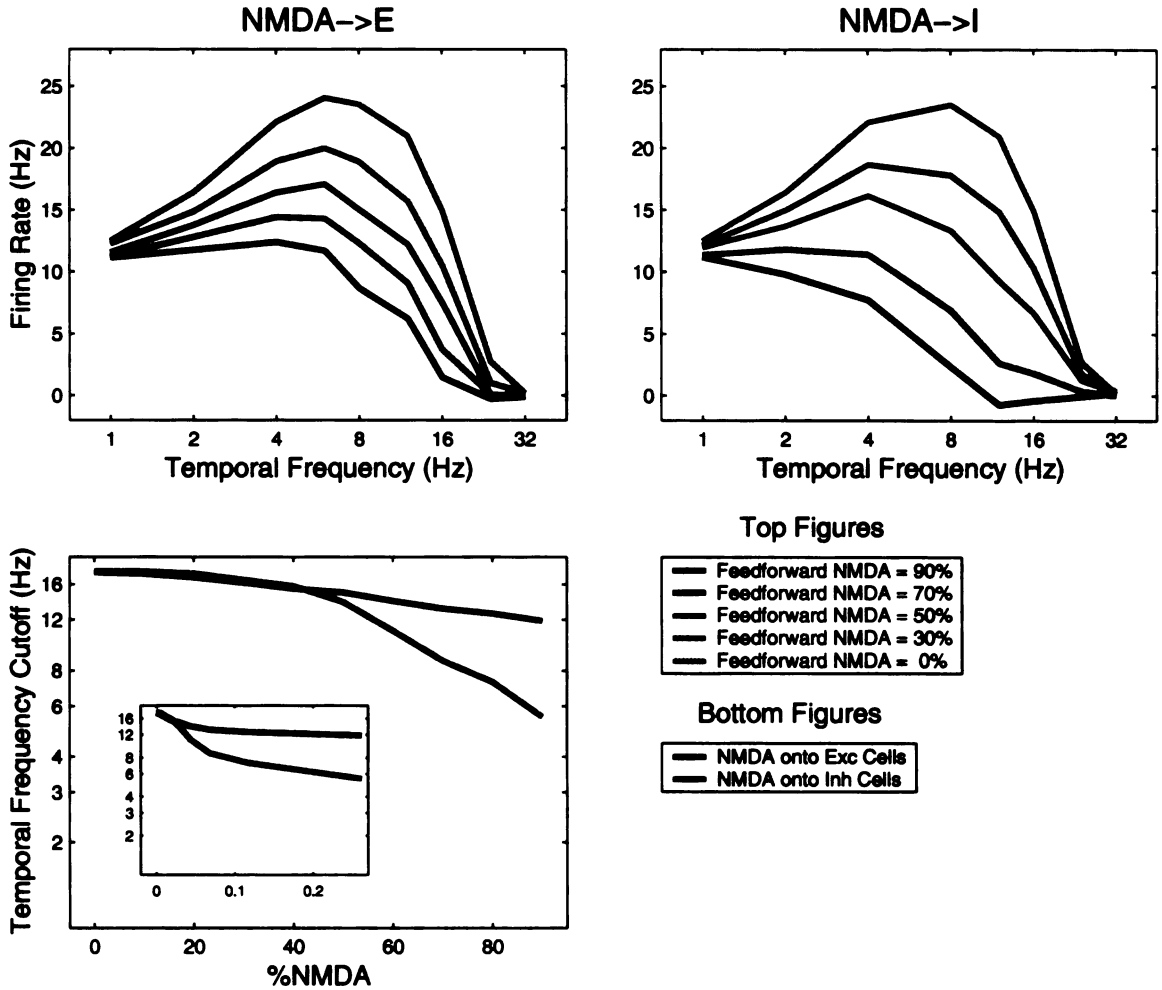

Figure 2.14:

Temporal frequency tuning as a function of fraction of feedforward NMDA onto excitatory and inhibitory cells. Top Left: temporal frequency tuning curves of the model for different levels of percentage NMDA of integrated thalamocortical currents onto excitatory cells, no NMDA onto inhibitory cells. Top Right: temporal frequency tuning curves of the model for different levels of percentage NMDA of integrated thalamocortical currents onto inhibitory cells, no NMDA onto excitatory cells. Percentage NMDA for both top figures is represented by color as shown in the key. Bottom panels summarize the tuning curves of the top panels by plotting the cutoff temporal frequency (the temporal frequency when the response reaches 50% of the peak) as a function of fraction NMDA onto excitatory cells (blue) and onto inhibitory cells (green). Inset figure is temporal frequency cutoff versus NMDA/AMPA amplitude ratios. Total integrated thalamocortical current at threshold is kept constant as the percentage NMDA is varied. Default parameter set with no synaptic depression and no amplifier.

ure 2.16, upper right) and continues to lower the high frequency cutoff (figure 2.16, lower left, red trace). The full shape of the high frequency cutoff as a function of thalamocortical NMDA onto

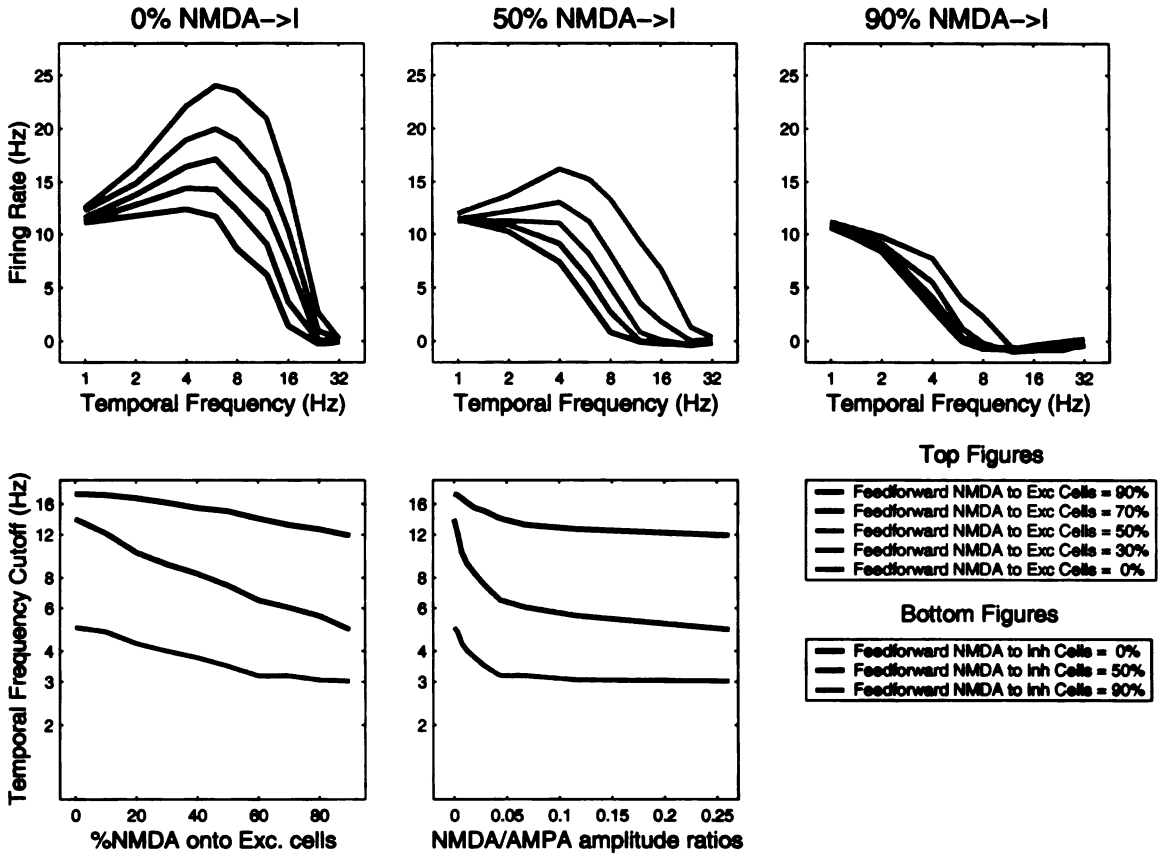

Figure 2.15:

Temporal frequency tuning as a function of fraction of feedforward NMDA onto excitatory cells. In each of the top three panels the temporal frequency tuning curves of the model are shown for different levels of percentage feedforward NMDA onto excitatory cells, represented by different colors as shown in the key. Top left panel has no NMDA onto inhibitory cells (repeat of figure 2.14, top left). Top middle panel has 50% NMDA onto inhibitory cells. Top right panel has 90% NMDA onto inhibitory cells. Bottom panels summarize the tuning curves of the top panels by plotting the high frequency cutoff as a function of fraction NMDA onto excitatory cells. The different colors represent the different levels of NMDA onto inhibitory cells (blue is no NMDA onto inhibitory cells, corresponding to the top left panel; green is 50% NMDA onto inhibitory cells, top middle panel; red is 90% NMDA onto inhibitory cells, top right panel. Default parameter set with no synaptic depression and no amplifier.

excitatory and inhibitory cells is shown in figure 2.17. The contours showing the equal levels of high frequency cutoff are all roughly parallel to each other, with a slope that shows a consistently stronger effect of NMDA onto inhibitory cells than onto excitatory cells, except at the lowest levels of NMDA, where NMDA onto excitatory and inhibitory cells are roughly equally effective. The thicker, blue-green line represents a temporal frequency cutoff of 6Hz, which is a representative experimental cutoff value for cortical cells (see experimental review above), and is equal to the cutoff

for the population tuning curve of layer 4 simple cells shown in figure 2.1. With no amplifier, this cutoff value cannot be achieved without thalamocortical NMDA onto inhibitory cells (figure 2.17, y-axis), but it can be achieved without NMDA onto excitatory cells (figure 2.17, x-axis).

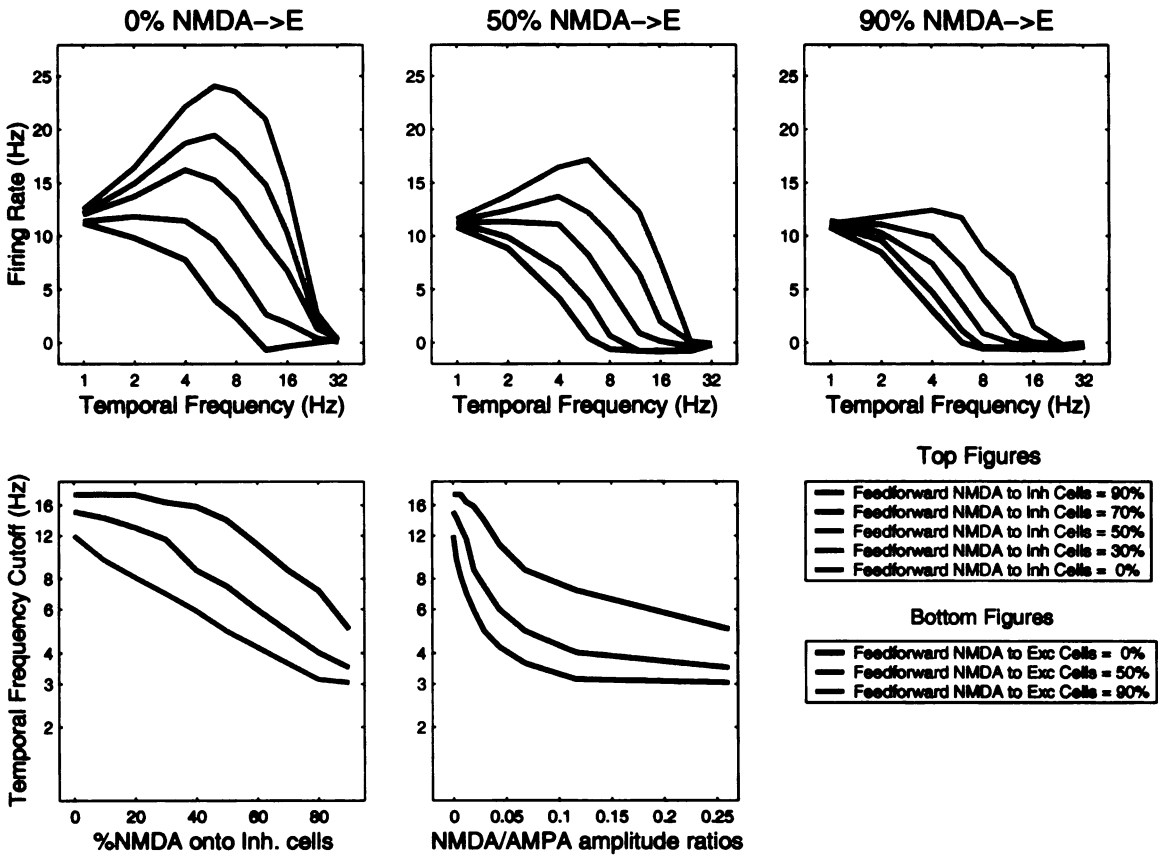

Figure 2.16:

Temporal frequency tuning as a function of fraction of feedforward NMDA onto inhibitory cells. This is the same format as figure 2.15, only the roles of NMDA onto excitatory and NMDA onto inhibitory cells have been reversed. In each of the top three panels the temporal frequency tuning curves of the model are shown for different levels of percentage feedforward NMDA onto inhibitory cells, represented by different colors as shown in the key. Top left panel has no NMDA onto excitatory cells (repeat of figure 2.14, top right). Top middle panel has 50% NMDA onto excitatory cells. Top right panel has 90% NMDA onto excitatory cells. Bottom panels summarize the tuning curves of the top panels by plotting the high frequency cutoff as a function of fraction NMDA onto inhibitory cells. The different colors represent the different levels of NMDA onto excitatory cells (blue corresponds to top left panel, green to top middle panel; red to top right panel). Default parameter set with no synaptic depression and no amplifier.

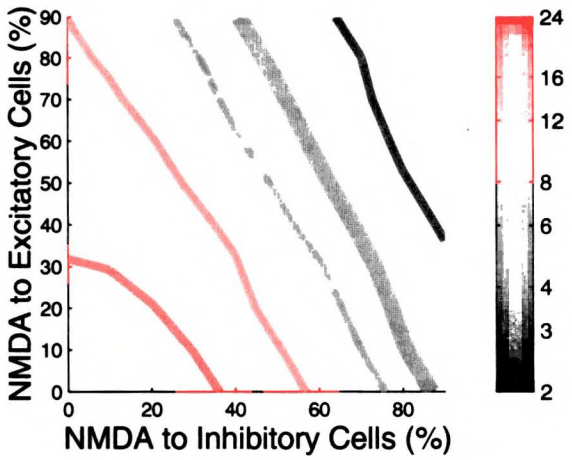

Figure 2.17:

Contour plot of temporal frequency cutoffs (frequencies at which response equals half the maximal response) shown as a function of percentage NMDA in thalamocortical synapses onto excitatory and inhibitory cells, with no amplifier. The colors of the contours represent the value of the temporal frequency cutoff, as shown in the color bar. The thicker line (blue-green) at 6Hz represents a typical experimental value of temporal frequency cutoffs for cat simple cells (see experimental review section above). In particular, the pool of layer 4 simple cells recorded by Saul and Humphrey whose population tuning curve is shown in figure 2.1 has a cutoff frequency of 6.1Hz. Default parameter set with no synaptic depression with no amplifier.

2.4.5 Comparison to a Simple Analytic Model

We have tested the intuition of thalamocortical NMDA acting as a low-pass filter on the feedforward conductances that demodulates the higher frequency input by comparing the results of the full model to a simple analytic model, described in detail in the methods section. In the simple model, we consider the responses of a single simple cell that receives direct input from the LGN and a single inhibitory simple cell of opposite spatial phase. Synaptic inputs are modeled as injected currents rather than time varying conductances, and spike rates are modeled as a linear function of the cell's membrane potential above a fixed threshold.

Despite the simplicity of this model, the resultant temporal frequency tuning curves are remarkably similar to those of the full model (figure 2.18). The largest difference is in the tuning without any NMDA, where the analytic model predicts a larger low pass shift, relative to the tuning of the LGN inputs, than the full model. The effects of NMDA are quite similar in the two models. As discussed in the robustness section above, for the full model, including NMDA in the thalamocortical connections to inhibitory cells alone is more effective as a low pass filter than including it onto excitatory cells alone, and this holds for the simple model as well. The simple model suggests the reason: the thalamocortical to inhibitory cell synaptic filter gets the added bonus of going through

the extra membrane filter of the inhibitory cells and an extra synaptic filter (albeit the fast filter due to Gaba-A currents), in addition to the membrane filter of the excitatory cell. In the frequency domain, all of these filters multiply each other, so the NMDA filter is strengthened by these extra filters, even if individually the extra filters are relatively weak.

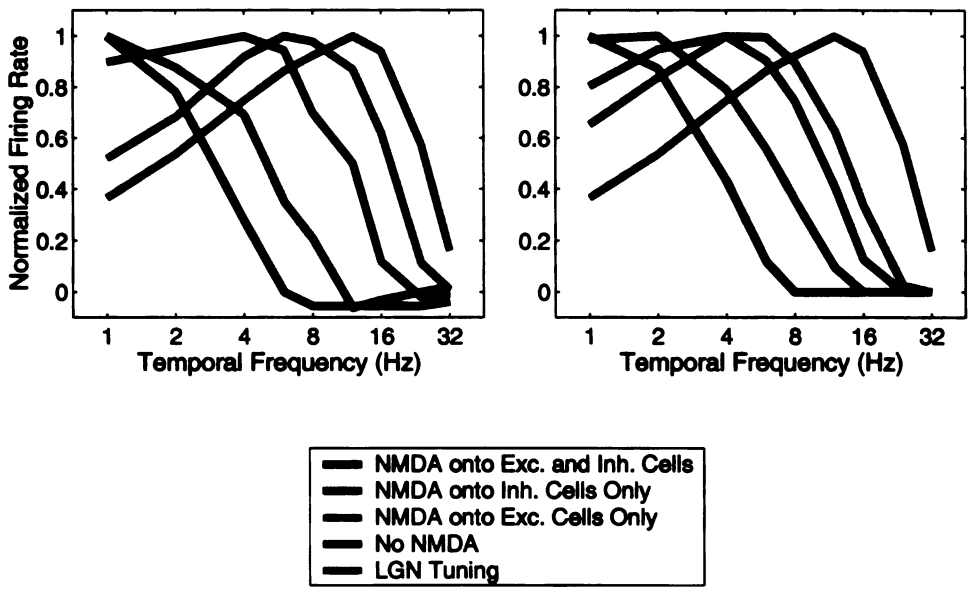

Figure 2.18:

Comparison of temporal frequency between the full model and the simple analytic model. Each figure compares the temporal frequency tuning of the mean firing rates of the LGN inputs in the model, adapted from Sclar [1987], to the tuning of the cortical cells under four conditions: no NMDA in any of the thalamocortical synapses, 90% NMDA in thalamocortical synapses onto excitatory cells only, 90% NMDA in thalamocortical synapses onto inhibitory cells only, and 90% NMDA onto both excitatory and inhibitory cells. Left figure is for the full integrate and fire circuit model, and the right figure is for the simple analytic model. As in previous figures, this is with the default parameter set with no synaptic depression and no amplifier.

While there are a number of assumptions of the simple model that distinguish it from the full detailed model, we have found that the following assumptions are sufficient to explain the differences in behavior of the two models:

1. Ignoring spike-rate adaptation conductances in the analytic model. The adaptation conductance is strongest when the excitatory cell is responding, which tends to suppress the modulation of the membrane potential.

2. Fixing the input resistances of the excitatory and inhibitory cell in the analytic model to the mean level achieved, when the cells are responding, in the full model. In the full model, the

total synaptic conductance for cells at the preferred orientation is relatively strong, which modulates the input resistance over the course of the stimulus presentation, for both the excitatory and inhibitory cells. For inhibitory cells, the input resistance is lowest during the peak of the response, and highest during the trough of the response, so the analytic model should overestimate the voltage modulation of inhibitory cells. For excitatory cells, the inhibitory conductance dominates over the excitatory conductance, so we have the opposite situation: the input resistance is highest during the peak of the response, and lowest during the trough, so the voltage modulation should be underestimated.

3. Modeling synaptic inputs as currents, and ignoring reversal potential effects, in the analytic model. The excitatory reversal potential effect is not that significant, since the voltage modulations are small compared to the distance of the threshold voltage to the excitatory reversal potential. However, it is relatively close to the inhibitory reversal potential, so ignoring the inhibitory reversal potential will overestimate the modulation due to inhibitory inputs.

4. Fixing the gains and thresholds of the linear threshold conversion from membrane potential to spike rates. While the linear threshold model seems to be fairly accurate given a fixed set of stimulus and parameter conditions, the actual gains of the excitatory and inhibitory cells in the full model are parameter dependent. Fixing them to those found with a certain set of conditions to be implemented under all conditions creates further distinctions between the two models.

5. Eliminating the voltage dependence of the NMDA conductance. The results of the full model are not strongly dependent on the voltage dependence of NMDA, because the voltage modulations of the isopotential somatic compartment are relatively small, as compared to the modulations of membrane potential in the dendrites of real neurons (where NMDA conductances are largely located) which we are not modeling. However, there are some subtle differences with and without the voltage dependence in our full model. For example, without the voltage dependence, the mean level of input is independent of the fraction of NMDA, while with the voltage dependence it is slightly dependent on fraction NMDA.

We have verified that these are the primary factors distinguishing the models, by comparing the analytic model to an artificial version of the full model in which these differences were eliminated. In the artificial full model, we set the excitatory and inhibitory reversal potentials to be very large while preserving the same amount of total synaptic threshold current. That is, synaptic conductances became very small to compensate for the large excitatory and inhibitory driving forces, to keep the currents the same size. This almost entirely eliminated both the fluctuations in membrane conductance and the reversal potential effects mentioned above. We also eliminated spike-rate adaptation and the voltage dependence of the NMDA conductances. Finally, we looked only at the direct thalamocortical input onto excitatory cells, to eliminate the problem of fitting a

fixed gain to the linear threshold conversion from voltage to the firing rate of the inhibitory cells. With this boiled-down version of the detailed model, the means and modulations of the membrane potential matched very well between the two models, over a wide range of temporal frequencies and levels of NMDA in the thalamocortical synapses.

2.4.6 Effects of the Feedback Amplifier

The results shown so far have only included the anti-phase inhibitory component of the circuit shown in figure 2.7. The effect of the same-phase excitatory connections is to amplify the phase-specific component of the input that succeeds in driving the excitatory cells. This feedback amplification increases firing rates of the excitatory cells at the preferred (and neighboring) orientations. As the strength of the excitatory connections is increased, the response gets stronger, and eventually starts to spread to neighboring orientations, broadening the orientation tuning. Eventually, with increasing excitatory strength, the amplifier runs out of control, and the cells respond at all orientations.

While the focus of this chapter is on the effect of NMDA in the feedforward excitatory connections, it is known that intracortical connections also have a significant NMDA component. In fact, there is some experimental evidence that the intracortical connections have a stronger NMDA component than thalamocortical connections. Previous models have argued that slow conductances in the intracortical excitatory connections can account for the low-pass shift in temporal frequency tuning.[Maex and Orban, 1992]

The effect of the including a feedback amplifier that is predominantly mediated by NMDA (95% NMDA in terms of integrated threshold current) on temporal frequency tuning, in conjunction with different levels of thalamocortical NMDA, is explored in the figures to follow. In general, the feedback excitatory connections selectively amplify the lower temporal frequencies, and do not significantly affect the higher frequencies (figures 2.19 - 2.21). For the case of feedforward NMDA onto excitatory cells alone, including the amplifier does lower the high frequency cutoff significantly (figure 2.19, lower left). But for either NMDA onto inhibitory cells alone (figure 2.20), or for NMDA onto both excitatory and inhibitory cells (figure 2.21), including the amplifier lowers the high frequency cutoff marginally. In general, it is clear that the main source of lowering the cutoff is from including the NMDA in the feedforward connections.

Overall, the NMDA-dominated feedback amplifier has the effect of making the role of thalamocortical NMDA onto excitatory or inhibitory cells more symmetrical than without the amplifier. Figures 2.22- 2.24 reexamine the synergistic effect of including thalamocortical NMDA onto both excitatory and inhibitory cells simultaneously, now in the presence of the amplifier. When there is 90% TC NMDA onto excitatory cells, including NMDA onto inhibitory cells no longer has much of an effect on the temporal frequency tuning curves (figure 2.23, top right) or on the high frequency cutoff (figure 2.23, bottom figures, red traces). The full shape of the surface of temporal frequency

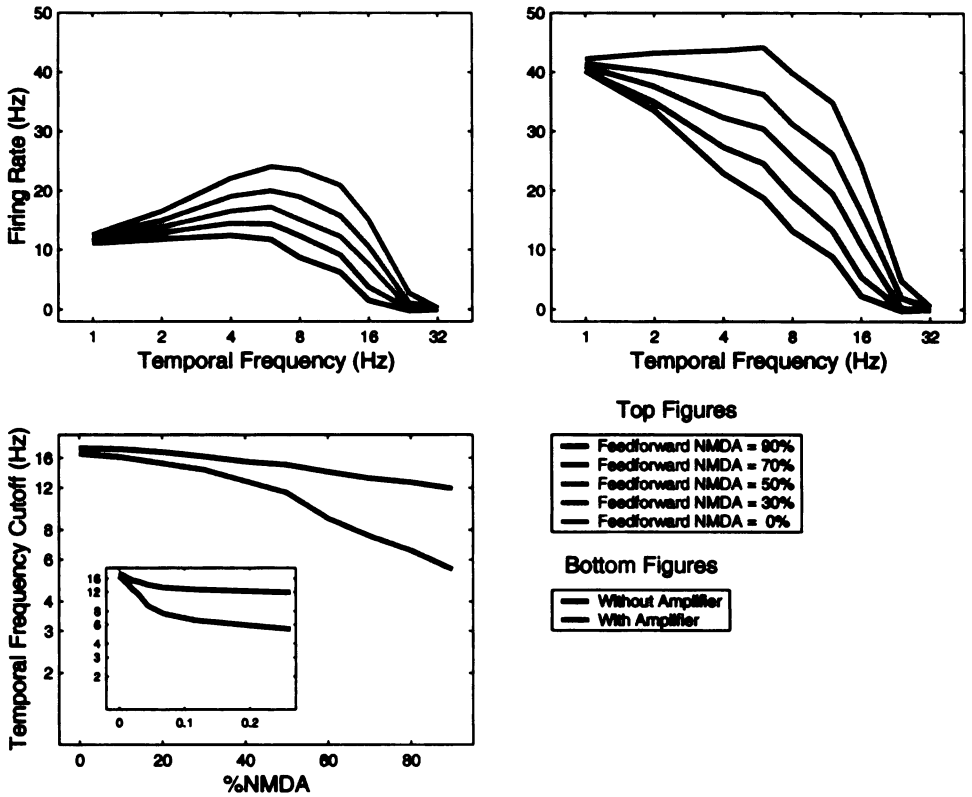

Figure 2.19:

Temporal frequency tuning of the model for different levels of thalamocortical NMDA onto excitatory cells, with and without the feedback amplifier that is 95% mediated by NMDA. Top figures, different colors represent percentage NMDA in thalamocortical synapses onto excitatory cells. Top left: No amplifier (repeat of figure 2.14, top left). Top right: With amplifier. Bottom left: High frequency cutoff vs percentage TC NMDA onto Excitatory cells. Inset: High frequency cutoff vs NMDA/AMPA amplitude ratio. Default parameter set with no synaptic depression with no amplifier (inhibition = 3.5 nA msec; intracortical excitation = 0 nA msec; thalamocortical excitation = 5 nA msec as in previous figures) and with amplifier (inhibition = 3.5 nA msec; intracortical excitation = 4 nA msec; thalamocortical excitation = 5 nA msec).

cutoff versus TC NMDA, in the presence of the amplifier, is more symmetrical with respect to TC NMDA onto excitatory and inhibitory cells (figure 2.24), than without the amplifier (figure 2.17). The contours have a slope closer to -1, which would be the slope for perfect symmetry between NMDA onto excitatory and inhibitory cells. Furthermore, all the contours have been shifted towards the origin, representing the fact that overall the cutoffs are lower at all levels of NMDA. In particular, with the amplifier, it is possible to achieve a cutoff of 6Hz without any NMDA onto inhibitory cells; that is, the thick, blue-green contour now intersects the y-axis of figure 2.24.

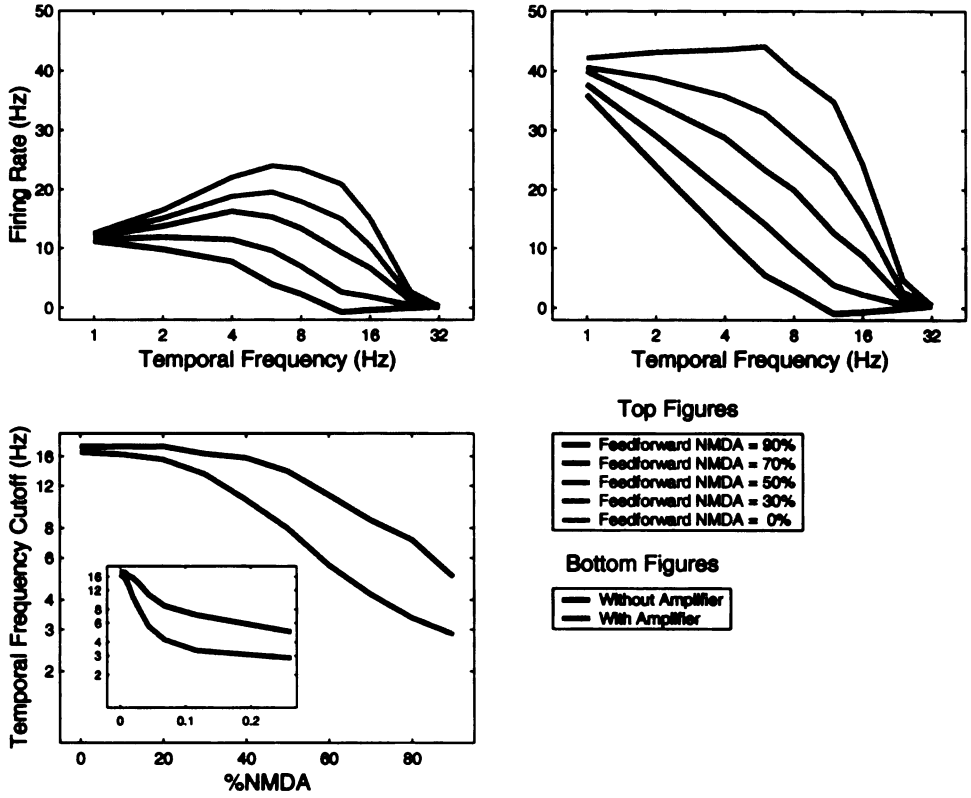

Figure 2.20:

Temporal frequency tuning of the model for different levels of thalamocortical NMDA onto inhibitory cells, with and without the feedback amplifier that is 95% mediated by NMDA. Top figures, different colors represent percentage NMDA in TC synapses onto inhibitory cells. Top left: No amplifier (repeat of figure 2.14, top right). Top right: With amplifier. Bottom left: High frequency cutoff vs percentage TC NMDA onto Inhibitory cells. Inset: High frequency cutoff vs NMDA/AMPA amplitude ratio. Default parameter set with no synaptic depression with and without amplifier.

What is the main effect of including NMDA in the feedback excitatory synapses? In general, feedback NMDA acts as a differential amplifier of lower temporal frequencies. We explored the effect of varying the amount of NMDA in the feedback connections, while keeping the total connection strength fixed (figure 2.25). We consider two situations, with thalamocortical NMDA onto excitatory cells, and without.

With thalamocortical NMDA onto excitatory cells, adding NMDA into the feedback connections preserves the amplification at low frequencies, while it reduces the amount of amplification at higher frequencies (figure 2.25, top right). The amount of amplification at different temporal frequencies is shown explicitly in the middle right plot figure 2.25, which shows the ratio of the firing rates

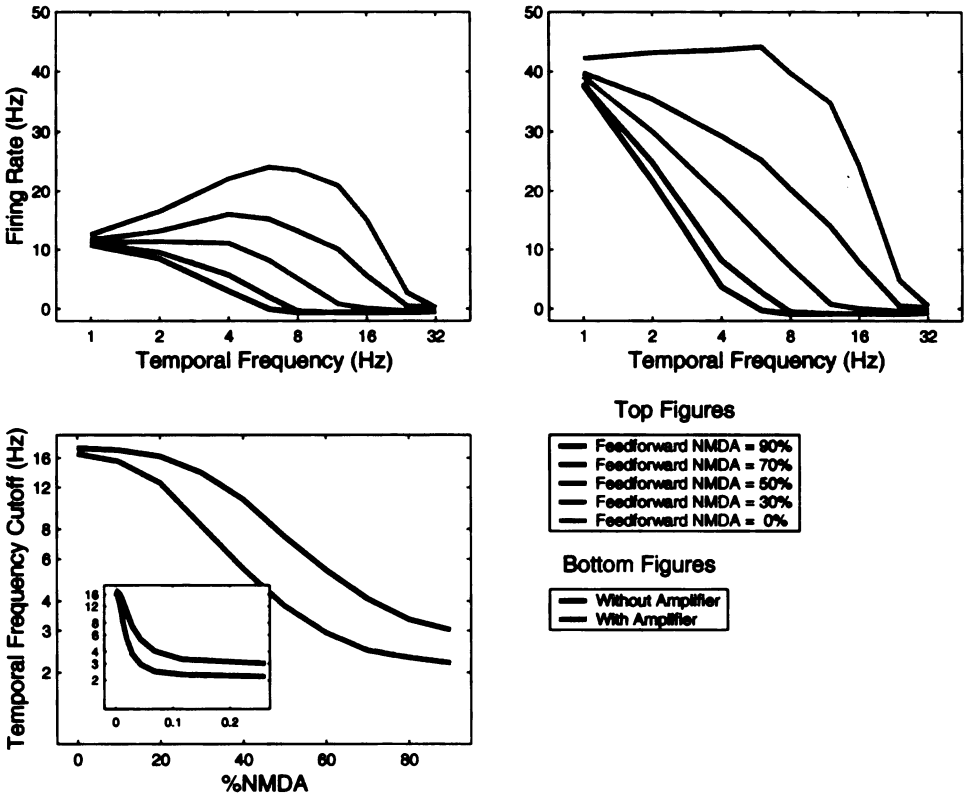

Figure 2.21:

Temporal frequency tuning of the model for different levels of thalamocortical NMDA onto both excitatory and inhibitory cells, with and without the feedback amplifier that is 95% mediated by NMDA. Top figures, different colors represent percentage NMDA in TC synapses onto all cortical cells. Top left: No amplifier. Top right: With amplifier. Bottom left: High frequency cutoff vs percentage TC NMDA onto all cells. Inset: High frequency cutoff vs NMDA/AMPA amplitude ratio. Default parameter set with no synaptic depression with and without amplifier.

with feedback excitation to the firing rate without feedback connections. The amplification with predominantly AMPA feedback connections is roughly flat out to 8Hz (purple trace). Beyond 8Hz, the cells are not responding well without any amplification, so including the amplifier does not have much of an effect. With NMDA in the feedback connections, the degree of amplification falls off well before 8Hz. This selective amplification does lower the high frequency cutoff with increasing levels of feedback NMDA, to some extent (figure 2.25, lower figure, blue traces).

With no feedforward NMDA (figure 2.25, top and middle left), the cells respond out to 16Hz, and the predominantly AMPA feedback connections amplify all frequencies almost equally. The degree of amplification falls slightly with temporal frequency, probably due to the spike-rate adaptation, which restricts how much amplification is possible at higher frequencies. But with NMDA in the

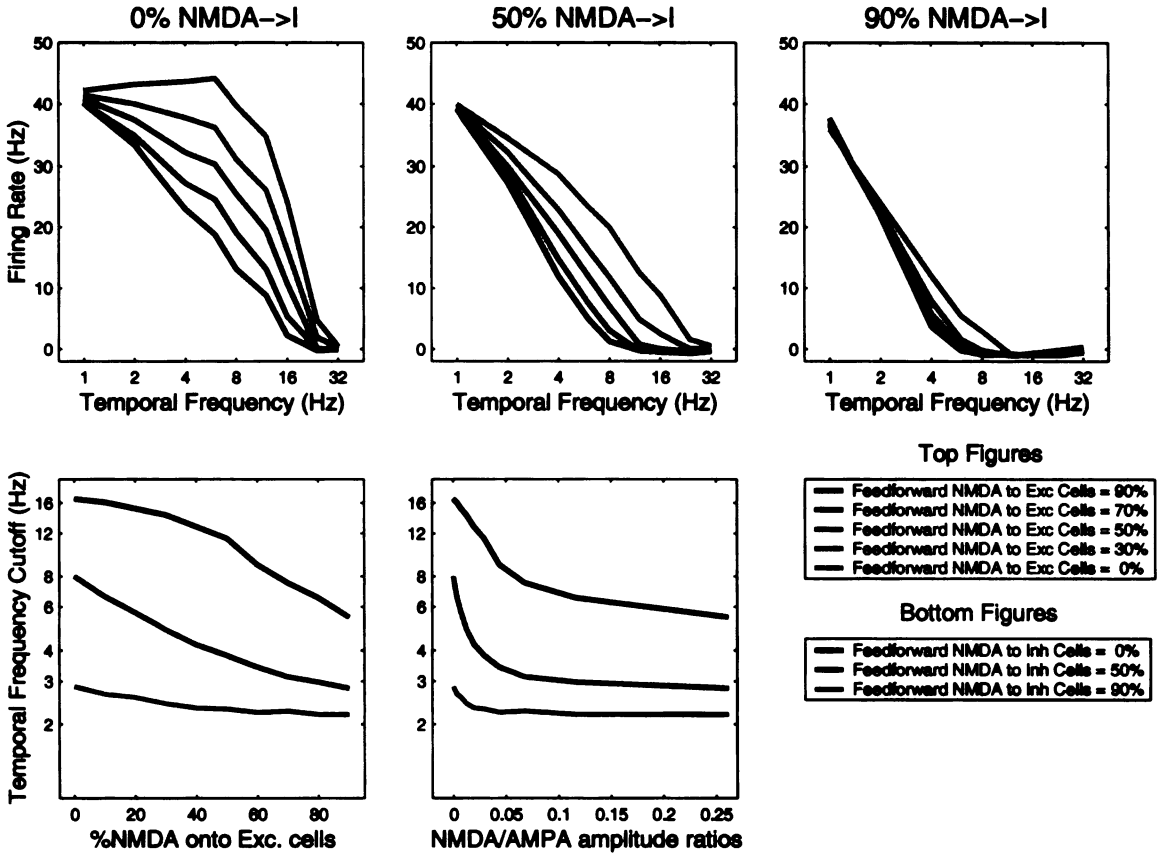

Figure 2.22:

Temporal frequency tuning as a function of fraction of feedforward NMDA onto excitatory cells, including feedback excitatory connections. Format of figure is identical to figure 2.15. In each of the top three panels the temporal frequency tuning curves of the model are shown for different levels of percentage feedforward NMDA onto excitatory cells, represented by different colors as shown in the key. Top left panel has no NMDA onto inhibitory cells (repeat of figure 2.14, top left). Top middle panel has 50% NMDA onto inhibitory cells. Top right panel has 90% NMDA onto inhibitory cells. Bottom panels summarize the tuning curves of the top panels by plotting the high frequency cutoff as a function of fraction NMDA onto excitatory cells. The different colors represent the different levels of NMDA onto inhibitory cells (blue is no NMDA onto inhibitory cells, corresponding to the top left panel; green is 50% NMDA onto inhibitory cells, top middle panel; red is 90% NMDA onto inhibitory cells, top right panel. Default parameter set with no synaptic depression and with amplifier.

feedback connections, the amplification falls dramatically with temporal frequency. Even with this selective amplification, however, including NMDA in the feedback connections does not have any effect on the high frequency cutoff, without any thalamocortical NMDA (figure 2.25, lower figure, green traces).

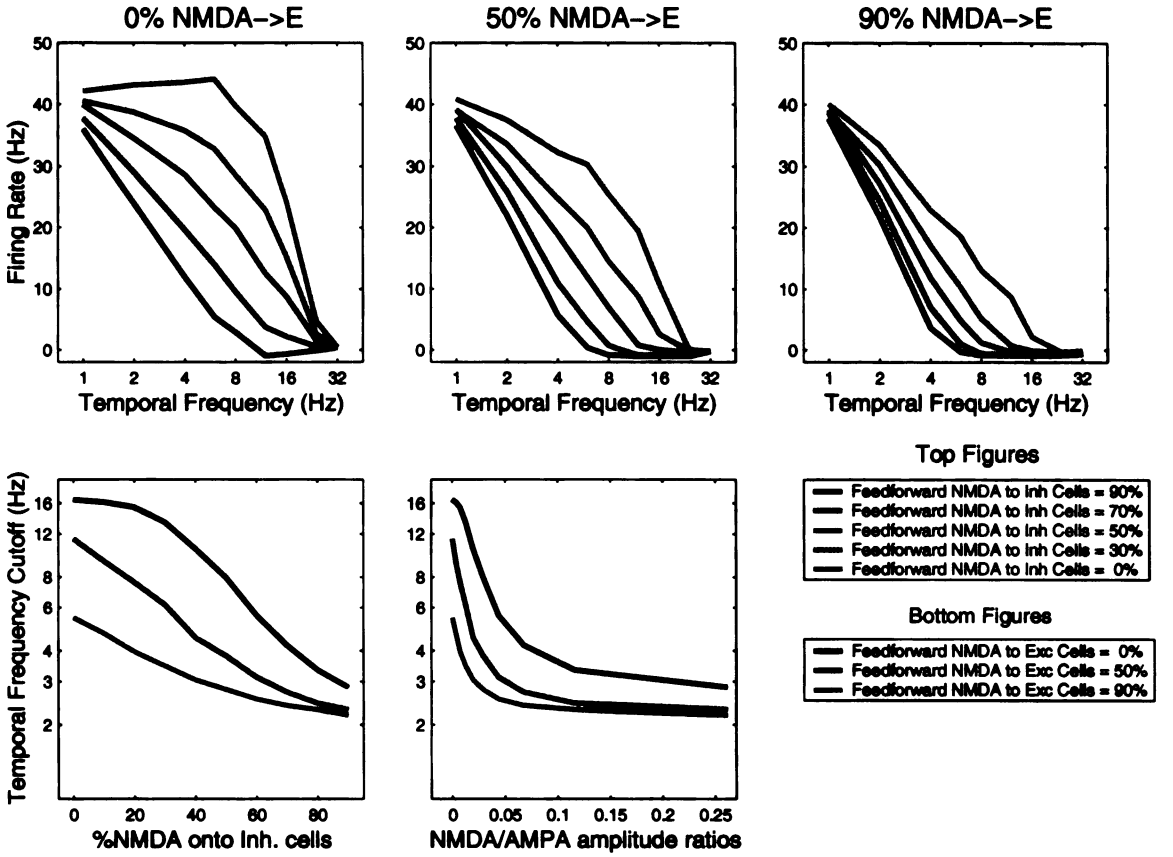

Figure 2.23:

Temporal frequency tuning as a function of fraction of feedforward NMDA onto inhibitory cells, including feedback excitatory connections. This is the same format figure as figure 2.16. In each of the top three panels the temporal frequency tuning curves of the model are shown for different levels of percentage feedforward NMDA onto inhibitory cells, represented by different colors as shown in the key. Top left panel has no NMDA onto excitatory cells (repeat of figure 2.14, top right). Top middle panel has 50% NMDA onto excitatory cells. Top right panel has 90% NMDA onto excitatory cells. Bottom panels summarize the tuning curves of the top panels by plotting the high frequency cutoff as a function of fraction NMDA onto inhibitory cells. The different colors represent the different levels of NMDA onto excitatory cells (blue is no NMDA onto excitatory cells, corresponding to the top left panel; green is 50% NMDA onto excitatory cells, top middle panel; red is 90% NMDA onto excitatory cells, top right panel. Default parameter set with no synaptic depression and with amplifier.

2.4.7 Short-Term Synaptic Depression

In addition to NMDA, another physiological property of cortical synapses that is known to exist experimentally and to act on appropriate time scales to potentially affect temporal frequency tuning is short-term synaptic depression. The basic phenomenon is that in response to a train of

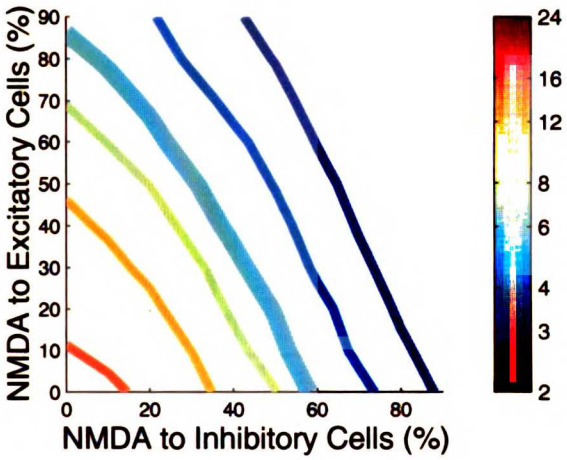

Figure 2.24:

Contour plot of temporal frequency cutoffs (frequencies at which response equals half the maximal response) shown as a function of percentage NMDA onto excitatory and inhibitory cells, with amplifier. Format the same as figure 2.17. The colors of the contours represent the value of the temporal frequency cutoff, as shown in the color bar. The thicker line (blue-green) at 6Hz represents a typical experimental value of temporal frequency cutoffs for cat simple cells. Default parameter set with no synaptic depression with amplifier.

presynaptic action potentials, the post-synaptic potentials measured are strongest in response to the initial spike, and become weaker with each subsequent spike. For sufficiently high rates, the amplitude of individual responses decays to a steady-state amplitude that is inversely proportional to the presynaptic firing rate, so that the total postsynaptic effect of the spike train is roughly independent of the presynaptic firing rate [Abbott et al., 1997, Tsodyks and Markram, 1997]. When the presynaptic cell stops firing, the postsynaptic amplitudes decay back to full strength with a time constant on the order of 100 msec or so.

We have implemented short term synaptic depression following the model described by Abbott et al. [1997]. After adjusting feedforward strengths to satisfy constraints on background firing rates, as described in the methods section, we find that including synaptic depression does not significantly affect the temporal frequency tuning of the model (figure 2.26). Synaptic depression acts as a weak high-pass filter, and moves the tuning slightly to higher frequencies, but the effect of NMDA dominates. Including an NMDA-mediated amplifier in conjunction with thalamocortical NMDA further reduces the temporal frequency cutoff, as was seen without synaptic depression previously (figure 2.27). Synaptic depression was included in these feedback excitatory synapses as well. As for the thalamocortical synaptic strengths, the overall strength of the feedback excitatory connections had to be increased to achieve sufficient amplification given the reduced strength of these synapses when in the depressed state. Otherwise, including synaptic depression in the feedback excitatory synapses had very little effect on the results of the model.

The main effect of synaptic depression in the thalamocortical synapses is to suppress the DC component of the LGN input, to some extent, at all temporal frequencies. A comparison of the

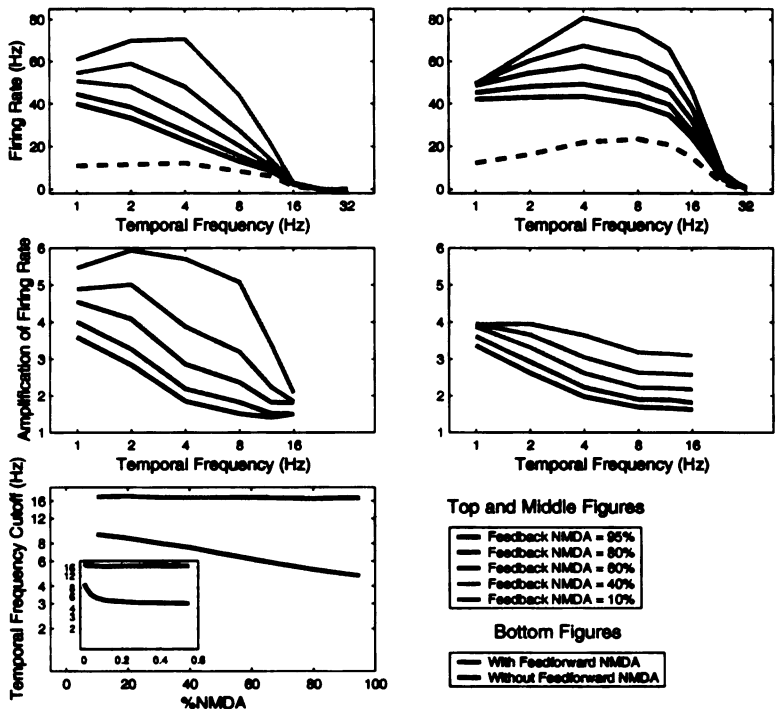

Figure 2.25:

Temporal frequency tuning of the model with different levels of NMDA in the feedback excitatory connections. For the top and middle figures, the colors represent different percentages of NMDA in the feedback connections, as shown in the key. Top figures show the temporal frequency tuning curves for different levels of feedback NMDA. Top left: With 90% NMDA in the thalamocortical connections onto excitatory cells. Top right: No NMDA in thalamocortical connections. For both top graphs, the blue dashed curve is without any feedback connections. Middle curves show the degree of amplification of the firing rate due to the amplifier, for different levels of feedback NMDA. The traces are the ratio of firing rates with the amplifier (the colored traces from the top figures) to the firing rates without the amplifier (the blue dashed figures from the top figures). Middle left: Ratios from top left figure, with 90% NMDA in the thalamocortical connections onto excitatory cells. Middle right: Ratios from top right figure, with no NMDA in thalamocortical connections. Bottom left: High frequency cutoff versus percentage NMDA in feedback connections. Inset: High frequency cutoff versus NMDA/AMPA amplitude ratios.

ratio of the stimulus-induced mean (DC) component over the modulation (F1) component of the LGN input with and without synaptic depression is shown in figure 2.28, both with and without thalamocortical NMDA. Without thalamocortical NMDA, this ratio is relatively flat as a function of temporal frequency. With NMDA, the ratio climbs steeply with temporal frequency, since the NMDA suppresses the modulation at higher frequencies, but leaves the DC component almost completely unchanged. In both cases the DC to F1 ratio is reduced at all temporal frequencies by the presence of synaptic depression; the stimulus-induced mean component of the input goes

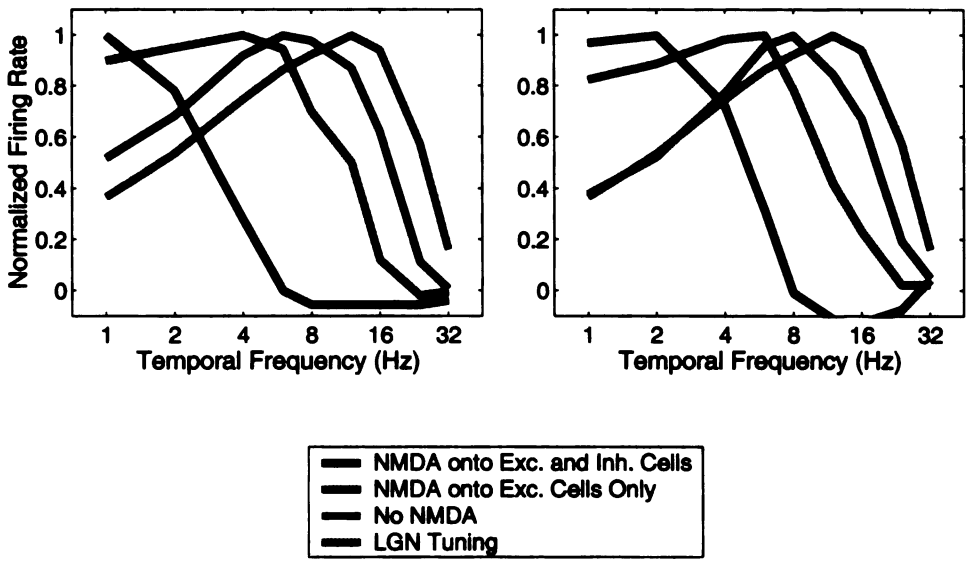

Figure 2.26:

Comparison of temporal frequency tuning with and without synaptic depression. Temporal frequency tuning of the population of cells found at the preferred orientation with and without thalamocortical NMDA, as compared to the tuning of LGN input, taken from Sclar [1987]. Left: Temporal frequency tuning without synaptic depression, as shown in figure 2.11. Default parameter set without synaptic depression and with no amplifier as in previous figures. Right: Temporal frequency tuning with synaptic depression. Default parameter set with synaptic depression and with no amplifier (inhibition = 5 nA msec; intracortical excitation = 0 nA msec; thalamocortical excitation = 20 nA msec).

almost to zero at the lowest frequencies. This DC filtering effect of the synaptic depression has some significant consequences when considering some experimental predictions of the model described below.

The suppression of the DC component is illustrated in a single cell model, for the case of no thalamocortical NMDA. The single cortical cell receives synapses from a large number of LGN cells, 500 ON cells and 500 OFF cells. We used many LGN cells in this case in order to provide good sampling. Just as in the full circuit model, each LGN cell fires as a rate-varying Poisson process, where the rate follows a rectified sine wave. All of the LGN cells modulate exactly in phase, an idealized version of the response to a preferred orientation grating.

Shown in figures 2.29 and 2.30 are the conductance traces averaged over several cycles of the grating for three different temporal frequencies. In figure 2.29, there is no synaptic depression, and it is clear that the mean conductance level is above background at all frequencies, due to LGN rectification. The mean level increases with temporal frequency, following the temporal frequency tuning of the LGN. When including synaptic depression (figure 2.30), the DC level

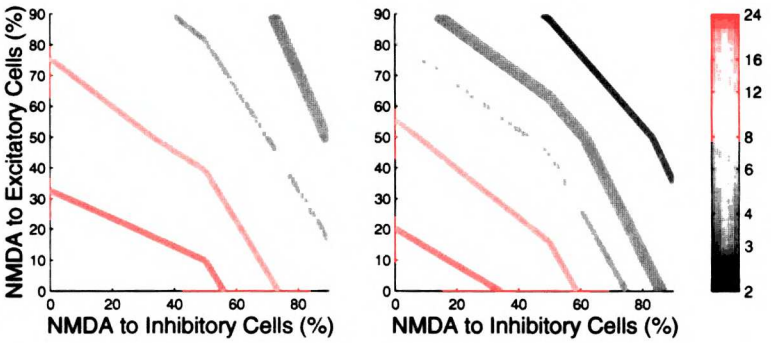

Figure 2.27:

Contour plot of temporal frequency cutoffs (frequencies at which response equals half the maximal response) shown as a function of percentage NMDA onto and excitatory and inhibitory cells, with synaptic depression. Format the same as figure 2.17. Left: Without feedback amplifier. Right: With feedback amplifier. The colors of the contours represent the value of the temporal frequency cutoff, as shown in the color bar. The thicker line (blue-green) at 6Hz represents a typical experimental value of temporal frequency cutoffs for cat simple cells. Default parameter set with synaptic depression, without amplifier (inhibition = 5 nA msec; intracortical excitation = 0 nA msec; thalamocortical excitation = 20 nA msec) and with amplifier (inhibition = 5 nA msec; intracortical excitation = 7 nA msec; thalamocortical excitation = 20 nA msec).

is suppressed completely at 1 Hz, and increases somewhat above background at higher frequencies, again, following the tuning of the LGN.

Figure 2.31 illustrates the origin of the shapes of the conductance traces in figure 2.30. At any given moment during the stimulus cycle, the total conductance can be thought of as the instantaneous presynaptic firing rate for all of the presynaptic LGN cells, multiplied by the mean weight of the presynaptic synapses. The green traces in figure 2.31 represent the normalized presynaptic firing rate, averaged over many cycles, which are determined by counting the number of spikes that occurred in each time bin. The red traces are the mean synaptic weights, averaged all the cells that happened to fire in each time bin. The simulation program does not keep track of synaptic weights at every moment in time; an individual weight is only updated whenever that particular synapse fires. As a result, when the presynaptic rates are low, or zero, there are no presynaptic cells that fire at all, and there are no statistics for the actual weights in that time window. The dotted red lines show the expected exponential recovery of a synapse that was precisely at the mean value of synaptic strength at the beginning of the inactive period (averaged over the synapses sampled in the time bin immediately preceeding the inactive period). The blue traces are the product of the green (rate) traces and the red (weight) traces, and closely match the shapes of the mean conductance

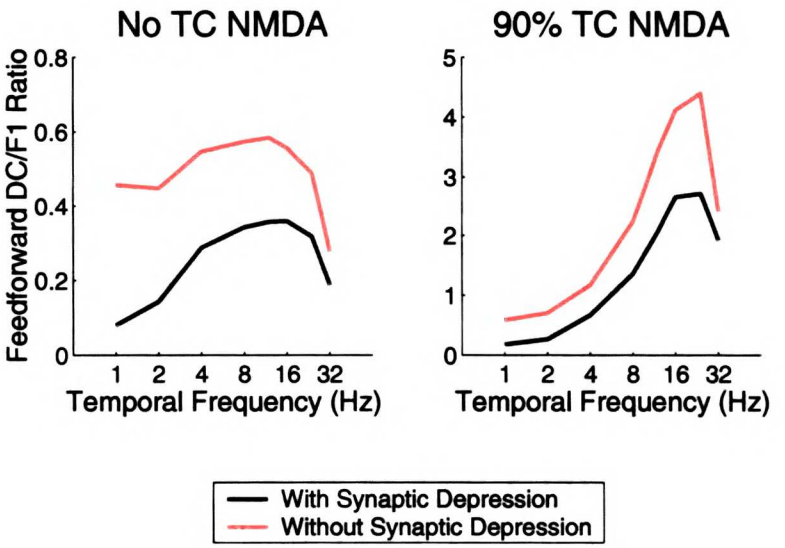

Figure 2.28:

Ratio of mean over modulation of the stimulus-induced thalamocortical conductances alone (without intracortical connections), for gratings of different temporal frequencies, with and without synaptic depression. Left: No thalamocortical NMDA. Right: 90% thalamocortical NMDA onto excitatory cells. Thalamocortical NMDA suppresses F1 at high frequencies, but not the DC, causing the ratio to increase sharply with temporal frequency: note different y-axis scale for right figure. Default parameters: without depression, thalamocortical excitation = 5 nA msec, with depression, thalamocortical excitation = 20 nA msec.

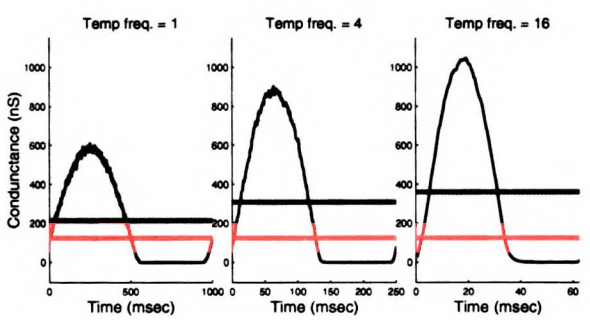

Figure 2.29:

For the single cell model with 1000 feedforward synapses: the average conductance traces at three different temporal frequencies, without synaptic depression. Horizontal red line shows the mean background conductance, and the horizontal green lines show the mean stimulus induced conductance. Synapses are purely AMPA mediated.

traces shown in figure 2.30. Note that the different peaks in presynaptic rates, due to the LGN temporal frequency tuning, have been normalized out in figure 2.31. Horizontal red lines show the mean weight during the active period of the cycle, and the horizontal blue line shows the mean conductance over the whole cycle.

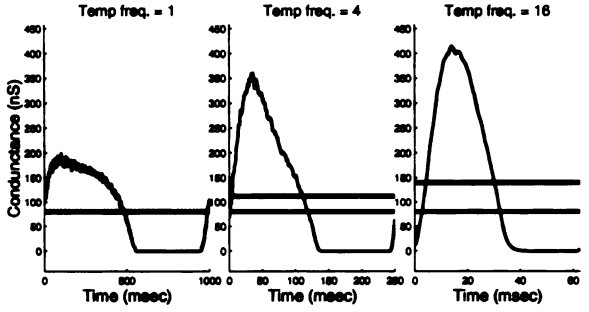

Figure 2.30:

For the single cell model with 1000 feedforward synapses: the average conductance traces at three different temporal frequencies, with synaptic depression. Synapses are purely AMPA mediated. Traces as in Fig. 2.29.

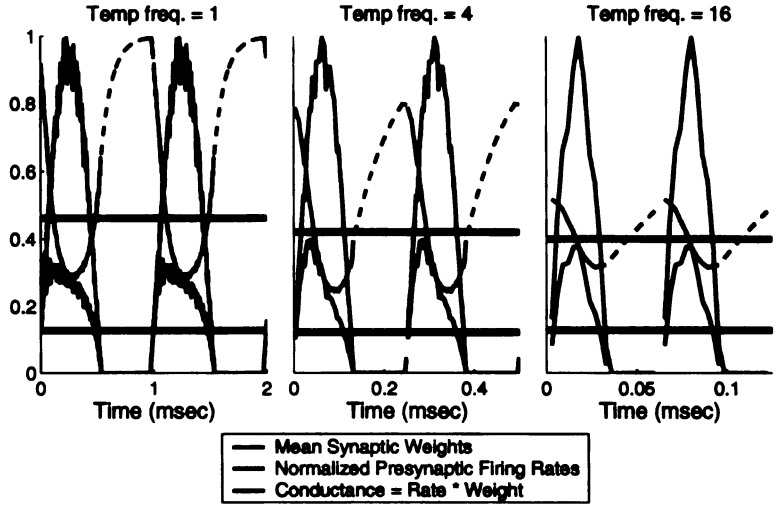

Figure 2.31:

Illustration of the shape of the conductance due to synaptic depression for three different temporal frequencies. Average presynaptic firing rates over one cycle at each frequency, normalized to the maximal rate, are shown in green. Note the different time scales for the three frequencies. Average synaptic weights are shown in red. The recovery of the average synaptic weight during the inactive period is shown in the dotted red lines. The average normalized conductance, calculated as the product of the normalized firing rate and the average synaptic weight is shown in blue. Horizontal red line is the mean synaptic weight averaged during the active period, and the horizontal blue line is the mean conductance averaged over the whole cycle.

At lower frequencies, the quiescent period is long enough for the synapses to almost fully recover from the synaptic depression to full strength. As the presynaptic rates start to climb, the mean weight starts to fall rapidly, and is at its minimum by the time the rates reach their peak. At higher frequencies, the quiescent period is very short; the weights do not have much of a chance to recover significantly. However, neither do they dip very far down because of the limited time each individual synapse has to spike, and therefore depress, during the active period. The mean weight during the active period, shown in the horizontal red lines, decrease slightly with temporal

frequency. But the mean total conductance when normalizing out the different peak values of presynaptic rates, shown in the blue horizontal lines, is almost identical at all frequencies. The remaining increases in the DC input seen in figure 2.30, therefore, are due to the LGN temporal frequency tuning.

2.4.8 Developmental changes in NMDA and Temporal Frequency tuning

As described in the experimental review section above, it has been shown that simple cells in young kittens are tuned to lower temporal frequencies than in adults. Furthermore, when kittens are reared under dark conditions, the period during which cells have lower tuning is extended. These developmental results parallel developmental changes in NMDA. The relative fraction of the decay of NMDA conductances that has a slower time constant is larger in juveniles than in adults, and this dominance of the slow component is prolonged by dark-rearing. We have tested the possibility that there is a connection between these parallel developmental changes by calculating the temporal frequency tuning responses expected in the model given different fractions of the faster and slower decay components of NMDA conductances.

Following the results of Carmignoto and Vicini [1992], we modeled NMDA conductances as having a double exponential decay term, with the slow component (200 msec time constant) representing 90% of the current for young animals. For adult animals, the NMDA decay term was dominated by the fast component (68 msec), with the slow component only contributing 12%.

With NMDA only in the thalamocortical connections onto excitatory cells, the temporal tuning curves were almost completely independent of the "age" of the model of NMDA decay (figure 2.32, top and middle left panels). For intermediate and high levels of thalamocortical NMDA onto inhibitory cells, the distinction between the fast and slow NMDA becomes more apparent (figure 2.32, top and middle, center and right panels). This is also reflected in the temporal frequency cutoffs (figure 2.32, bottom panels and figure 2.33). For regimes where the temporal frequency cutoff is relatively high, the slower NMDA does not significantly change the cutoff. In parameter regimes where the frequency cutoff is relatively low with the mature NMDA, switching to the slower NMDA lowers it further.

The dominating slow component of NMDA tends to demodulate the input at all temporal frequencies even more strongly than the fast-component-dominated NMDA. However, the remnant AMPA input, although it is a small component of the total current, is relatively independent of temporal frequency, and will dominate at higher frequencies. The difference between using the fast and slow decay components of the NMDA term will only manifest itself within the range of temporal frequencies for which the small AMPA component is not dominating the response. Beyond this frequency, the cortical response should be independent of how strong the NMDA filter is. This is why switching to slow NMDA does not affect the temporal tuning in the regime of weaker total NMDA.

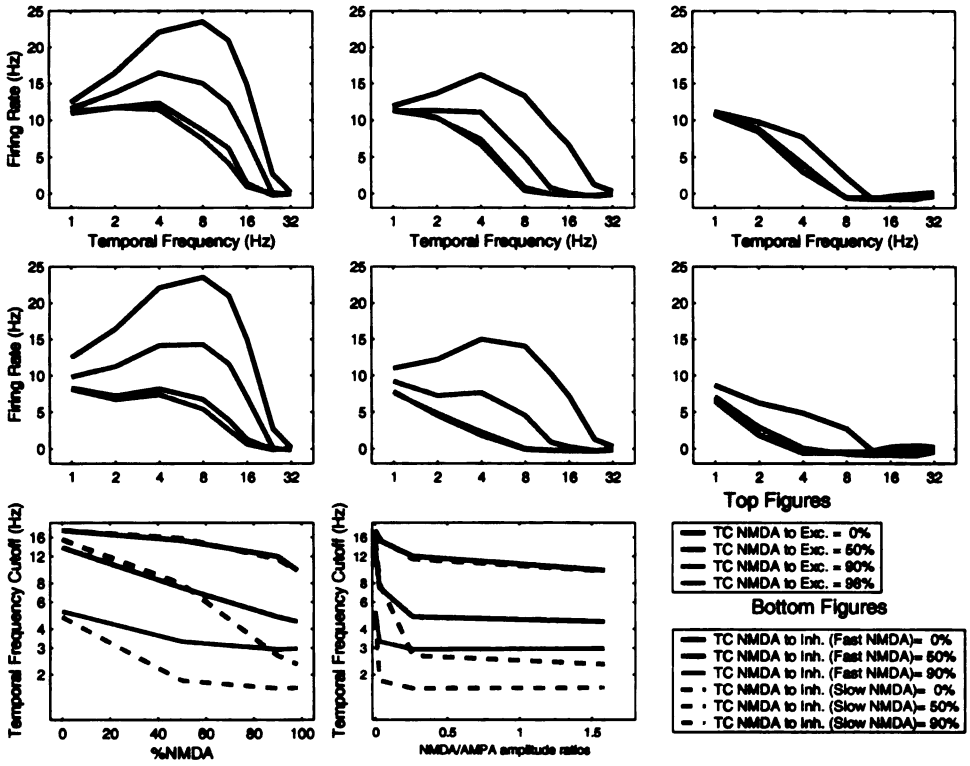

Figure 2.32:

Model of developmental changes in temporal frequency tuning, without the feedback amplifier. Top and middle panels: temporal frequency tuning as a function of fraction of feedforward NMDA onto excitatory cells, for different contribution levels of the slow component of the decay term NMDA. In each of the top and middle three panels the temporal frequency tuning curves of the model are shown for different levels of percentage feedforward NMDA onto excitatory cells, represented by different colors as shown in the key. Left panels have no NMDA onto inhibitory cells. Middle panels have 50% NMDA onto inhibitory cells. Right panels have 90% NMDA onto inhibitory cells. Top panels: adult NMDA (12% contribution of slow component) repeated from figure 2.15. Middle panels: young NMDA (90% contribution of slow component). Bottom panels summarize the tuning curves of the top and middle panels by plotting the high frequency cutoff as a function of fraction NMDA onto excitatory cells. Solid curves are for adult NMDA, taken from top panels. Dashed curves are for young NMDA, taken from the middle panels. The different colors represent the different levels of NMDA onto inhibitory cells (blue is no NMDA onto inhibitory cells, corresponding to the left panels; green is 50% NMDA onto inhibitory cells, middle panels; red is 90% NMDA onto inhibitory cells, right panels. Default parameter set with no synaptic depression and no amplifier.

It should be noted that Crair and Malenka [1995] report that NMDA is not only slower in younger animals, but also constitutes a greater proportion of the total excitation. Figures 2.32 and 2.34 include traces for 98% thalamocortical NMDA onto excitatory cells, to match the amplitude ratio of NMDA to AMPA (1.62) that they report for their youngest slices. This higher level of NMDA did not significantly change the results from the 90% maximal level (amplitude ratio of

approximately 0.3) normally used, since this is already in a saturated parameter regime, but it emphasizes the point that this strong NMDA regime, where switching to the slower NMDA does change the temporal tuning, may be the relevant regime for younger animals.

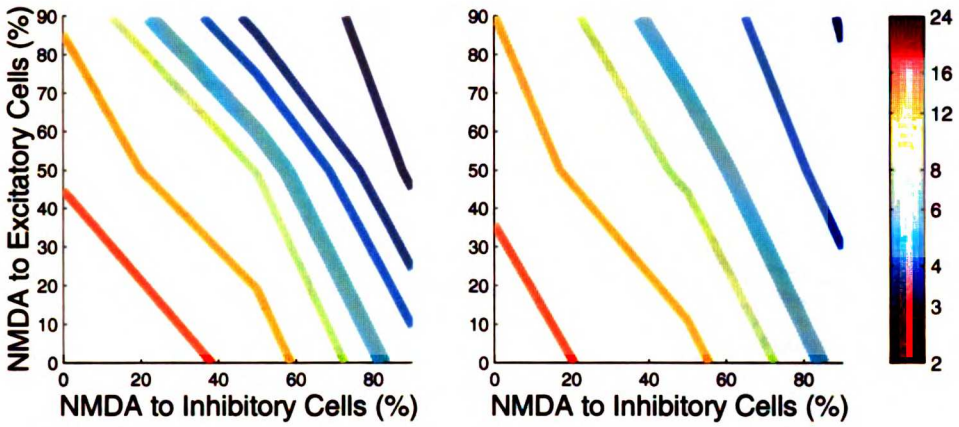

Figure 2.33:

Comparison of temporal frequency cutoff contours for young and adult NMDA, as a function of percentage NMDA onto and excitatory and inhibitory cells, with no amplifier. Left: Young NMDA (90% contribution of slow component). Right: Adult NMDA (12% contribution of slow component), as seen in figure 2.17. The colors of the contours represent the value of the temporal frequency cutoff, as shown in the color bar. The thicker line (blue-green) at 6Hz represents a typical experimental value of temporal frequency cutoffs for cat simple cells. Default parameter set with no synaptic depression and no amplifier.

The differential effect of slow versus fast NMDA is even more dependent on parameter regime when including the feedback amplifier in the model (figures 2.34 and 2.35). The slower NMDA is a much less effective amplifier than the fast NMDA, just as fast NMDA is a much less effective amplifier than pure AMPA. As a result, responses at low frequencies, where the fast NMDA is an effective amplifier, decrease significantly with the slow NMDA. At higher frequencies, where the fast NMDA is not an effective amplifier, there is much less difference between the responses with fast or slow NMDA. For low total levels of NMDA, this has the paradoxical effect of actually raising the cutoff frequency when switching from fast to slow NMDA (figure 2.34, bottom panels and figure 2.35). For higher levels of thalamocortical NMDA, the temporal frequency cutoff once again becomes lower with the slower NMDA, because the amplifier has less of an effect in determining where the response reaches half of its maximum.

2.4.9 Blocking NMDA

In the robustness section above, we analyzed the dependence of temporal frequency tuning on different ratios of NMDA and AMPA mediation of the various excitatory connections in the model. As

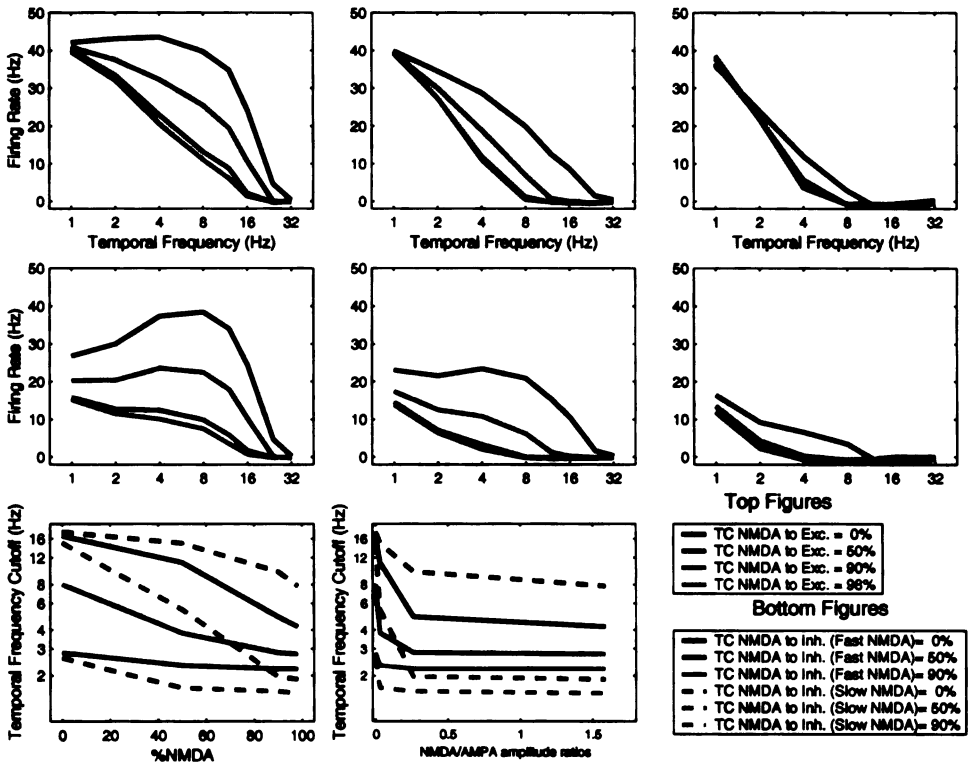

Figure 2.34:

Model of developmental changes in temporal frequency tuning, with the feedback amplifier. Format the same as for figure 2.32. Intracortical excitatory connections have 95% NMDA throughout. Top and middle panels: temporal frequency tuning as a function of fraction of feedforward NMDA onto excitatory cells, for different contribution levels of the slow component of the decay term NMDA. In each of the top and middle three panels the temporal frequency tuning curves of the model are shown for different levels of percentage feedforward NMDA onto excitatory cells, represented by different colors as shown in the key. Left panels have no NMDA onto inhibitory cells. Middle panels have 50% NMDA onto inhibitory cells. Right panels have 90% NMDA onto inhibitory cells. Top panels: adult NMDA (12% contribution of slow component) repeated from figure 2.22. Middle panels: young NMDA (90% contribution of slow component). Bottom panels summarize the tuning curves of the top and middle panels by plotting the high frequency cutoff as a function of fraction NMDA onto excitatory cells. Solid curves are for adult NMDA, taken from top panels. Dashed curves are for young NMDA, taken from the middle panels. The different colors represent the different levels of NMDA onto inhibitory cells (blue is no NMDA onto inhibitory cells, corresponding to the left panels; green is 50% NMDA onto inhibitory cells, middle panels; red is 90% NMDA onto inhibitory cells, right panels. Default parameter set with no synaptic depression and with amplifier.

the levels of NMDA varied, the levels of AMPA were adjusted to compensate and preserve the total integrated threshold current at each of the three types of excitatory connections (thalamocortical to excitatory cell, thalamocortical to inhibitory cell and intracortical). This was in an effort to see how robust the observed shifts in temporal frequency tuning are to a wide range of potential levels

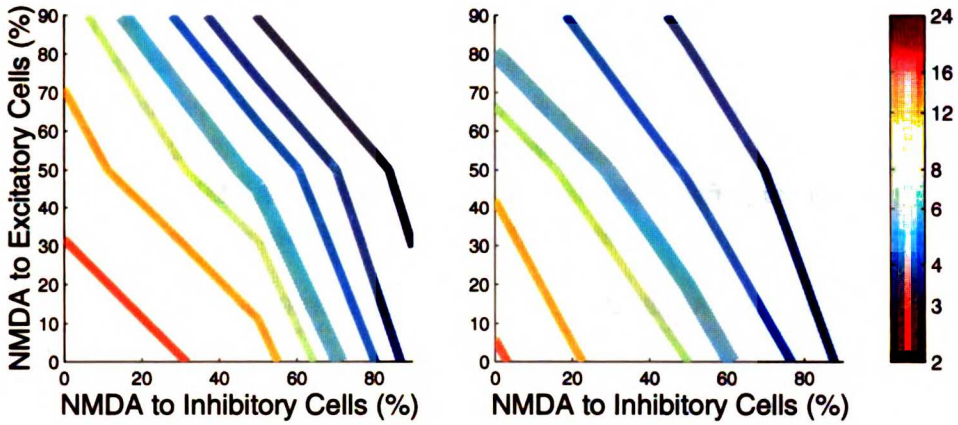

Figure 2.35:

Comparison of temporal frequency cutoff contours for young and adult NMDA, as a function of percentage NMDA onto and excitatory and inhibitory cells, with NMDA-mediated amplifier. Left: Young NMDA (90% contribution of slow component, for both thalamocortical and intracortical excitatory synapses). Right: Adult NMDA (12% contribution of slow component, for both thalamocortical and intracortical excitatory synapses), as seen in figure 2.24. The colors of the contours represent the value of the temporal frequency cutoff, as shown in the color bar. The thicker line (blue-green) at 6Hz represents a typical experimental value of temporal frequency cutoffs for cat simple cells. Default parameter set with no synaptic depression and no amplifier.

of NMDA that might actually exist in the mature cortex; the experimental evidence leaves open a range of possibilities.

It is also of interest to see how the model behaves in a simulation of sudden, artificially induced shifts in the levels of NMDA without compensating adjustments in AMPA levels, mimicing pharmocological experiments that block NMDA channels *in vivo*. As described in the experimental review section above, a number of NMDA blocking studies have been performed, with different results as to the extent that visual responses have been affected or eliminated. But the temporal frequency tuning of cortical cells under partial blockade of NMDA receptors, for which some visual response is preserved, has not been studied to date.

To study how such NMDA blockade experiments might affect temporal frequency tuning, we fixed the fraction of excitatory currents mediated by NMDA in the different excitatory connections, and then reduced the strength of all NMDA conductances by different factors, mimicing different levels of pharmocological antagonism. The effect on temporal frequency tuning depends on which synapses have NMDA. With no feedback amplifier, the shift in the cutoff frequency with NMDA blockade is in opposite directions for the case of thalamocortical NMDA onto excitatory cells alone, and NMDA onto both excitatory and inhibitory cells (figure 2.36). With TC NMDA onto both excitatory and inhibitory cells, increasing levels of NMDA blockade reveal more of the AMPA mediated temporal frequency tuning curve (figure 2.36, top left), and increase the high frequency

cutoff (figure 2.36, bottom left, blue trace). With TC NMDA onto excitatory cells alone, the total inhibitory strength stays the same with all levels of NMDA blockade. Increasing levels of NMDA blockade, therefore, effectively lower the entire tuning curve without significantly changing its overall shape (figure 2.36, top right). The high frequency cutoff, therefore, is actually lowered with increasing levels of NMDA blockade (figure 2.36, bottom left, green trace), as the weakest responses at higher frequencies are the first to be eliminated.

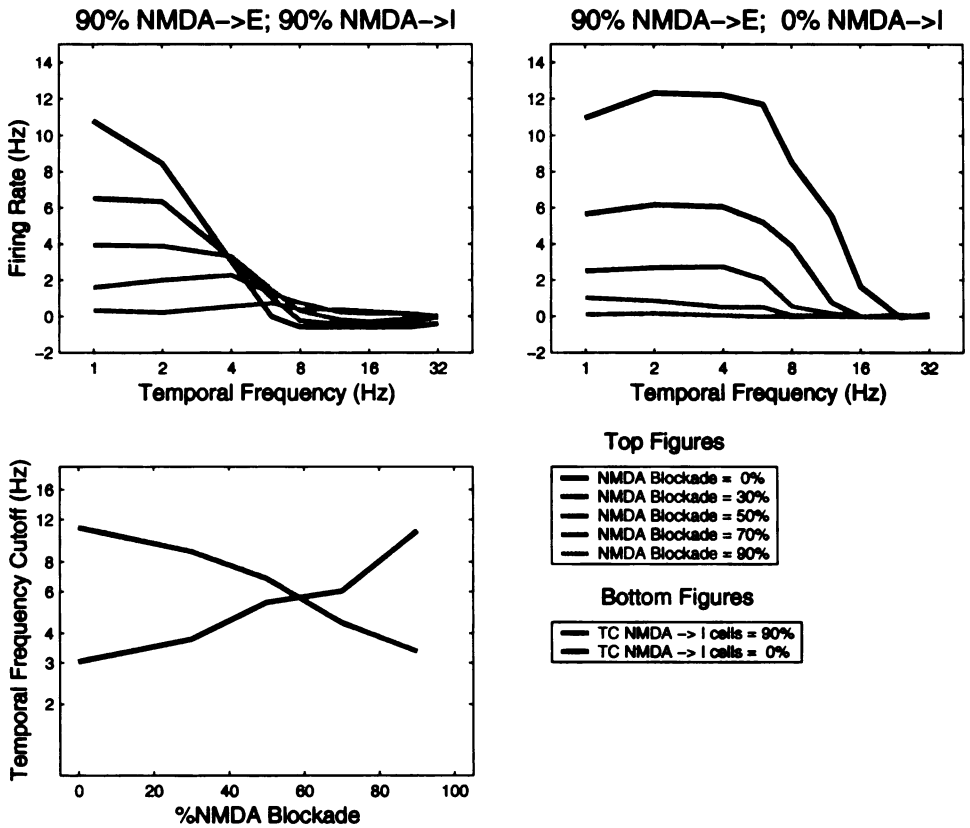

Figure 2.36:

Modeling NMDA blocking experiments, with no feedback amplifier. Temporal frequency tuning of the model for different levels of NMDA activity. Top figures, different colors represent percentage NMDA blockade as the fraction reduction from full strength, control levels. Top left: NMDA included in thalamocortical synapses onto both excitatory and inhibitory synapses (90% of integrated current for both at full strength). Top right: NMDA included in thalamocortical synapses onto excitatory cells only (90% of integrated threshold current at full strength). Thalamocortical synapses onto inhibitory cells are purely AMPA mediated and are unaffected by the NMDA blockade. Bottom left: High frequency cutoff vs percentage NMDA blockade. Blue trace is thalamocortical NMDA onto both excitatory and inhibitory cells. Green trace is thalamocortical NMDA onto excitatory cells only.

Including a predominantly NMDA mediated feedback amplifier changes the behavior of the shift in cutoff frequencies only at low levels of blockade (figure 2.37). Weak blockade of NMDA suddenly turns off the feedback amplifier, which had selectively amplified the lower frequencies. This increases the cutoff frequency to some degree, for both TC NMDA onto excitatory cells alone, and TC onto both excitatory and inhibitory cells (figure 2.36, bottom left). For stronger levels of blockade, the amplifier is effectively already shut off, and the trends are similar to those described above for the case of no amplifier.

2.4.10 Sharpening of Orientation Tuning with Temporal Frequency

Cortical cells in the model respond when there is sufficient modulation in the input to overcome the mean level of input that is dominated by the inhibition. Moving away from the preferred orientation reduces the modulation of the input, while preserving the same mean level of input. When the modulation is no longer sufficiently strong, the cells ceases to respond. The orientation at which this occurs determines the width of the orientation tuning curve.

When NMDA is not included in the model, the range of responding orientations for which the modulation is strong enough to evoke a response is roughly independent on temporal frequency. That is, the orientation tuning width does not depend of temporal frequency. However, including NMDA in the feedforward connections suppresses the modulation at higher temporal frequencies, while leaving the mean level of input roughly unchanged. This, in turn, reduces the range of orientations for which the modulation is sufficient to overcome the inhibitory-dominated mean, and narrows the orientation tuning curve at higher frequencies.

The change of orientation tuning width with temporal frequency is shown over a range of parameters in figure 2.38. The degree of sharpening is measured by determining the percentage reduction of the half-width at half height (HWHH) of the orientation tuning curve from its value at the optimal temporal frequency to its value at the highest temporal frequency where the response at the preferred orientation is still at least 20% of the response at the optimal temporal frequency: $\%Sharpening = (\mathrm{HWHH}_{\mathrm{optimalTF}} - \mathrm{HWHH}_{\mathrm{highTF}})/\mathrm{HWHH}_{\mathrm{optimalTF}}$. This measure is reported over a range of thalamocortical NMDA onto both excitatory and inhibitory cells, and over the four main parameter sets discussed above: with and without synaptic depression and with and without the feedback amplifier. With no NMDA in the thalamocortical connections onto either excitatory or inhibitory cells, there is very little sharpening in any of the four parameter sets (figure 2.38, bottom left). With NMDA included in the thalmocortical connections, the overall trend is for the degree of sharpening with temporal frequency to increase along with the strength of the low pass shifting. That is, including NMDA onto inhibitory cells, which is more effective at shifting temporal tuning to lower frequencies as observed above, is also more effective at sharpening orientation tuning at high frequencies. However, for parameter regimes where the low-pass shift is very strong, for example with high levels of NMDA onto both excitatory and inhibitory cells and including the

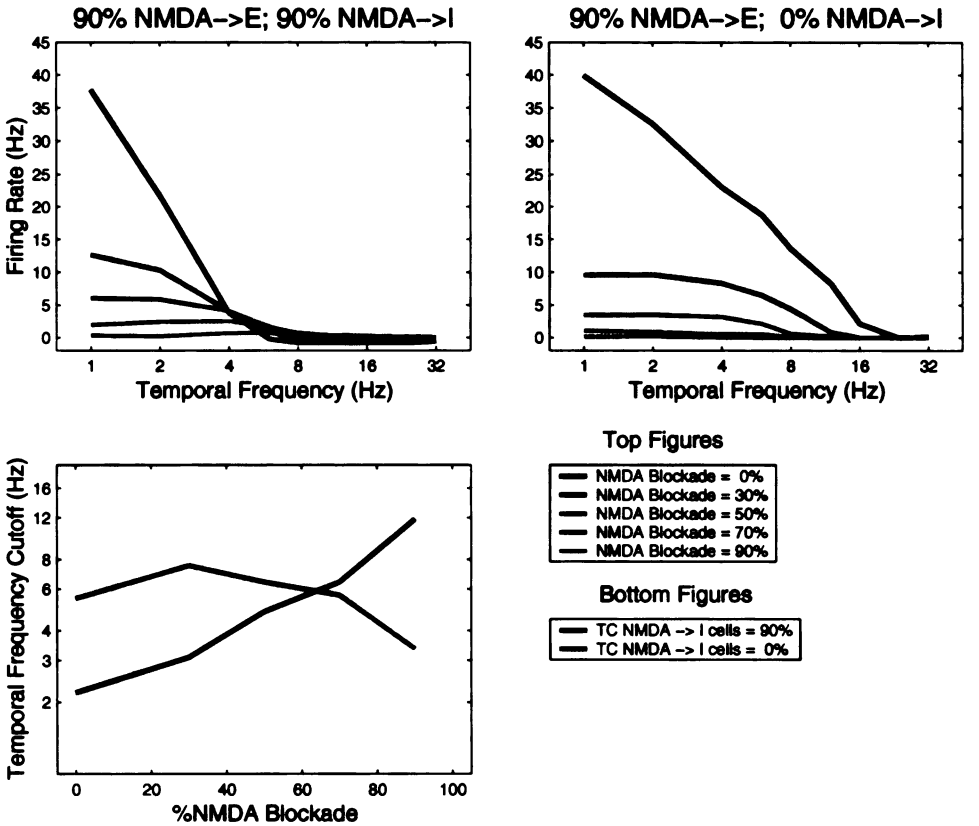

Figure 2.37:

Modeling NMDA blocking experiments, including mostly NMDA mediated feedback amplifier. Temporal frequency tuning of the model for different levels of NMDA activity, with the intracortical excitatory connections mediated by NMDA (95% of integrated current at full strength). Top figures, different colors represent percentage NMDA blockade as the fraction reduction from full strength, control levels. Top left: NMDA included in thalamocortical synapses onto both excitatory and inhibitory synapses (90% of integrated current for both at full strength). Top right: NMDA included in thalamocortical synapses onto excitatory cells only (90% of integrated threshold current at full strength). Thalamocortical synapses onto inhibitory cells are purely AMPA mediated and are unaffected by the NMDA blockade. Bottom left: High frequency cutoff vs percentage NMDA blockade. Blue trace is thalamocortical NMDA onto both excitatory and inhibitory cells. Green trace is thalamocortical NMDA onto excitatory cells only.

feedback amplifier, our measure of sharpening is low. This is because there are very few temporal frequencies among those sampled (1,2,4,6,8,12,16,24, and 32 Hz) for which the evoked response satisifed the requirement of being more than 20% of the peak response. When the temporal tuning is very sharp, the 4Hz response is already below 20% of the 1Hz response, and there is not much different in orientation tuning width between 1 and 2Hz, so we report that there is very little

sharpening. With finer sampling of temporal frequencies, it is possible that we might find more sharpening in these parameter regimes as well.

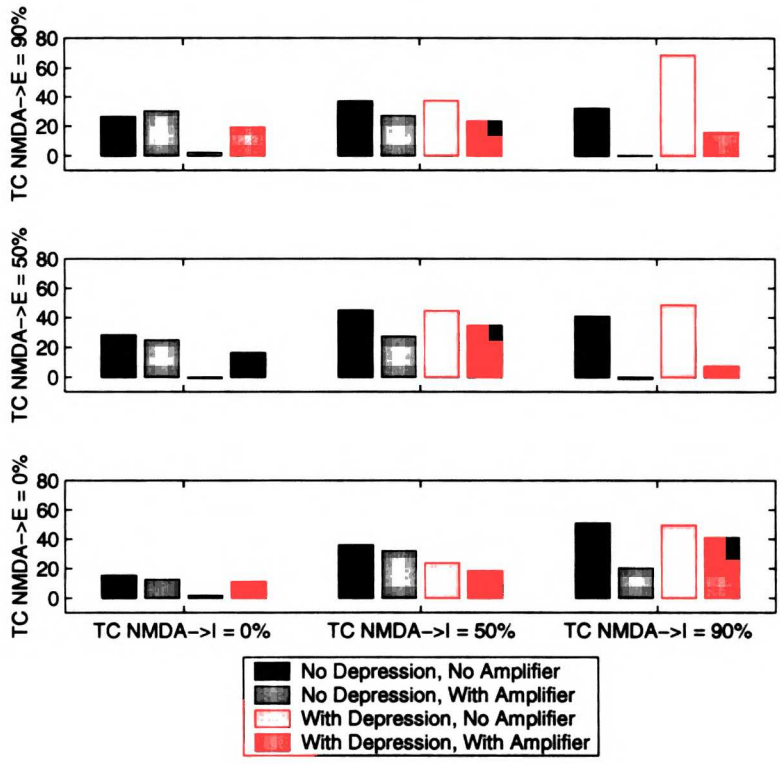

Figure 2.38:

Sharpening of orientation tuning with increasing temporal frequency. The parameter shown as the height of the bars is the degree of sharpening of orientation tuning from low to high temporal frequency, defined as the percentage reduction of the half-width at half height (HWHH) of the orientation tuning curve from its value at the optimal temporal frequency to its value at the highest temporal frequency where the response at the prefererred orientation is still at least 20% of the response at the optimal temporal frequency: $\%Sharpening = (\mathrm{HWHH}_{\mathrm{optimalTF}} - \mathrm{HWHH}_{\mathrm{highTF}})/\mathrm{HWHH}_{\mathrm{optimalTF}}$. The bars are presented in sets of four, representing the four sets of default parameters: no synaptic depression and no amplifier, no synaptic depression with the amplifier, with synaptic depression and no amplifier and with synaptic depression and with the amplifier. The sets are shown for three levels of %NMDA in the thalamocortical connections onto excitatory cells (0, 50 and 90% from bottom row to top) and onto inhibitory cells (0, 50 and 90% from left column to right).

2.4.11 Contrast Invariance of Orientation Tuning at Multiple Temporal Frequencies

Despite the fact that the orientation tuning overall depends on temporal frequency, we have found that the contrast invariance of orientation tuning is fairly robust over the whole range of temporal

frequencies for which there is a significant response, that is, a response of more than 20% of the peak response at high contrast (figure 2.39). Contrast invariance is very well preserved over all responsive temporal frequencies over a wide range of parameter space. The only exceptions are the case of high NMDA onto excitatory cells and no NMDA onto inhibitory cells. In this parameter regime, the coefficient of variation (CV) of the HWHHs of the orientation tuning curves, over the four contrasts tested, was very low except for the highest temporal frequency that was still within the 20% of peak response cutoff. At this borderline temporal frequency, the 10% contrast response was somewhat broader than at higher contrasts, creating a larger than normal CV.

2.4.12 Contrast Dependence of Temporal Frequency Tuning

We have also looked at the effect of contrast on the shape of temporal frequency tuning curves. It has been observed, although with only a few cells, that cortical cells will respond to higher temporal frequencies at high contrast than at low contrast [Albrecht, 1995]. This is also true in the LGN [Sclar, 1987], and it is unclear whether or not the degree of the shift in temporal frequency tuning with contrast is larger in the cortex than in the LGN, requiring an explanation of cortical origin. Kayser et al. [1999] provide more details of what is known experimentally on this issue to date.

Including NMDA in the thalamocortical connections actually reverses the trend of temporal tuning shifting to higher frequencies with increasing contrast that is observed in the LGN. Increasing contrast increases both the modulation (which yields net excitation) and the mean (which yields net inhibition) of the input. In parameter regimes where the modulation is strongly suppressed at high temporal frequencies by the presence of NMDA, these effects can come to nearly balance. In these regimes, the difference in response between contrasts is decreased with increasing temporal frequency, until the responses are almost the same at all contrasts. The result is that the shift of temporal tuning towards higher frequencies with increasing contrast, that is observed in the LGN inputs, is lost, and can be reversed. This is reflected in the log of the ratio of the temporal frequency cutoff at 80% contrast over the cutoff at 10% contrast shown in figure 2.40. With no thalamocortical NMDA onto either excitatory or inhibitory cells, the cortical cutoff ratio is comparable to, but slightly smaller than the cutoff ratio of the LGN input (figure 2.40, bottom left). Increasing levels of thalamocortical NMDA reduces the ratio and eventually flips the sign of the log of the ratio, particularly for the case of high levels of thalamocortical NMDA onto inhibitory cells.

2.4.13 High Inhibitory Regime Without NMDA

It is not absolutely necessary to include NMDA in the feedforward connections to achieve appropriate temporal frequency tuning in this cortical circuit. Even with all fast excitatory (AMPA) and inhibitory (Gaba-A) conductances, there is a small amount of overlap of excitatory and inhibitory current at fast temporal frequencies. If the inhibition is extremely strong, then this small amount of

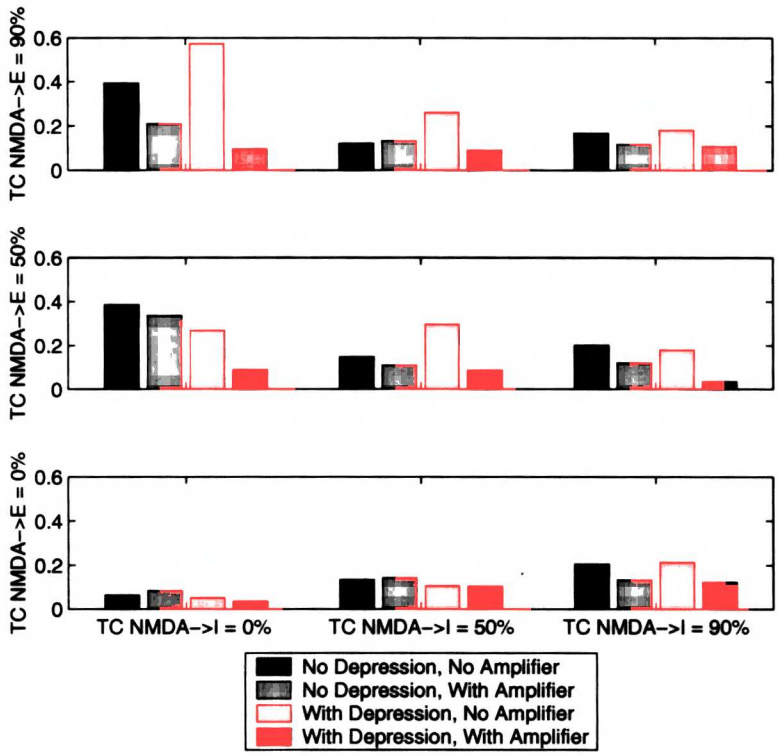

Figure 2.39:

Contrast invariance of orientation tuning over full range of temporal frequencies. The parameter used for determining the degree of contrast invariance of orientation tuning is the coefficient of variation (CV), or the ratio of the standard deviation to the mean, of the half width at half height (HWHH) of orientation tuning over the 4 contrasts used (10,20,40 and 80%). The height of the bars represents the maximal CV of HWHH over the range of temporal frequencies for which the response at the preferred orientation and high contrast was more than 20% of the response at the optimal temporal frequency. As in figure 2.38, the bars are presented in sets of four, representing the four sets of default parameters: no synaptic depression and no amplifier, no synaptic depression with the amplifier, with synaptic depression and no amplifier and with synaptic depression and with the amplifier. The sets are shown for three levels of %NMDA in the thalamocortical connections onto excitatory cells (0, 50 and 90% from bottom row to top) and onto inhibitory cells (0, 50 and 90% from left column to right).

temporal overlap can be sufficient to suppress the response of the cell at high temporal frequencies, as can be seen in figure 2.41.

We can therefore identify two general regimes of parameter space that can achieve a low-pass shift in the cortical circuit, one that we will refer to as the feedforward NMDA regime, and the other as the strong inhibitory regime. A summary of the tuning achieved in these two regimes is shown in figure 2.42. In order to distinguish between these two regimes, it is necessary to consider other experimental constraints. These constraints and a detailed search of parameter space are

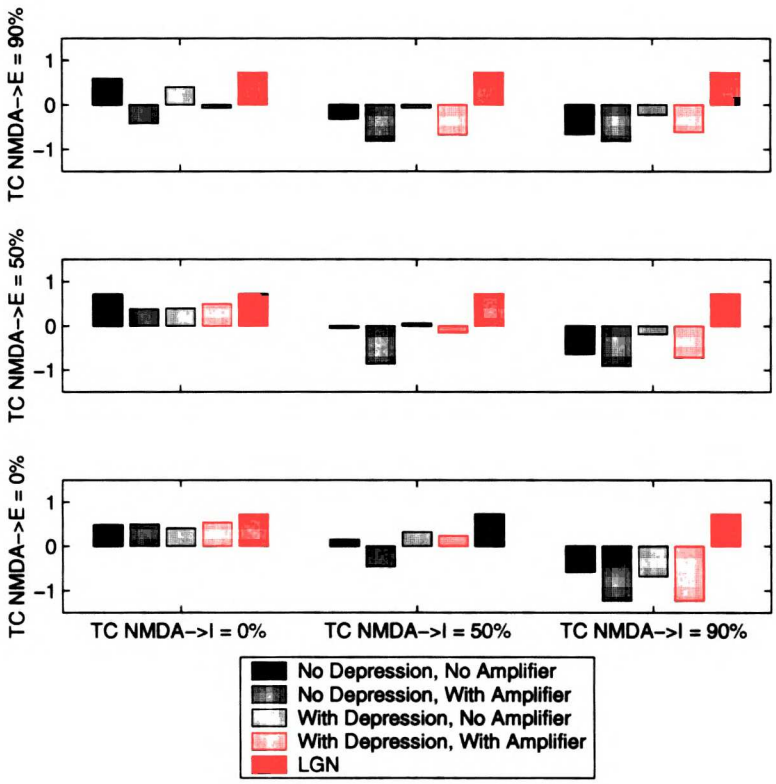

Figure 2.40:

Effect of contrast on temporal frequency cutoffs. The high temporal frequency cutoff, TF_{cut}, where the response reaches 50% of the maximal response, was calculated at both high (80%) and low (10%) contrast. The property shown in the height of the bars shown represents the ratio of TF_{cut} at high over low contrast measured in octaves, that is, $TF_{\text{ratio}} = log_2(TF_{\text{cut,80\%}}/TF_{\text{cut,10\%}})$. The bars are presented as in figures 2.38 and 2.39 with the cortical cutoff ratios presented in sets of four, representing the four sets of default parameters: no synaptic depression and no amplifier, no synaptic depression with the amplifier, with synaptic depression and no amplifier and with synaptic depression and with the amplifier. For the sake of comparison to the input, a fifth bar in each set, representing the cutoff ratio for the LGN data [Sclar, 1987] is included with each set of cortical ratios: $TF_{\text{ratio,LGN}} = 0.72$. The sets are shown for three levels of %NMDA in the thalamocortical connections onto excitatory cells (0, 50 and 90% from bottom row to top) and onto inhibitory cells (0, 50 and 90% from left column to right).

presented in the following chapter. It will be shown that the strong inhibitory parameter regime is not as consistent with other experimental constraints as is the feedforward NMDA regime, simply in that it requires a much higher level of inhibitory conductance than is normally observed. We, therefore, argue that the presence of NMDA in the thalamocortical connections is a more reasonable explanation for the observed shift in temporal frequency tuning.

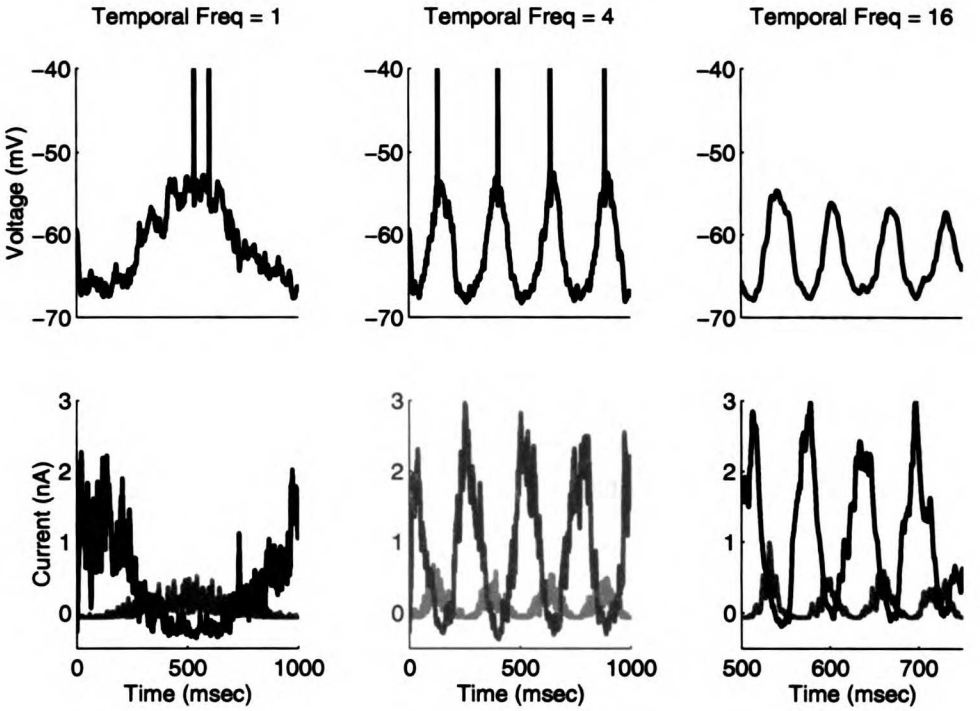

Figure 2.41:

Excitatory and inhibitory current inputs to and voltage response of a single cell in the model with no NMDA, strong inhibition. Traces are as in figure 2.10. High inhibition parameter set with no synaptic depression and no amplifier (inhibition = 15 nA msec; intracortical excitation = 0 nA msec; thalamocortical excitation = 5 nA msec).

2.5 Discussion

2.5.1 Summary

We have demonstrated that in a simple model of cortical circuitry that has previously accounted for the contrast invariance of orientation tuning of layer 4 simple cells, the presence of NMDA-mediated conductances in the thalamocortical synapses onto these cells, in proportions that have been measured experimentally, is sufficient to robustly account for the low-pass shift in temporal frequency tuning from the LGN to the cortex. Such a shift in temporal tuning appears to be a general trend in the difference between thalamic and cortical cells, which suggests that this may be a general role of NMDA in these connections in various species and various cortical regions. It was demonstrated that NMDA-mediated conductances in thalamocortical synapses were somewhat more effective at causing a shift in tuning when located in synapses onto inhibitory cells than when located in synapses onto excitatory cells. However thalamocortical NMDA onto excitatory

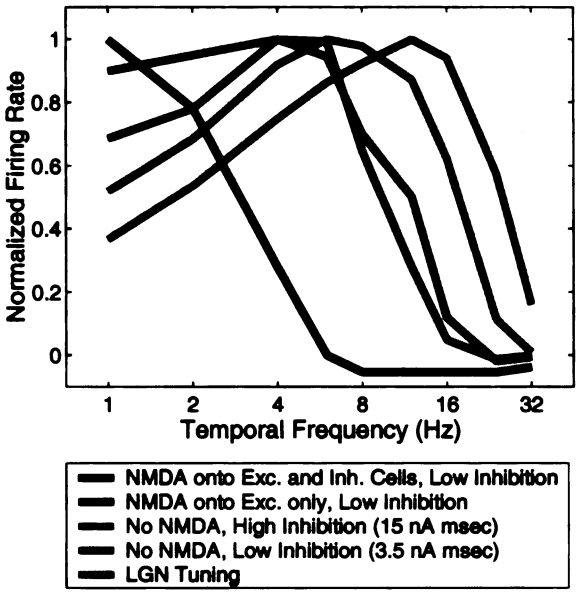

Figure 2.42:
Temporal frequency tuning of the population of cells with preferred orientation at the stimulus orientation, in two regimes: the high inhibitory regime with no NMDA, and the low inhibitory regime with thalamocortical NMDA onto either excitatory cells alone or onto both excitatory and inhibitory cells. Default parameter set with no synaptic depression and no amplifier.

cells alone did cause a significant shift from the tuning of the LGN, and, in conjunction with an NMDA-mediated amplifier, could account for the total shift that is observed experimentally. Including synaptic depression in the thalamocortical and intracortical excitatory synapses did not significantly change the shifts in temporal frequency tuning induced by NMDA.

2.5.2 Developmental Implications and Experimental Predictions

We have used the model to test the potential connection between parallel developmental changes in temporal frequency tuning of cortical cells and in the timing of the decay of NMDA conductances. Cortical cells in kittens are tuned to lower temporal frequencies than in adults cats, and NMDA-mediated EPSPs have longer decay time-constants when recorded in slices from young animals than from mature animals. Both of these developmental shifts are delayed by rearing kittens under dark conditions. Although the cells in the LGN also show a shift in temporal tuning with development, the LGN shift is not as strong as the cortical shift, indicating that there some of the developmental change is occuring at the level of the cortex.

At low levels of thalamocortical NMDA, the slower decay of NMDA observed in younger animals did not lower the temporal frequency tuning curve. In fact, when including an NMDA-dominated amplifier in the model, the slower NMDA significantly raised the temporal cutoff frequency by significantly weakening the amplifier, which was most effective at low frequencies with the mature NMDA. However, with higher levels of thalamocortical NMDA, the slower NMDA did lead to a lowering of the temporal frequency cutoff, as compared to the mature NMDA. Since the higher levels of NMDA are more in keeping with the proportion of NMDA observed in young animals,

these results are supportive of a connection between the timing of NMDA and the temporal tuning of cortical cells.

A strong test of this connection would be to test the temporal tuning of cortical cells in mature mice that have been genetically engineered to overexpress the NR2B subunit of the NMDA receptor [Tang et al., 1999]. These animals exhibit a prolonging into adulthood of the slow timing and large amplitude of NMDA-mediated EPSPs observed in wild-type younger mice. Our model would predict that these engineered mice should also exhibit a prolonging into adulthood of the slow cortical temporal frequency tuning of normal younger animals.

We have also used the model to predict the effects of partial pharmacological blockade of NMDA receptors on the temporal frequency tuning of cortical cells. The results of these tests depended on where NMDA conductances were included in the model. For thalamocortical NMDA onto both excitatory and inhibitory cells, the temporal frequency cutoff increased steadily with increasing amounts of NMDA blockade, whether or not the NMDA-dominated amplifier was included in the model. If thalamocortical NMDA was included onto excitatory cells only, and no amplifier was present, the cutoff frequency decreased with increasing amounts of blockade; including an NMDA-dominated amplifier caused an increase of the cutoff frequency at low levels of blockade, followed by a lowering of the cutoff frequency at higher levels of blockade. These results do not point to a simple, clear prediction for the proposed experiment. Although there is some experimental evidence for a higher proportion of NMDA onto excitatory than inhibitory cells, it seems more likely to observe a steady increase in temporal frequency cutoff with increasing levels of NMDA blockade. However, other results would not be inconsistent with the model, and might indicate a very weak presence of NMDA-mediated conductances in thalamocortical synapses onto inhibitory cells.

Finally, little is known to date about the interactions between temporal frequency and orientation tuning. Our model predicts that, over a wide range of parameters, levels of thalamocortical NMDA that induce a shift in temporal frequency tuning also lead to a sharpening of orientation tuning with increasing temporal frequency. McLean and Palmer [1994] collected three dimensional spectral data, varying spatial frequency, temporal frequency and orientation, for thirty simple cells from areas 17 and 18 of the cat. They therefore collected the data necessary to analyze any potential relationship between the width of orientation tuning and temporal frequency, but they do not discuss such an analysis in their paper. For the one cell for which two dimensional contour plots in the spatial frequency/orientation plane are shown for a range of temporal frequencies (see [McLean and Palmer, 1994], figure 4), it appears that there might be a slight trend towards sharper orientation tuning with increasing temporal frequency, but clearly further analysis and/or more experiments need to be performed to pursue this question. We have also shown that the model predicts that the width of orientation tuning should remain roughly contrast invariant over a wide range of temporal frequencies. This, also, has yet to be tested experimentally.

2.5.3 Comparison with Other Models

Our model is in many ways an extension and more detailed elaboration of the model of Maex and Orban [1992]. They present a model circuit, similar to ours, to account for the the shift in temporal tuning of cortical cells. They argue, as we do, that the membrane time constant of cortical cells is too low to account for the degree of shifting in temporal tuning that is observed experimentally. They discuss three potential mechanisms for this shift in tuning:

1. Feedback excitation effectively increases the membrane time constant, as long as it is significantly stronger than the feedforward excitation. They predict that the ratio of intracortical to thalamocortical excitatory conductances should have a ratio of approximately 8 to 1 to achieve the sufficient increase in time constants. They discuss the fact that such strong amplification causes stability problems, and they mention implementing some form of uniform inhibition to maintain stability, although the details of this implementation are not fully described.

2. They suggest that "mutual inhibition between neurons receiving input of opposite polarity," (which is essentially a form of push-pull inhibition, similar to ours, between simple cells of single subregions), "causes neurons to integrate the difference between their inputs with a time constant that depends on the weight of the inhibition." This description appears similar to what we refer to as the high inhibitory regime, which can cause a shift in temporal tuning even with short synaptic time constants. They say that this mechanism should only be effective for cells with a high spontaneous firing rate, and that they were only able to shift their velocity tuning with this mechanism in a simpler version of their model, and not in the detailed neuronal model presented in the paper.

3. The use of synapses with a long time course should also have a low-pass effect. In contrast to the results we have found here, they argue that temporal integration by slow synapses will be more effective in causing a low-pass shift in tuning in intracortical synapses than in thalamocortical synapses; their simulation results studied only the case of slow synapses in the intracortical synapses. They mention NMDA receptors as a potential biological mechanism for the slow intracortical excitatory synapses that their model calls for, but they say they were not explicitly attempting to model the role of NMDA receptors, because they did not model voltage-dependent conductances.

We have presented results demonstrating that including NMDA synapses in the intracortical synapses alone is not effective at suppressing the high temporal frequency LGN input and does not cause a significant shift in the temporal cutoff frequency, which is distinct from the observations of Maex and Orban [1992]. In our model, instead, the major effect of slow synapses on temporal frequency tuning occurs when such synapses are in the thalamocortical connections, although an

NMDA-mediated amplifier can cause further low-pass shifts when added to NMDA-mediated thalamocortical conductances. It would seem that the first two mechanisms mentioned above would call for strong intracortical synapses, and, in particular, strong inhibition, either to cause a temporal shift in and of itself, or to control a very strong amplifier. In the following chapter, we argue that these strong-amplifier and/or strong-inhibition parameter regimes are inconsistent with other experimental constraints.

Chance et al. [1998] explore the effects of synaptic depression on various aspects of the temporal response properties of primary visual cortical cells. They discuss the fact that the responses of V1 cortical cells shut down at frequencies lower than their LGN inputs, however, they do not implicate synaptic depression as being the primary mechanism responsible for achieving this shift in tuning. To model the LGN input feeding into a single cell, they use Poisson spike trains with rates varying as half-wave rectified sinusoids with the same amplitude at all temporal frequencies. They show that the modulation of the membrane potential of their model V1 cell falls off with increasing temporal frequency, with or without synaptic depression. The frequency at which the modulation reaches half of the peak appears to be somewhere between 10 and 20Hz. They say that this is simply due to the equivalent circuit model of their neuron. The degree of filtering that they observe simply from the membrane time constant is consistent with the fact that their membrane time constant is set to $30msec$, somewhat higher than is normally reported experimentally. This should provide more filtering than in our model cells, whose time constants during stimulus presentation fall from a rest level of $20msec$ to approximately $12msec$. In their model, including synaptic depression suppresses the modulation at low frequencies (less than 1Hz), due to the onset of depression kicking in before the input reaches its peak at low frequencies. Between 1 and 8Hz, depression does not have much of an effect, and at higher frequencies it causes the modulation to fall off slightly more rapidly than without depression. This is due to the fact that the synapses do not have time to recover fully between pulses of input at high frequencies. The fact that we do not see this effect in the range of frequencies that we studied, where LGN cells respond to input, is consistent with the fact that they use a longer recovery time constant of synaptic depression ($300msec$) than ours ($99msec$).

McCabe [1999] uses synaptic depression to explain similar shifts in temporal tuning between the thalamus and primary auditory cortex. Auditory cortical neurons respond only transiently to stimuli beyond 20Hz, even though their thalamic input neurons were able to follow stimuli up to approximately 200Hz. McCabe [1999] found that a circuit incorporating thalamocortical synaptic depression was able to reproduce this phenomenon, with an explanation similar to that described above: at high frequencies, the synapses did not have time to recover between pulses of input. Again, it is possible that we did not observe this effect of synaptic depression in our model because, for the range of frequencies we were studying, the time constant of recovery we used did not have much of an effect. Thus, the filtering we observed was dominated by the presence of NMDA mediated conductances.

2.5.4 Future Directions

While we have demonstrated that thalamocortical NMDA-mediated conductances are sufficient to account for the response of simple cells to simple temporal stimuli, it would be a strong test of the model to see if it was consistent with a number of experimental studies that have demonstrated that simple cells exhibit non-linear temporal summation in response to stimuli more complex than single sinusoids. A model that can account for the general temporal response properties of primary visual cortex should be able to explain these failures of linear summation.

Tolhurst et al. [1980] showed that the transient response to a two second flashed grating was too brief to accurately predict the temporal frequency tuning curve generated by the responses to the same standing grating whose contrast was modulated at different temporal frequencies. The predicted tuning curve, in general, underestimated the response to low frequencies relative to the response at high frequencies. Conversely, a linear prediction using the temporal frequency tuning curves would suggest that these cells should show a strong maintained response to a two second long flash, while in fact the flashed responses were dominated by a brief transient.

Dean et al. [1982] measured the response of simple cells to a stationary grating whose contrast was modulated by a waveform which was the sum of two sinusoids, at 1.25 and 7.75 Hz. The component of the response at 7.5 Hz was increased relative to the response to a stimulus modulated at that frequency alone, while the 1.25 Hz component of the response to the sum of sinusoids was reduced compared to the low frequency presented alone. They argued that this type of behavior can not be explained by a simple threshold non-linearity. The relative contrasts of the two components were adjusted to try to elicit similar responses to the single sinusoids on their own. With the single frequency responses matched in this way, a threshold non-linearity should increase each component equally in response to the compound stimulus, but that was not what was observed.

Finally, Reid et al. [1992] found similar properties using a broadband stimulus, generated by summing eight sinusoids. The lower frequency components of the response to the broadband stimulus were relatively reduced, as compared to those frequencies presented alone, while the higher frequency components were relatively increased as compared to the responses to single sinusoid gratings at high frequencies. Again, they argued that this behavior could not be accounted for based on a simple threshold non-linearity. They found that the integration time, or total effective delay of responses, measured as the slope of the line fitting the phase of the response to temporal frequency, was reduced with the sum of sinusoid stimulus as compared to single sinusoids.

Chance et al. [1998] use their model of synaptic depression to account for a number of these non-linear temporal properties of cortical cells. While we have performed some preliminary studies to determine if our model is consistent with some of these properties, further studies are required to ensure that the model can robustly account for them while preserving appropriate contrast-invariant orientation and temporal frequency tuning.

Furthermore, the model as it currently stands reverses the shift in temporal frequency tuning

towards high frequencies with increasing contrast that is observed in the LGN. There is only a small amount of experimental data demonstrating that cortical cells preserve, or perhaps even extend, the shift seen in the LGN, so it is not clear that the reversal that we observe is completely inappropriate. However, it would be useful to further study the model, to determine what mechanisms could be applied to preserve the contrast shift in temporal frequency tuning of the LGN while still maintaining an overall shift towards lower frequencies.

A potentially related problem is the fact that our cortical responses seem to be over rectified. They respond to less than half of the stimulus cycle, while simple cells are often observed to be closer to half-wave rectified, and sometimes even less than half-wave rectified [Tolhurst and Dean, 1990]. In general, the fact that the inhibitory dominance in our model causes all DC input from the LGN to be purely inhibitory seems to generate some of these putative problems. It is possible that a small phase-non-specific component of the feedback excitation, which connected cells of similar orientation but of all phases, could create some spillover in response phase, and solve the over-rectification problem.

Chapter 3

Constraining the Parameter Space

3.1 Introduction

The two parameter regimes mentioned at the end of the previous chapter – the strong inhibitory regime and the feedforward NMDA regime – both satisfy the experimental constraints of appropriate temporal frequency and orientation tuning. However they combine excitatory and inhibitory conductances in different ways to achieve that tuning. We have searched through the parameter space spanned by the levels of total intracortical excitatory and inhibitory synaptic strengths, with and without synaptic depression and with and without thalamocortical NMDA onto excitatory cells, in order to test how well the model satisfies various extracellularly and intracellularly determined experimental constraints in different parameter regimes. These constraints include:

1. Reasonably high firing rates to a preferred stimulus at high contrast.
2. Sufficiently narrow orientation tuning width of the spike rate response as compared to the tuning width of the intracellular voltage modulation.
3. The contrast-invariance of the width of the orientation tuning.
4. Sufficiently low cutoff frequency of the response to increasing temporal frequencies.
5. Reasonable levels of amplification of the voltage modulation due to the presence of the cortical circuit.
6. A small change ($\sim 20 - 40\%$) in the total conductance during stimulation with a null orientation stimulus.
7. Appropriate levels of total inhibitory and excitatory conductance.

Overall, it is possible to satisfy a remarkable number of these constraints simultaneously in a wide range of parameter space. For example, we originally assumed that the amplification factor

of the voltage modulation would have been a strong constraint on the maximum strength of both excitatory and inhibitory conductances; this did not turn out to be the case, as will be described below. The main constraint that favors the feedforward NMDA regime over the strong inhibitory regime is simply the total inhibitory conductance. Even the feedforward NMDA regime requires a level of inhibitory conductance that is higher than has been observed in a recent study by Anderson et al. [2000], but it is significantly closer, and we propose possible extensions of the model that could allow for lower total inhibitory conductances.

3.2 Experimental evidence for the various constraints

3.2.1 Firing Rate

There are a range of results in the literature with regards to the maximal firing rate of cat V1 simple cells in response to high contrast gratings at optimal orientation, spatial and temporal frequencies. Albrecht [1995] reported a population contrast response curve, measuring the amplitude of the first harmonic (F1) of the spike rate response, averaged over 116 cat V1 simple cells, that had a maximum saturated response of 40.2 spikes/sec. These curves were recorded at or near the optimal temporal frequency for each cell; the frequencies used ranged from 1-10Hz, and were between 4.0 and 8.0 Hz for 75% of the cells. At 80% contrast, the maximum contrast we use in our modeling, the response is close to the saturated response at $\approx$ 37 spikes/sec. Assuming a ratio of F1 to DC of 1.57 [Skottun et al., 1991], this results in a mean spike rate of $\approx$ 24 Hz. For our parameter searches, we will use a mean rate of 20 Hz as our lower bound.

3.2.2 Sharpening of Orientation Tuning: Iceberg Effect

A recent paper by Carandini and Ferster [2000] demonstrated that the orientation tuning width of the spike rate of simple cells is sharper than the tuning width of the membrane potential, a phenomenon that has been refered to as an "iceberg effect"; the shape of the portion of an iceberg that is above the surface of the water is usually much narrower than the shape of the full, mostly submerged, iceberg. They recorded intracellularly from 28 cortical cells (21 simple and 7 complex cells) in response to drifting sinusoidal gratings at preferred spatial and temporal frequencies (2 and 4Hz for the two individual cells shown), at various orientations. Orientation tuning curves were generated from the mean and amplitude of modulation (at the frequency of the stimulus) of both the firing rate and the membrane potential. They determined half-width at half height (HWHH) of these curves by fitting Gaussian functions to their recorded responses, and constrained the Gaussians for the mean and modulation data to have the same width; this procedure results in one value for the HWHH for spike rate and one for the membrane potential for each cell. They say that constraining the pairs of Gaussians in this way did not significantly worsen the fit, implying that, in general, the mean and modulation had similar orientation tuning for both experimental

measures. The mean HWHH for the membrane potential was 38 ± 15^o and for the firing rates it was 23 ± 8^o. The authors point out that this difference in tuning widths is probably even larger, given that the HWHH measure for firing rates was most likely overestimated. They sampled gratings at 30^o intervals, which prevented them from resolving orientation tuning curves with widths less than 17^o. Half of their cells reached this limit, and they simply assigned them an HWHH of 17^o, although the actual HWHHs for each of those cells could have been smaller.

The input receptive fields that we have used were set to match the tuning width of the membrane potential that Carandini and Ferster [2000] reported, as described in the methods section of the previous chapter. The degree of sharpening is determined mainly by the strength of the inhibitory connections, as described in Troyer et al. [1998]. The stronger the inhibition, the higher the effective spiking threshold and the sharper the orientation tuning of the spike rate. For our parameter searches, we will use 25^o HWHH as an upper bound for the orientation tuning of the spike rate.

3.2.3 Contrast-Invariance of Orientation Tuning

There are two studies on the contrast-invariance of orientation tuning that were referred to in the first chapter [Sclar and Freeman, 1982, Skottun et al., 1987]. Both were performed with gratings at optimal spatial and temporal frequencies. Neither study provides an ideal quantitative measure for comparing the degree of contrast invariance obtained in the model to what is observed experimentally. We have been using the coefficient of variation (CV, standard deviation divided by the mean) of the HWHHs of the orientation tuning curves obtained from mean firing rates over a range of contrasts. For the most part, the CV does not vary substantially over the region of parameter space for which the network is stable (usually less than 0.1, creeping up to 0.2-0.3 for parameter points that were on the edge of the stable region), so we will not be showing contours, but simply reporting range of values over the regions mapped.

3.2.4 Temporal Frequency Tuning

The experimental literature on the temporal frequency tuning of cortical cells was reviewed in detail in the previous chapter. While there are a range of high frequency cutoffs reported in the literature, in the previous chapter we referred to 6Hz as a reasonable estimate of what is generally observed. For our parameter searches in this chapter we have used a cutoff frequency of 8Hz as an upper bound.

3.2.5 Amplification Factor

One possible experimental contraint on the strength of the intracortical connections, both excitatory and inhibitory, comes from the cortical cooling experiments reported in Ferster et al. [1996]. Under the cooled condition, cortical cells could not fire action potentials, effectively eliminating the

intracortical circuitry. However, the layer 4 cells could still receive thalamocortical input, albeit in a reduced and slowed form. By recording intracellularly from these cells, and stimulating the LGN electrically, they were able to estimate the degree to which the cooling protocol dampened the thalamocortical input to these cells. By measuring the voltage response in the two conditions, they calculated the ratio of the amplitude of the monosynaptic EPSP in the warm condition to the amplitude in the cold condition: $\frac{R_{\text{electrical,warm}}}{R_{\text{electrical,cold}}}$. This ratio is their estimate of the reduction of the strength of the thalamocortical synapses under the cooling condition.

They also recorded from these cells during visual stimulation with drifting sinusoidal gratings. They found that the modulation of the membrane potential at the frequency of the grating was significantly reduced by the cooling protocol, more so than the reduction in response to the electrical stimulation. Under the assumption that cooling reduces the thalamic component of the visually induced response by the same factor as the reduction of the electrically evoked EPSP, they then used the ratio above to calculate the fraction of the normal visual response that is due to thalamic input. That is, $R_{\text{visual,warm(geniculatecomponent)}} = R_{\text{visual,cold}} \frac{R_{\text{electrical,warm}}}{R_{\text{electrical,cold}}}$. We define the amplification factor, A, to be the ratio of the amplitude of visually enduced modulation of the membrane potential with the intracortical circuitry intact, to the amplitude that is due to the geniculate input. That is

$$A = \frac{R_{\text{visual,warm}}}{R_{\text{visual,warm(geniculatecomponent)}}} \tag{3.1}$$

$$= \frac{R_{\text{visual,warm}}}{R_{\text{visual,cold}}} \frac{R_{\text{electrical,cold}}}{R_{\text{electrical,warm}}} \tag{3.2}$$

The mean value of A for the ten cells presented in Ferster et al. [1996] is approximately 2.7. For the parameter searches we have used an amplification factor of 3.5 as an upper bound.

Another measure of the degree of amplification has been estimated from experiments using electrically evoked silencing of the intracortical circuit [Chung and Ferster, 1998]. Cortical cells near the site of electrical stimulation experience a long lasting inhibition that prevents them from firing in response to visual stimulation. The authors recorded intracellularly from simple cells that received monosynaptic input from the LGN, and measured the amplitude of the evoked synaptic potential in response to a grating flashed during the period of electrically evoked suppression of the cortical circuitry. They found that amplitude of the response was reduced by 54%, as compared to the the amplitude under normal conditions. This suggests a degree of amplification of the LGN input by the cortical circuit by a factor of approximately 2. While this is clearly a very different effect – particularly since this was measuring the amplitude of the transient response to a brief stimulus, rather than the steady-state modulation to a drifting grating – it was a different type of measurement that providable a comparable estimate of the relative size of the intracortical to thalamocortical, visually evoked synaptic input to simple cells.

3.2.6 Total Conductance Change

There have been several experimental studies suggesting that the presence of a null stimulus does not significantly increase the total conductance of simple cells above background [Anderson et al., 2000, Hirsch et al., 1998, Borg-Graham et al., 1998, Berman et al., 1991, Douglas et al., 1991, 1988].

Although simple cells do not spike in response to null oriented gratings, the examination of the nature of LGN inputs to simple cells in the previous chapter makes it clear that there should be an increase in the amount of input from the LGN in response to the null grating; the phase non-specific, or unmodulated, component of the input is present at all orientations. Furthermore, the model predicts that there should be a compensating increase in intracortical inhibition to assure that the cells do not spike in response to this increase in feedforward input. Together, the increase in feedforward and inhibitory input would be observed as an increase in total conductance.

In order for this total increase in conductance to be small, the feedforward connections and the intracortical inhibitory connections can not be too strong. It turns out that this experimental constraint is not strict when the effect of synaptic depression is included, particularly at the low temporal frequencies that are normally used for these experiments. Short-term synaptic depression filters out a portion of the DC input from the LGN, as was shown above. At the lowest temporal frequencies, it filters out nearly all of the DC component, and should therefore eliminate most of the conductance increase that would otherwise be observed at low frequencies. This will be shown in the model in sections to follow.

3.2.7 Excitatory and Inhibitory Conductances

A recent study has used intracellular techniques to determine the total excitatory and inhibitory synaptic conductance changes experienced by cortical cells when responding to drifting sinusoidal gratings [Anderson et al., 2000]. They recorded the membrane potential of V1 cells during visual stimulation, and while injecting different levels of steady currents. The strength of the current-induced shifts in the membrane potential serves as a measure of the total conductance. To quantify this effect, a line can be fit to the relationship between membrane potential and level of injected current, for each point in time during a stimulus presentation: $V(t) = V_{\text{visual}}(t) + I_{\text{inj}}/g(t)$, where $V(t)$ are the measured voltage response, I_{inj} are the levels of injected current. The two fit parameters are $V_{\text{visual}}(t)$, the zero-injected-current intercept, and $g(t)$, which is the total conductance. By assuming that the conductance $g(t)$ is made up only of a rest conductance and synaptic excitatory and inhibitory conductances, all with known reversal potentials, it is possible to estimate the total excitatory and inhibitory conductances as a function of time, simply using the measurements of $V_{\text{visual}}(t)$ and $g(t)$.

With this technique they calculated the mean and modulations (F1 component at the frequency of the stimulus) of excitatory and inhibitory conductances impinging upon seven simple cells. These values, in response to a high contrast grating at the preferred orientation, are shown in table 3.1. For

our parameter searches, we have used 5nS as an upper bound for the mean excitatory conductance. The inhibitory conductances are in general, not in the range observed here, so instead we discuss what range of values are observed in the model in different parameter regimes.

	Mean Excit.	Modul. Excit	Mean Inhib.	Modul. Inhib.
Cell #1	4	10	6	10
Cell #2	2	1.3	6	1.5
Cell #3	.7	.4	1.5	1.8
Cell #4	7	6	13	8
Cell #5	1	.5	2	3.8
Cell #6	.5	1	2	5
Cell #7	1	.25	3	1

Table 3.1:
Anderson et al. [2000] Excitatory and Inhibitory Conductances Induced by High Contrast, Preferred Orientation Gratings
All values are in nS. These are estimated values, drawn by hand from Anderson et al. [2000] figure 17. Average values ($\pm$ 1 standard deviation) over the seven cells are as follows: Mean excitation is 2.3 ($\pm$ 2.4) nS. Mean inhibition is 4.8 ($\pm$ 4.1) nS. Modulation of excitation is (2.8 $\pm$ 3.8) nS. Modulation of inhibition is (4.4 $\pm$ 3.5) nS.

3.2.8 Summary

The experiments described above put multiple restrictions on the parameters that govern our model of visual cortex. Almost all of the experiments were performed with high contrast drifting gratings at or near the preferred temporal frequency of the cells being tested (except, of course, the temporal frequency tuning experiments). Since the precise temporal frequencies used were often not mentioned, or when they were mentioned they covered a range of values, we have chosen to use a 2Hz grating to test all of these various constraints, as it is near the preferred temporal frequency in our model, and is consistent with the range of frequencies that are used in these various experiments. As in the previous chapter, we measure the value of most of the various constraints averaged over the pool of cells in the preferred orientation bin. The exceptions are the change in conductance, which is averaged over the cells in the null orientation bin, and the constraints that are a function of orientation (the orientation tuning width and the contrast invariance of the tuning width), which are calculated over the range of orientation bins.

3.3 Parameter Searches

We ran the model of the previous chapter in a series of simulations over a range of parameter space in order to explore how the main parameters of the model affect the various experimental constraints described above. The two main free parameters are the excitatory strength and the inhibitory strength. The question to be addressed in this section is over how wide a range of these two parameters is it possible to satisfy all of the above constraints simultaneously.

The general behavior of the model will first be illustrated in a series of figures showing the results for this two-dimensional parameter search for the case of no synaptic depression and no thalamocortical NMDA, but with NMDA included in the feedback connections. (figures 3.1- 3.7). Each figure shows the iso-constraint contours for a particular experimental constraint. Depicted is a wide range of the space of excitatory and inhibitory strength, including a region with low inhibition and strong excitation for which the network goes out of control, at least at some temporal frequencies. We have defined the parameter regime where the network is unstable as the region in which the orientation tuning (halfwidth at halfheight) at any temporal frequency is broader than 50 degrees. This region is marked in white in the upper left corner of the following figures, bounded by the thick black contour. The "allowed" stable region, it seems, can extend indefinitely along the upper right diagonal, as increasing excitation is kept under control by increased inhibition.

3.3.1 Firing Rates

The black contours of figure 3.1 connect points with the same mean firing rate of cells at the preferred orientation. While the main effect of inhibition is to prevent firing of cells away from the preferred orientation, the small degree of overlap between excitatory and inhibitory currents at the preferred orientation, even at slow temporal frequencies, in this case 2Hz, causes the firing rate at the preferred orientation to decrease somewhat with increasing inhibitory strengths. Increasing the feedback excitation can counteract this effect of inhibition. Therefore, the contours slope up and to the right, with simultaneously increasing inhibition and excitation. The contours are closer and closer together with increasing levels of excitation, because of the feedback behavior of the amplifier. That is, the higher the firing rate, the stronger the effect of the feedback excitation, and the less the excitatory strength required to further increase the firing rate.

3.3.2 Sharpening of Orientation Tuning

The dark green contours in figure 3.2 connect points with matching orientation tuning widths, defined as having equal halfwidth at halfheight. Increasing inhibition effectively raises the firing threshold, restricting the region of responding cells to a narrower range of orientations. The main parameter, therefore, that determines the orientation tuning width is the inhibitory strength, and the contours connecting points of equal halfwidth are essentially vertical. However, if the feedback

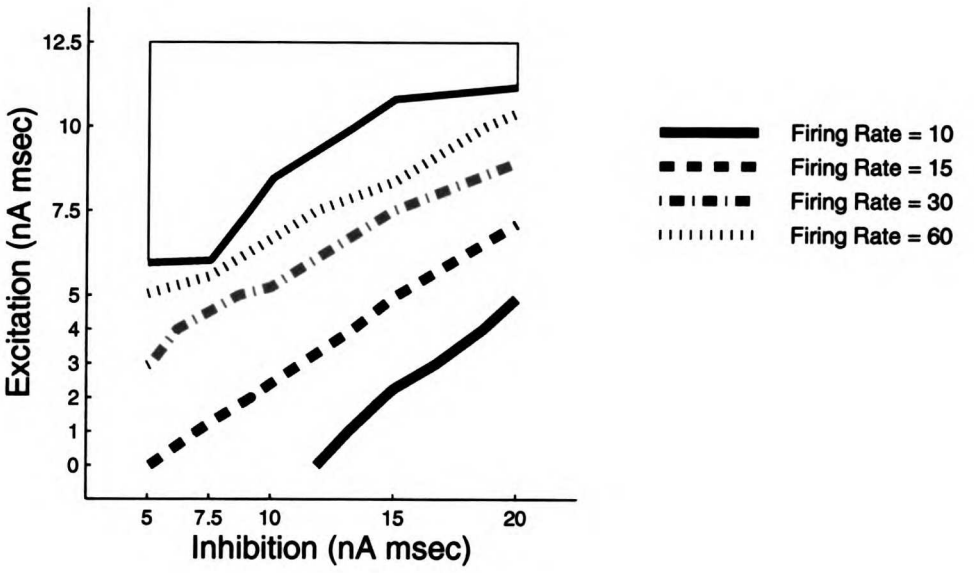

Figure 3.1:

Contour plots for mean stimulus-induced firing rates in Hz, varying total inhibitory and total intracortical excitatory strength. Each contour illustrates where in parameter space the mean firing rate, averaged over the cells in the preferred orientation bin, is equal to a particular value, as shown in the key. Responses are to a high contrast (80%), 2 Hz grating. The parameter regime where the network is unstable, defined as the region wherein the orientation tuning at any temporal frequency is broader than 50 degrees HWHH, is "whited-out," bounded by the thick black contour, in the upper left corner. Using no synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Default background parameters for no synaptic depression: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec.

excitation is too strong, the region that responds starts to spread out in orientation, broadening the tuning curve. If excitation is increased beyond a critical level, the feedback becomes unstable, and cells at all orientations respond. Increasing inhibition stabilizes the amplifier to some extent, raising the critical level of excitation that yields instability. That is why the the green contour starts to turn to the right at higher levels of excitation.

3.3.3 Temporal Frequency Cutoffs

The light blue contours of figure 3.3 connect the points with equal temporal frequency cutoffs, defined as the frequency at which the response is equal to 50% of the peak response. There are two ways of lowering the cutoff frequency: either by selectively suppressing response at high frequencies, or by selectively amplifying the response at low frequencies. For high levels of inhibition, increasing inhibition suppresses responses more strongly at high frequencies, where excitation and

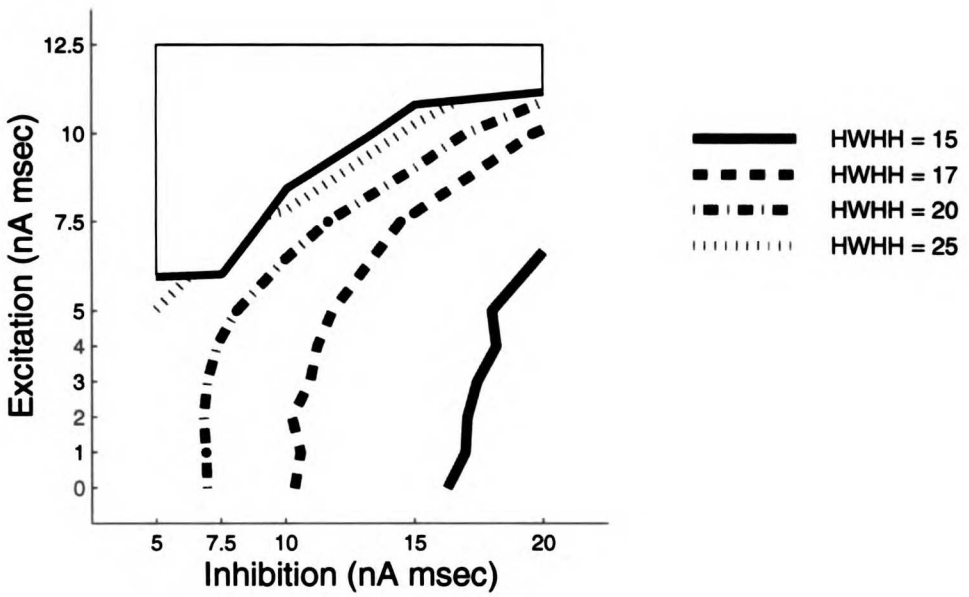

Figure 3.2:

Contour plots for half-width at half height (HWHH) of orientation tuning curves derived from mean stimulus-induced firing rates, varying total inhibitory and total intracortical excitatory strength. Each contour illustrates where in parameter space the HWHH is equal to a particular value, as shown in the key. Responses are to a high contrast (80%), 2 Hz grating. Unstable region, as defined in legend of figure 3.1, is "whited-out." Using no synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Default background parameters for no synaptic depression: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec.

inhibition are significantly overlapped, than at low frequencies. Thus, at high inhibitory levels, the cutoff frequency decreases with increasing inhibition. The cutoff frequency also tends to decrease with increasing excitation, because the feedback excitatory connections, which are predominantly mediated by NMDA, selectively amplify the low frequencies over the high frequencies. At low levels of excitation, where the effect of the amplifier is neglible, the contour is roughly vertical. But at higher levels of the amplifier, the temporal selectivity of the amplifier dominates, and the contours become horizontal. In fact, at high levels of excitation the contours start to slope slightly down to the left, indicating that when the amplifier is dominating the shape of the temporal frequency tuning curve, increasing inhibition can actually slightly raise the cutoff frequency, because it is reducing the effect of the amplifier.

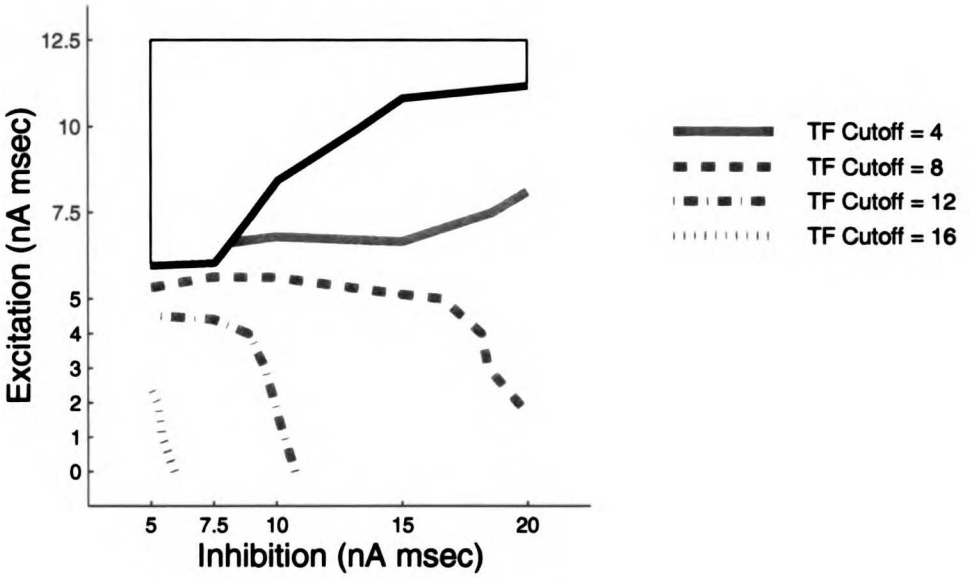

Figure 3.3:

Contour plots for temporal frequency cutoff (cycles/sec) derived from mean stimulus-induced firing rates averaged over the cells in the preferred orientation bin, varying total inhibitory and total intracortical excitatory strength. Each contour illustrates where in parameter space the cutoff frequency, defined as the frequency at which the response equals half of the maximal response, is equal to a particular value, as shown in the key. Responses are to a high contrast (80%) grating. Unstable region, as defined in legend of figure 3.1, is "whited-out." Using no synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Default background parameters for no synaptic depression: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec.

3.3.4 Null Conductance Change

The purple contours of figure 3.4 connect the points with equal total conductance change at the null orientation. This constraint is almost completely independent of the excitatory strength, since the excitatory cells at the null orientation are not responding at all. The exception is when excitation is so strong that the amplifier is unstable, and cells at all orientations are responding, as mentioned above. But, over the region of interest, this contour is vertical, and sets a maximum level of inhibition. Increasing inhibition increases the stimulus induced conductance at the null relative to background, increasing the conductance change due to the presence of the stimulus.

3.3.5 Amplification Factor

The orange contours of figure 3.5 connect the points with matching amplification factors of the voltage modulation due to the presence of the intracortical circuitry, at the preferred orientation. In

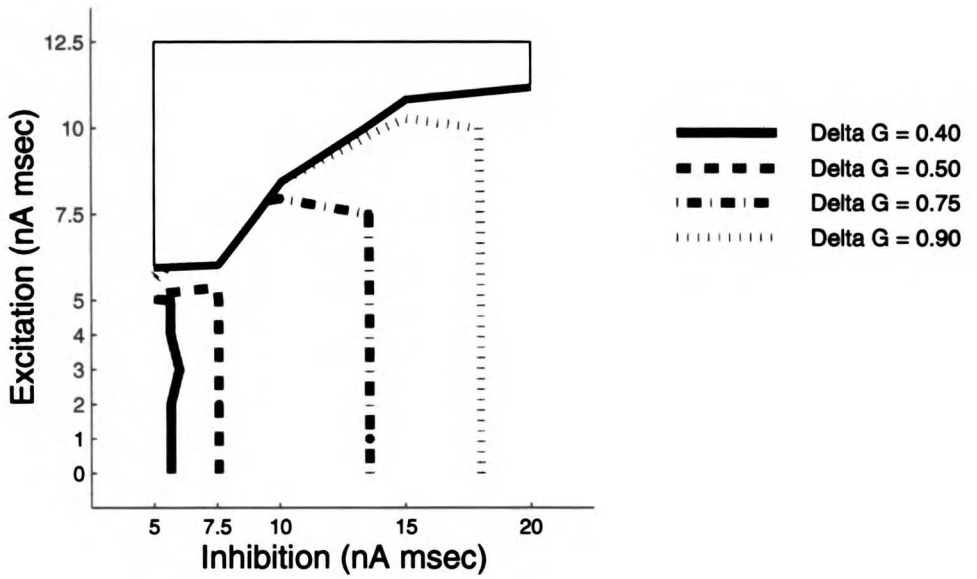

Figure 3.4:

Contour plots for total conductance change averaged over the cells in the null orientation bin, varying total inhibitory and total intracortical excitatory strength. Each contour illustrates where in parameter space the percentage change in total conductance, defined as the stimulus-induced conductance (subtracting background conductance) divided by the background conductance, is equal to a particular value, as shown in the key. Responses are to a high contrast (80%), 2 Hz grating. Unstable region, as defined in legend of figure 3.1, is "whited-out." Using no synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Default background parameters for no synaptic depression: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec.

general, the amplification factor is affected by both inhibition and excitation. Increasing inhibition amplifies the voltage modulation by deepening the hyperpolarizations while excitation raises the depolarizations. (Note: we measure amplification in terms of the "shadow voltage," the voltage trace a cell would have if it did not spike and simply integrated all the currents it receives, including currents due to spike-rate adaptation from the cell's actual spiking (see Methods of Chapter 2 (Model of NMDA conductances) for a full description). Thus, the spike threshold does not place an upper limit on depolarizations for this measurement. This is done because the somatic voltage in the intracellular measurements being modelled does not seem to be limited by a fixed spike threshold, presumably due to threshold increases (Na^+-channel inactivation) during spike responses and/or to voltage differences (resistance) between the soma and the spike-initiation zone in the axon hillock.) It would seem, therefore, that the amplification factor would be an effective experimental constraint to set maximum levels of both excitation and inhibition. It turns out, however, not to be the case. The hyperpolarizations saturate at the inhibitory reversal potential, which in

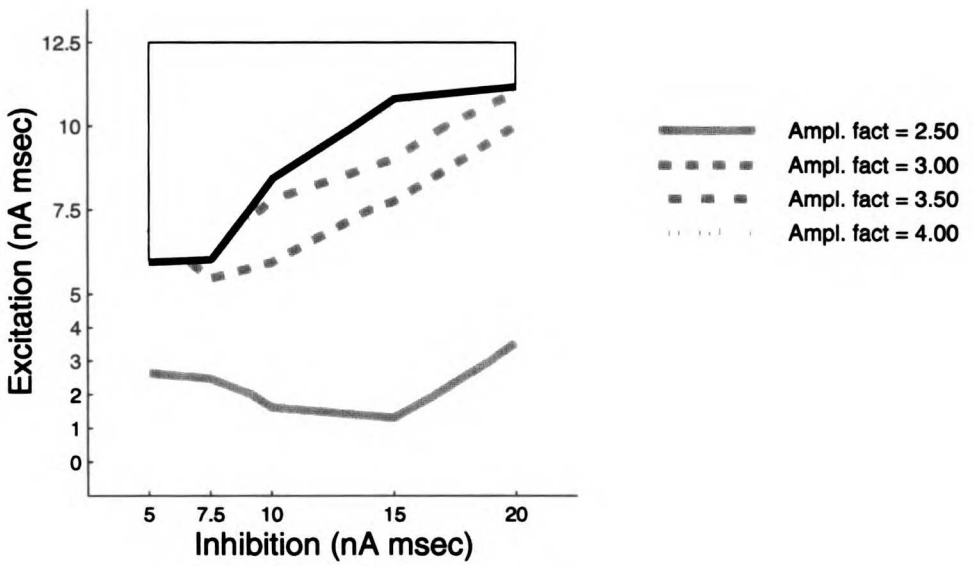

Figure 3.5:

Contour plots for voltage amplification factor averaged over the cells in the preferred orientation bin, varying total inhibitory and total intracortical excitatory strength. The amplification factor is defined as the amplitude of the first Fourier (F1) component of the membrane potential with cortical circuitry turned on divided by the F1 amplitude of membrane potential without cortical circuitry, due only to the thalamocortical connections. Each contour illustrates where in parameter space the amplification factor is equal to a particular value, as shown in the key. Responses are to a high contrast (80%), 2 Hz grating. Unstable region, as defined in legend of figure 3.1, is "whited-out." Using no synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Default background parameters for no synaptic depression: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec.

our case of purely GabaA inhibition is at -70mV, limiting the range over which inhibition can increase the amplification factor. Furthermore, excitatory firing rates within the region of interest are limited in how high they can go before the network becomes unstable. The total amount of feedback excitation is essentially the product of the excitatory firing rates and the $E \Rightarrow E$ connection strength. Therefore, the amplification factor contours are roughly parallel to the firing rate contours of figure 3.1 and do not clamp the parameters along that diagonal.

3.3.6 Excitatory and Inhibitory Conductances

We have also considered the absolute levels of excitatory and inhibitory conductance as a way of restricting the region of interest. The feedforward excitatory conductance from the LGN dominates over the feedback conductances in the region of stability, and the total mean excitatory conduc-

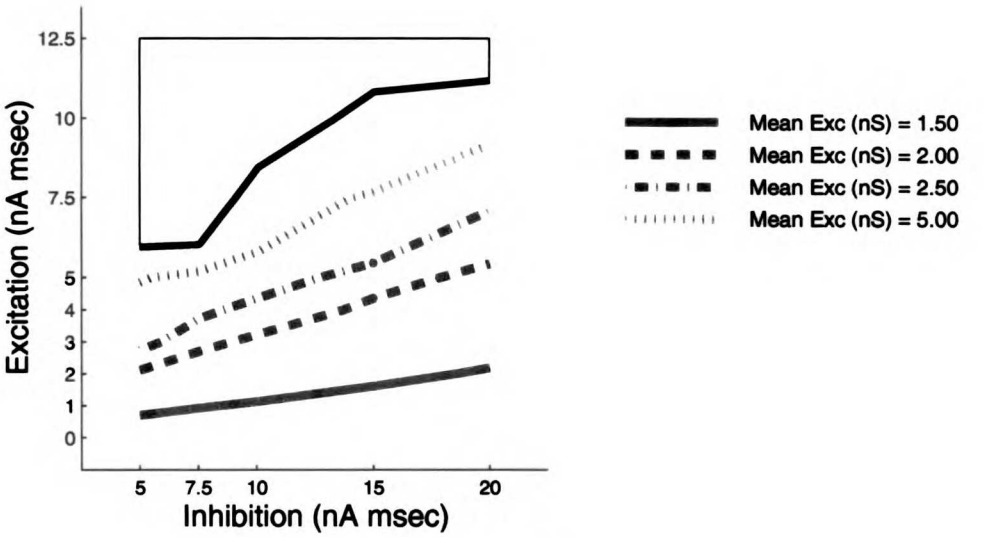

Figure 3.6:

Contour plots for mean stimulus-induced excitatory conductance, varying total inhibitory and total intracortical excitatory strength. Each contour illustrates where in parameter space the excitatory conductance, averaged over the cells in the preferred orientation bin, is equal to a particular value, as shown in the key. Responses are to a high contrast (80%), 2 Hz grating. Unstable region, as defined in legend of figure 3.1, is "whited-out." Using no synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Default background parameters for no synaptic depression: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec.

tance, therefore, does not vary significantly over the region of interest (figure 3.6). It increases with increasing intracortical excitatory strength, but stays within reasonable limits as long as the amplifier is maintained under control. The mean inhibitory conductance, however, increases dramatically over this region, and is therefore a strong constraint on the total inhibitory strength (figure 3.7).

3.4 Mapping regions of interest

3.4.1 No thalamocortical NMDA and no synaptic depression

We can now explore the various allowed regions of parameter space with no thalamocortical NMDA and no synaptic depression by combining sample contours of figures 3.1- 3.7 simultaneously, as shown in figure 3.8. The region with a sufficiently low temporal frequency cutoff is above and to the right of the light blue curve. With either a strong NMDA mediated amplifier, or very strong inhibition, the temporal frequency cutoff is sufficiently low. This region is consistent with having sharp orientation tuning (to the right of the green contour), low amplification factor (below the

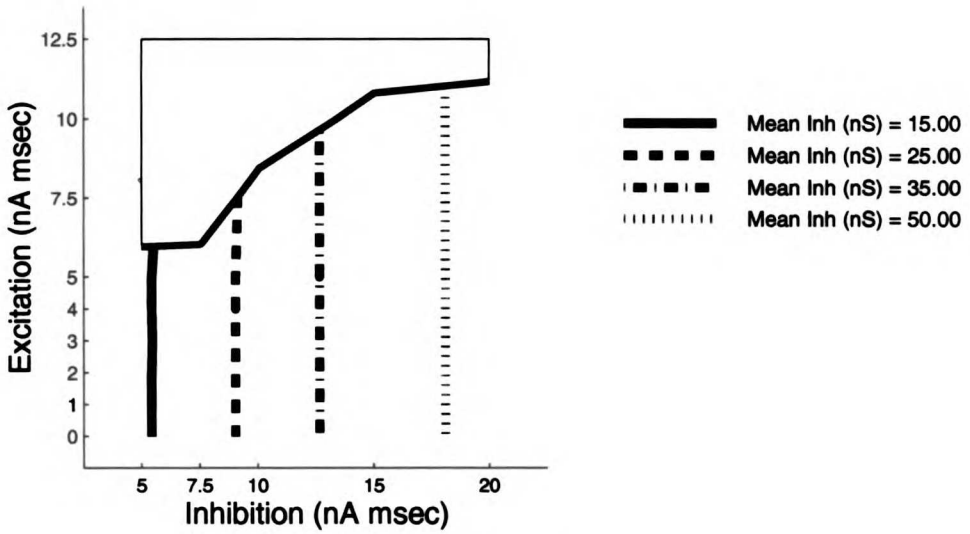

Figure 3.7:

Contour plots for mean stimulus-induced inhibitory conductance, varying total inhibitory and total intracortical excitatory strength. Each contour illustrates where in parameter space the inhibitory conductance, averaged over the cells in the preferred orientation bin, is equal to a particular value, as shown in the key. Responses are to a high contrast (80%), 2 Hz grating. Unstable region, as defined in legend of figure 3.1, is "whited-out." Using no synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Default background parameters for no synaptic depression: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec.

orange contour) and can also have low mean excitatory conductance while still having a sufficiently high mean firing rate (between the black and red contours). However, the entire region shown has a high mean inhibitory conductance and high change in null conductance, as illustrated by the 75% conductance change (purple) and 30 nS inhibitory conductance (blue) contours that cut through the center of this region. In order for the amplifier to be strong enough to give a low temporal frequency cutoff, the inhibition has to be very strong - too strong according to these conductance constraints.

3.4.2 No thalamocortical NMDA with synaptic depression

Including synaptic depression solves the issue of the null conductance change, but not of the inhibitory conductance at the preferred orientation (figure 3.9). As was discussed in the previous chapter, synaptic depression effectively eliminates much of the mean level of LGN input, particularly at low temporal frequencies (figure 2.28). When the mean level of LGN input is reduced, the conductance change at the null orientation is reduced along with it. Remapping the contours of

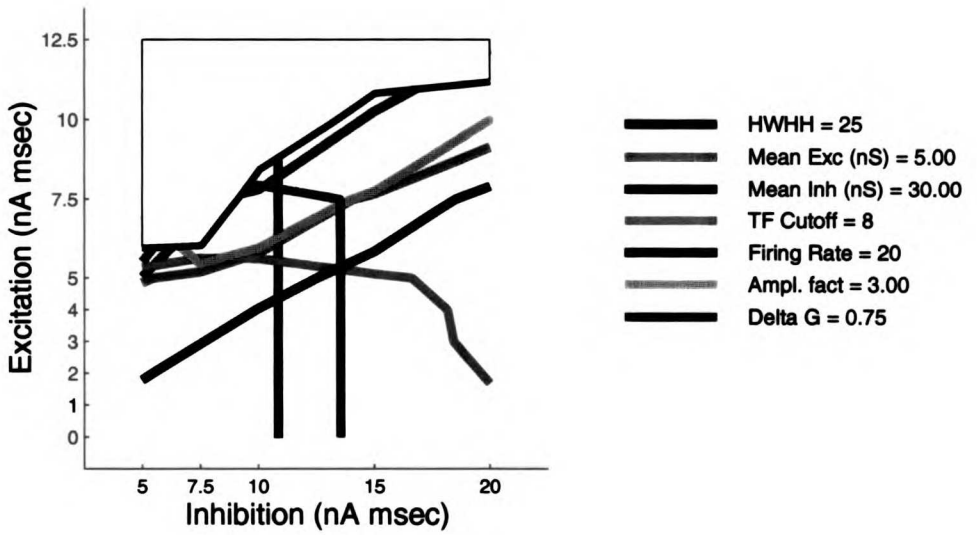

Figure 3.8:

Multiple contour plots for no depression and no thalamocortical NMDA. Each contour illustrates where in parameter space an experimental constraint has a particular value. All of the constraint parameters are determined in response to a high contrast (80%) grating, of the preferred orientation in all cases except for the conductance change, which is measured for the null orientation grating. All are at 2Hz temporal frequency except for the temporal frequency cutoff, for which the mean firing rate is measured at a range of temporal frequencies. Using no synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Default background parameters for no synaptic depression: $T = 6000$, $K = 1$, thalamocortical excitation = 5 nA msec. The range of values for the various constraints, including background excitatory and inhibitory rates and the contrast-invariance parameter (CV of HWHH over different contrasts) over the allowed region are as follows: BG Exc rates, max: 0.48, min: 0.00; BG Inh rates, max: 19.27, min: 19.18; HWHH CVs, max: 0.32, min: 0.04; HWHH, max: 23.94, min: 13.81; TF Cutoff, max: 17.12, min: 1.38; Firing Rate, max: 117.94, min: 6.51; Ampl. Fact, max: 3.88, min: 2.20; Mean Exc Cond., max: 18.83, min: 1.33; Mean Inh Cond., max: 55.24, min: 13.69; Null Cond. Change, max: 0.96, min: 0.34.

these various constraints, now including synaptic depression in both the thalamocortical and intracortical excitatory connections, we find that the change in null conductance, which is calculated at the relatively low frequency of 2 Hz, is no longer problematic (figure 3.9). Over this entire region of parameter space, the null conductance change varies only from 15 to 25%. However, the mean level of inhibitory conductance at the preferred orientation required to achieve a sufficiently low temporal frequency cutoff is still very high. Shown in blue of figure 3.9 is the 25nS contour. The other constraints can be satisfied simultaneously. The excitatory strength has to be fairly strong to achieve a low temporal frequency cutoff, since the intracortical synaptic depression has weakened

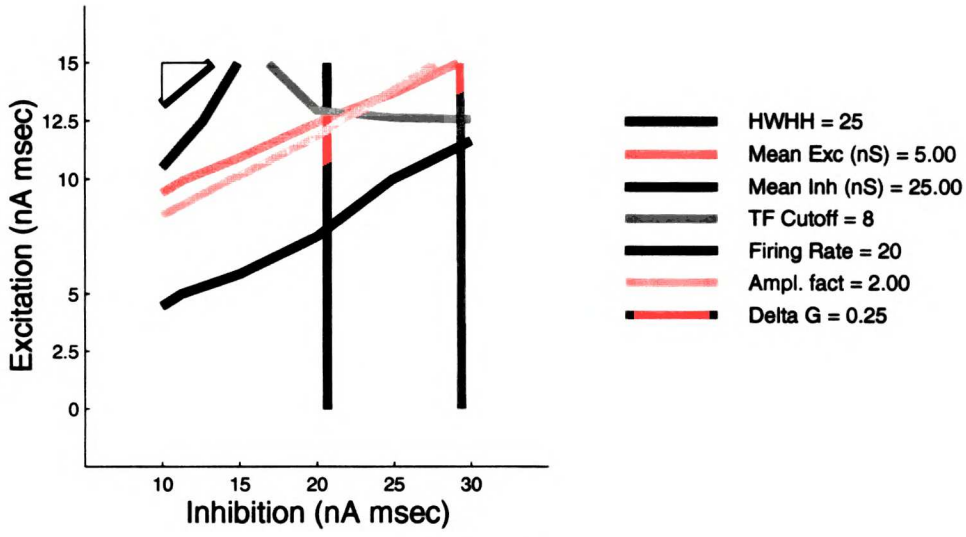

Figure 3.9:

Multiple contour plots with synaptic depression and no thalamocortical NMDA. Each contour illustrates where in parameter space an experimental constraint has a particular value. All of the constraint parameters are determined in response to a high contrast (80%) grating. All are in response to a grating at the preferred orientation except for the conductance change which is measured for the null orientation grating. All are at 2Hz temporal frequency, except for the temporal frequency cutoff, where the mean firing rate is measured at a range of temporal frequencies. Using synaptic depression, no thalamocortical NMDA, but with 95% NMDA in the intracortical excitatory connections. Unstable region, as defined in legend of figure 3.1, is "whited-out." Default background parameters for synaptic depression: $T = 4000$, $K = 8$, thalamocortical excitation = 20 nA msec. The range of values for the various constraints in the stable region are as follows: BG Exc rates, max: 0.75, min: 0.04; BG Inh rates, max: 29.57, min: 29.57; HWHH CVs, max: 0.10, min: 0.02; HWHH, max: 27.90, min: 14.48; TF Cutoff, max: 15.49, min: 5.42; Firing Rate, max: 92.00, min: 6.97; Ampl. Fact, max: 2.60, min: 1.46; Mean Exc Cond., max: 11.04, min: 0.67; Mean Inh Cond., max: 36.27, min: 12.09; Null Cond. Change, max: 0.25, min: 0.15.

the overall effect of the amplifier. But both the amplification factor, which is quite low with synaptic depression (ranging 1.46 to 2.83 over the entire region), and the total excitatory conductance are within reasonable limits.

3.4.3 Thalamocortical NMDA and no synaptic depression

By including NMDA in the thalamocortical to excitatory simple cell connections, all of these constraints can be satisfied simultaneously (figure 3.10). When including thalamocortical NMDA onto excitatory cells, much less inhibition is required to provide a low temporal frequency cutoff, as

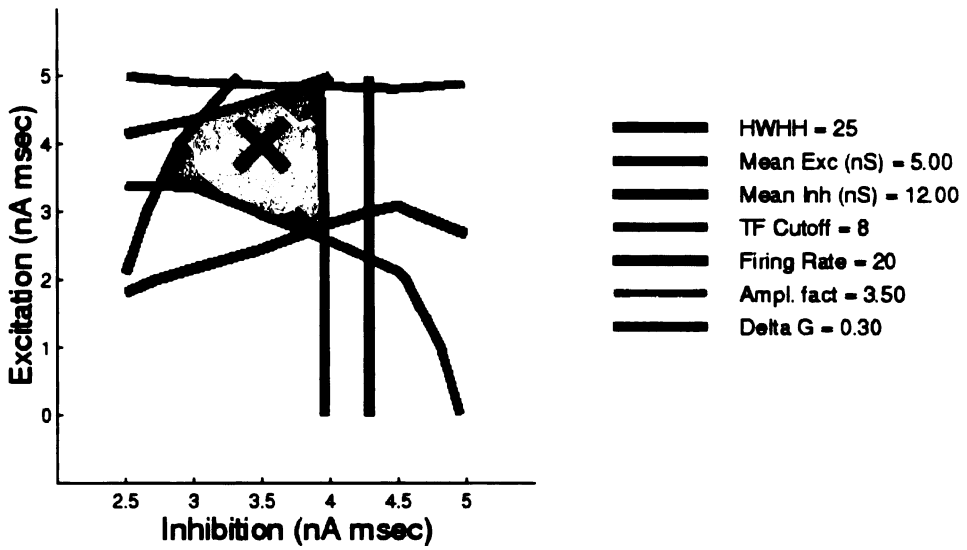

Figure 3.10:

Multiple parameter contour plots for no depression and thalamocortical NMDA onto E only. Default background and thalamocortical parameters for no depression (see methods). TC NMDA onto E cells = 90%. TC NMDA onto I cells = 0%. IC NMDA = 95%. Each contour illustrates where in parameter space an experimental constraint has a particular value. All of the constraint parameters are determined in response to a high contrast (80%) grating. All are in response to a grating at the preferred orientation except for the conductance change which is measured for the null orientation grating. All are at 2Hz temporal frequency, except for the temporal frequency cutoff, where the mean firing rate is measured at a range of temporal frequencies. Large black X indicated the default parameters for no depression, with amplifier, used in the previous sections. BG Exc rates, max: 0.92, min: 0.11; BG Inh rates, max: 19.64, min: 19.18; HWHH CVs, max: 0.19, min: 0.01; HWHH, max: 32.10, min: 18.33; TF Cutoff, max: 13.46, min: 1.79; Firing Rate, max: 92.51, min: 10.63; Ampl. Fact, max: 3.55, min: 2.35; Mean Exc Cond., max: 9.49, min: 1.23; Mean Inh Cond., max: 13.82, min: 6.91; Null Cond. Change, max: 0.36, min: 0.21.

was discussed in detail in the previous chapter. This reveals an "allowed" region (shaded in grey), bounded above by the mean excitatory conductance contour (red), to the left by the orientation tuning width contour (green) from below by the temporal frequency cutoff contour (light blue) and to the right by the null conductance change contour (purple). Note that this region also has mean inhibition below 12nS (dark blue contour), mean firing rate above 20Hz (black), and amplification factor below 3.5 (orange). The X marks the point that we have been considering our default parameter set, with the amplifier on and synaptic depression off. The mean inhibitory conductance is still larger than what was generally observed by Anderson et al. [2000], but it is significantly closer to the range of values that they observed than the regions shown in the previous figures,

without thalamocortical NMDA. Finally, the contrast invariance is maintained (CV of HWHH less than 0.2) throughout the entire region.

3.4.4 Thalamocortical NMDA with synaptic depression

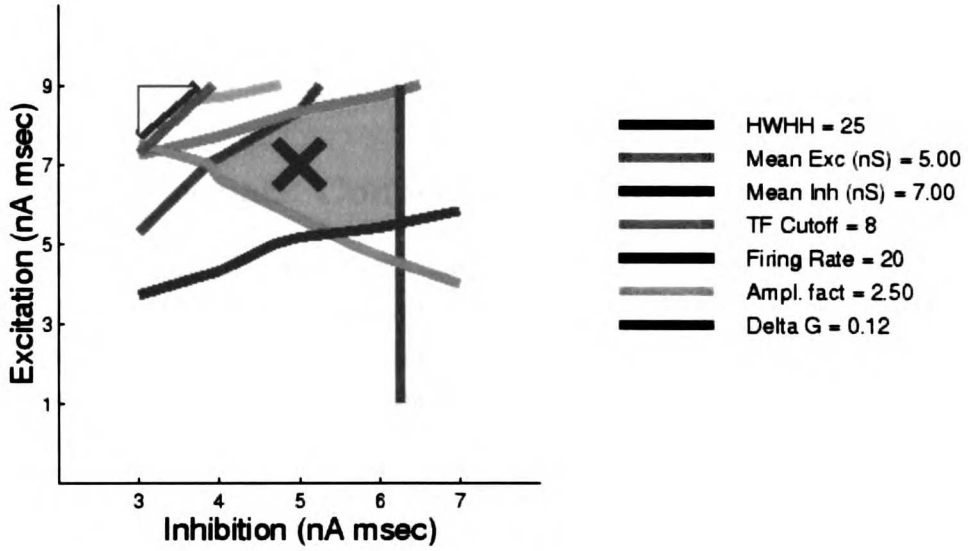

Figure 3.11:

Multiple parameter contour plot with synaptic depression and NMDA onto E only. Default background and thalamocortical parameters with depression (see methods). Each contour illustrates where in parameter space an experimental constraint has a particular value. All of the constraint parameters are determined in response to a high contrast (80%) grating. All are in response to a grating at the preferred orientation except for the conductance change which is measured for the null orientation grating. All are at 2Hz temporal frequency, except for the temporal frequency cutoff, where the mean firing rate is measured at a range of temporal frequencies. Large black X indicated the default parameters with depression and with amplifier, used in the previous sections. Unstable region, as defined in legend of figure 3.1, is "whited-out." BG Exc rates, max: 2.98, min: 0.39; BG Inh rates, max: 30.05, min: 29.57; HWHH CVs, max: 0.27, min: 0.04; HWHH, max: 31.06, min: 19.49; TF Cutoff, max: 13.33, min: 4.39; Firing Rate, max: 68.57, min: 9.26; Ampl. Fact, max: 2.56, min: 1.60; Mean Exc Cond., max: 7.12, min: 1.00; Mean Inh Cond., max: 7.82, min: 3.35; Null Cond. Change, max: 0.10, min: 0.06.

We can similarly find a region that satisfies all of the above constraints with NMDA in the thalamocortical connections to excitatory cells and synaptic depression (figure 3.11). As was the case when adding depression without thalamocortical NMDA, the conductance change becomes very small at a temporal frequency of 2Hz. Aside from the one data point with the lowest inhibition and highest exctiation that is "whited-out" in the unstable region, the conductance change is less

than 10% throughout the entire region mapped. The "allowed" region (shaded in grey), is bounded above by the mean excitatory conductance contour (red), to the left by the orientation tuning width contour (green) from below by the temporal frequency cutoff contour (light blue) and the firing rate contour (black) and to the right by the mean inhibitory contour at 7nS (dark blue). The mean inhibitory conductances are even closer to the range of values observed by Anderson et al. [2000] in this case, with synaptic depression and thalamocortical NMDA. Furthermore, the amplification factor is low throughout, as is the contrast-invariance measure. The X marks the point that we have been considering our default parameter set, with the amplifier and synaptic depression on.

3.5 Effect of Feedback Connections to Inhibitory Cells

There are two sets of cortical connections that would complete the correlation-based circuit described above that have not yet mentioned: the connections from excitatory cells to inhibitory cells of the same phase, and the connections between inhibitory cells of opposite phase. These connections have not been included in the modeling of the previous chapter; but the effects of these connections will be discussed here.

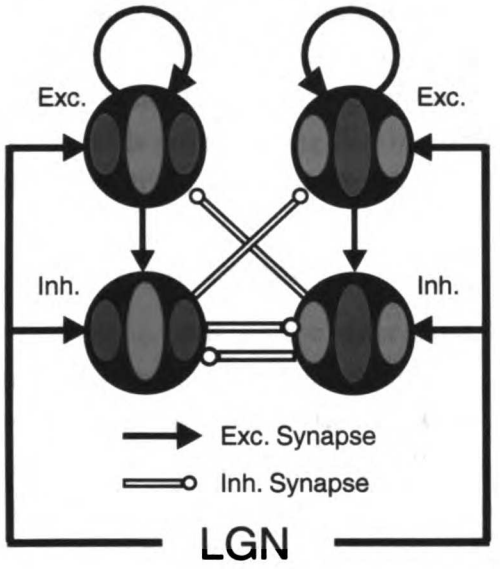

Figure 3.12:
Cartoon of the full correlation-based-connectivity circuit, including the feedback connections to inhibitory cells: same-phase E⇒I connections and anti-phase I⇒I connections.

3.5.1 E⇒I connections

The E⇒I connections have a minimal effect on the behavior of the network. They simply amplify the inhibition onto excitatory cells of one phase when excitatory cells of the opposite phase are active. Since excitatory cells of the same orientation and two opposite phases are never simultaneously active in this network, the effect is simply to increase the inhibition onto cells that are already

inactive. Thus, adding E⇒I connections causes essentially no change in spike responses of excitatory cells (see Chapter 1).

However, there are a few effects of the E⇒I connections worth mentioning. First, they alter the spike responses, and in particular the orientation tuning, of the inhibitory cells. Without the E⇒I connections, the mean response of the inhibitory cells is essentially untuned for orientation, following the untuned DC component of the LGN input. The E⇒I connections increase the activity of the inhibitory cells at the orientations to which the corresponding excitatory cells respond. This confers upon their orientation tuning a tuned hump around the preferred orientation. This effect will be discussed in the following chapter.

The E⇒I connections also make it more difficult to satisfy the mean inhibitory conductance constraint. Figure 3.13 is a repeat of the parameter search without synaptic depression and with thalamocortical NMDA onto excitatory cells only (figure 3.10), but now including the E⇒I connections. Many of the contours have not changed positions much at all; however the mean inhibitory conductance contour (dark blue) bends to the left as intracortical excitatory strength increases, rather than being vertical and independent of excitatory strength. As the amplifier increases in strength, the excitatory cells respond more strongly and provide stronger excitation to the inhibitory cells of the same phase and orientation. Thus, at their preferred orientation and phase, the inhibitory cells respond more strongly and provide stronger modulating inhibition onto the excitatory cells of the opposite phase. Furthermore, with the E⇒I connections turned on, the inhibitory cells are rectifying (hitting zero firing rate during the troughs of the modulations of their LGN input). This also increases the mean level of inhibition received by the excitatory cells.

3.5.2 I⇒I connections

The I⇒I connections, on the other hand, have a dramatic effect on the network. Full strength I⇒I connections, that is, equal to the strength of the I⇒E connections, can not be tolerated. They create a "winner take all" competition between inhibitory cells of different phases at the null orientation, which then disinhibits excitatory cells of the opposite phase of the loser. With the amplifier turned on, this makes the network become unstable, even with weak levels of the I⇒I connections. Very weak I⇒I connections can be tolerated by the network, and essentially weaken the effective I⇒E strengths. Figure 3.14 is another repeat of the parameter search shown in figure 3.10, but now including anti-phase I⇒I connections that are only 5% of the strength of the I⇒E synapses. The contours all look very similar to those in figure 3.10, only having translated the whole figure slightly to the right; implying that the intracortical inhibitory strength needs to be slightly stronger to have the same effect as without the weak I⇒I connections.

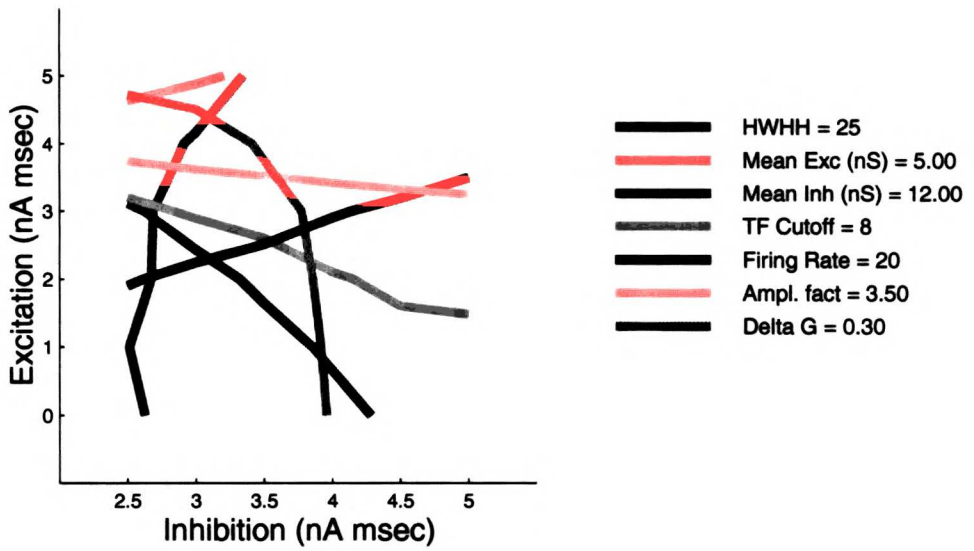

Figure 3.13:

Multiple parameter contour plots, including E⇒I connections, with no depression and thalamocortical NMDA onto E only. Default background and thalamocortical parameters for no depression (see methods). TC NMDA onto E cells = 90%. TC NMDA onto I cells = 0%. IC NMDA = 95%. Each contour illustrates where in parameter space an experimental constraint has a particular value. All of the constraint parameters are determined in response to a high contrast (80%) grating. All are in response to a grating at the preferred orientation except for the conductance change which is measured for the null orientation grating. All are at 2Hz temporal frequency, except for the temporal frequency cutoff, where the mean firing rate is measured at a range of temporal frequencies. BG Exc rates, max: 0.78, min: 0.10; BG Inh rates, max: 21.73, min: 19.58; HWHH CVs, max: 0.00, min: 0.00; HWHH, max: 31.26, min: 17.46; TF Cutoff, max: 13.87, min: 2.51; Firing Rate, max: 53.14, min: 10.29; Ampl. Fact, max: 4.21, min: 2.37; Mean Exc Cond., max: 5.75, min: 1.26; Mean Inh Cond., max: 29.61, min: 7.00; Null Cond. Change, max: 0.39, min: 0.21.

3.6 Intracellular Constraints that Vary with Temporal Frequency

3.6.1 Amplification Factor

If thalamocortical synapses onto excitatory cells but not inhibitory cells are mediated significantly by NMDA, the model predicts a dramatic increase in the amplification factor described above with increasing temporal frequency. With increasing temporal frequency, the excitation-induced voltage modulation is suppressed by the NMDA. However the inhibition-induced modulation will not be suppressed at high frequencies, creating a much larger ratio of full-circuit F1 over LGN-induced F1 (figure 3.15). A summary of the prediction of the amplification factor versus temporal frequency,

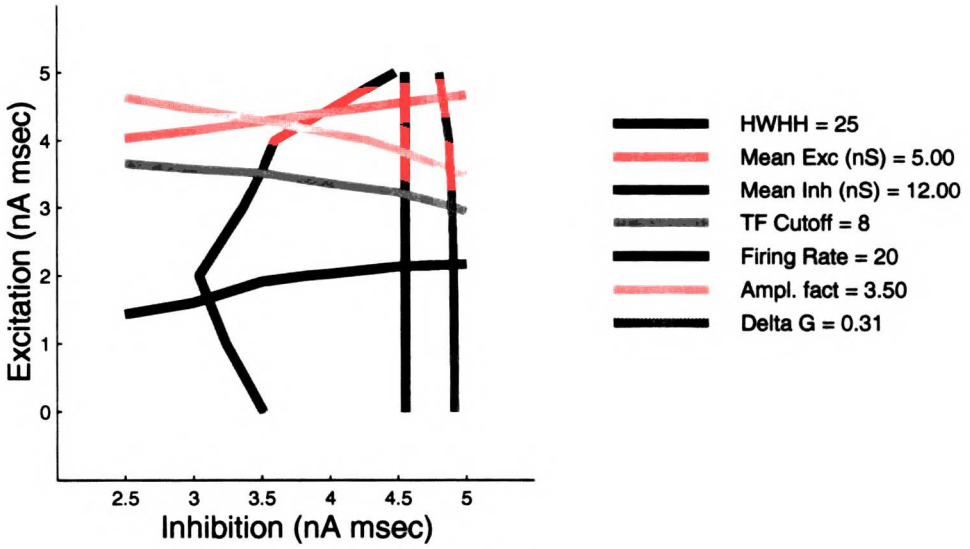

Figure 3.14:

Multiple parameter contour plots, including very weak I⇒I connections (5% of the strength of the I⇒E connections), with no depression and thalamocortical NMDA onto E only. Default background and thalamocortical parameters for no depression (see methods). TC NMDA onto E cells = 90%. TC NMDA onto I cells = 0%. IC NMDA = 95%. Each contour illustrates where in parameter space an experimental constraint has a particular value. All of the constraint parameters are determined in response to a high contrast (80%) grating. All are in response to a grating at the preferred orientation except for the conductance change which is measured for the null orientation grating. All are at 2Hz temporal frequency, except for the temporal frequency cutoff, where the mean firing rate is measured at a range of temporal frequencies. BG Exc rates, max: 1.10, min: 0.20; BG Inh rates, max: 18.71, min: 17.82; HWHH CVs, max: 0.00, min: 0.00; HWHH, max: 51.09, min: 23.07; TF Cutoff, max: 14.30, min: 1.90; Firing Rate, max: 120.11, min: 12.91; Ampl. Fact, max: 4.02, min: 2.41; Mean Exc Cond., max: 12.32, min: 1.29; Mean Inh Cond., max: 13.08, min: 6.74; Null Cond. Change, max: 0.32, min: 0.20.

with the standard set of parameters and with NMDA onto the excitatory cells only, is shown in figure 3.16.

However, for the case in which thalamocortical synapses onto inhibitory but not excitatory cells are mediated by NMDA, the model predicts the opposite trend for the amplification factor with increasing temporal frequency (figure 3.17). In this case, the feedforward modulation is not supressed with increasing temporal frequency, because the LGN synapses onto excitatory cells are AMPA mediated and are able to follow the modulation of the stimulus out to high frequencies. The cortical inhibition, on the other hand, can not follow the high frequency modulation. In fact, the amplification factor can actually go to less that 1 at high frequencies, because the DC

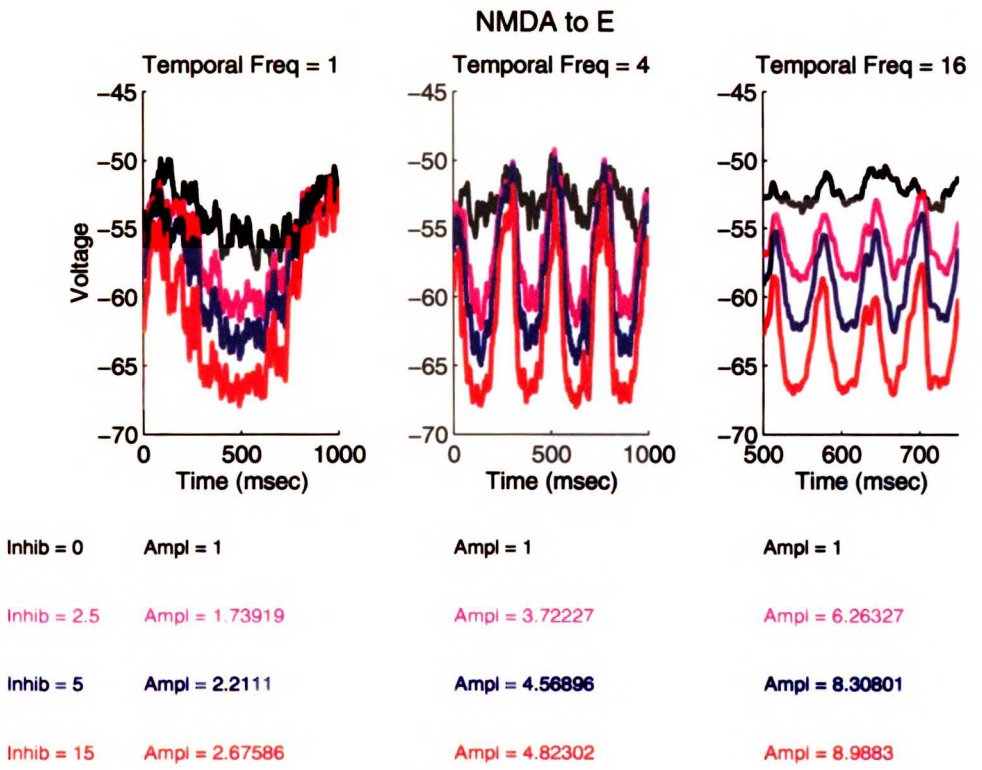

Figure 3.15:

Traces showing voltage modulation for different temporal frequencies and different levels of inhibition, with NMDA synapses included in thalamocortical connections onto excitatory but not inhibitory cells, and with no intracortical excitation. Below graphs are shown inhibitory strength (nA-msec) corresponding to each color, and the inhibition-induced amplification of the F1 of the voltage relative to the zero-inhibition case.

of the inhibition, which increases with temporal frequency, moves the voltage traces closer to the inhibitory reversal potential, leaving less room for the modulatory troughs at high frequencies.

3.6.2 Conductance Change at Null Orientation

The conductance change at the null is predicted to follow the DC of the LGN input, which increases with temporal frequency (figure 3.18). As described earlier, the effect of synaptic depression is to suppress the DC at all frequencies by a roughly subtractive mechanism. The conductance change, therefore, still increases with frequency, but it is extremely small at low frequencies where conductance measurements are usually performed.

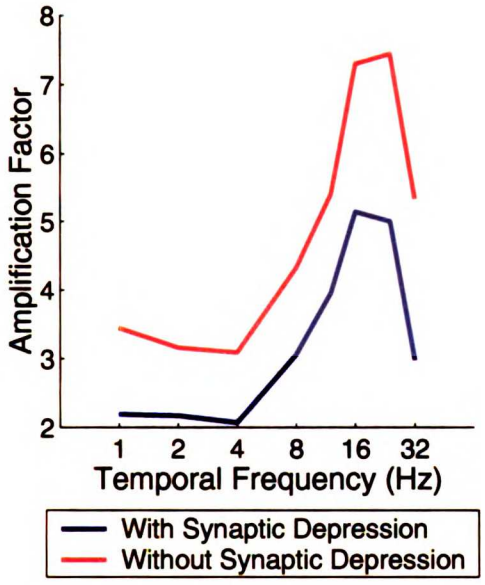

Figure 3.16:

Amplification factor of modulation of membrane potential due to intracortical circuitry plotted as a function of temporal frequency, with and without synaptic depression, with thalamocortical synapses having 90% NMDA onto excitatory cells and 0% NMDA onto inhibitory cells. Intracortical excitatory synapses have 95% NMDA. Default parameters: without depression, thalamocortical excitation = 5 nA msec, inhibition = 3.5 nA msec, intracortical excitation = 4 nA msec; with depression, thalamocortical excitation = 20 nA msec, inhibition = 5 nA msec, intracortical excitation = 7 nA msec.

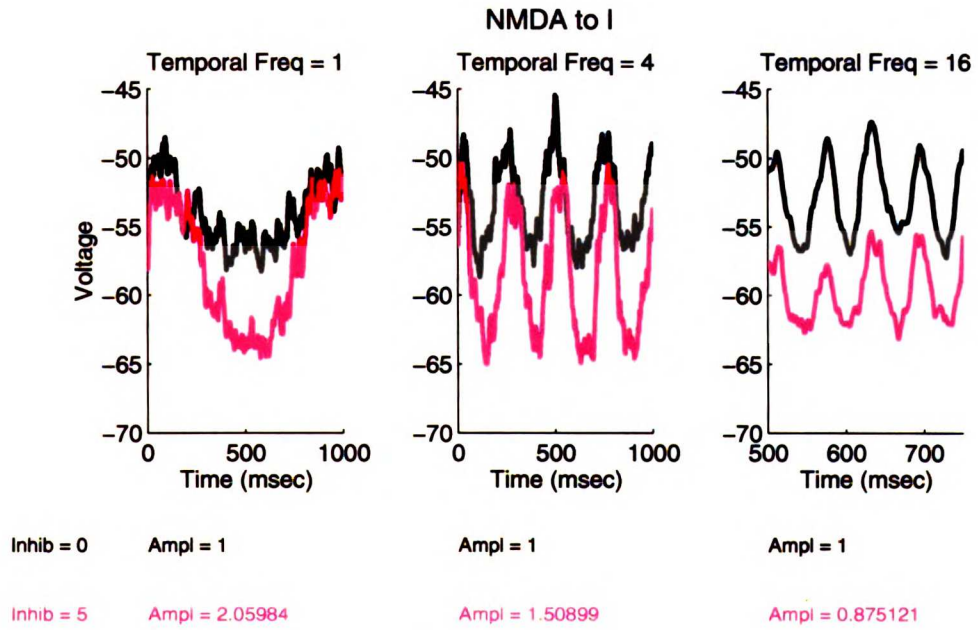

Figure 3.17:

Traces showing voltage modulation for different temporal frequencies and different levels of inhibition, with NMDA synapses included in thalamocortical connections onto inhibitory but not excitatory cells.

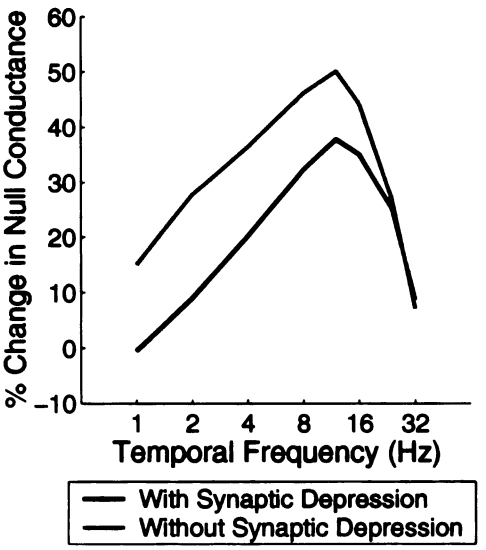

Figure 3.18:

Total conductance change at null orientation plotted as a function of temporal frequency, with and without synaptic depression, with 90% NMDA onto excitatory cells. Default parameters: without depression, thalamocortical excitation = 5 nA msec, inhibition = 3.5 nA msec, intracortical excitation = 4 nA msec; with depression, thalamocortical excitation = 20 nA msec, inhibition = 5 nA msec, intracortical excitation = 7 nA msec.

3.7 Discussion

From the parameter searches described above, we have shown that while both parameter regimes mentioned at the end of the previous chapter – the strong inhibitory regime and the feedforward NMDA regime – satisfy a number of experimental constraints, they can be distinguished by the total level of inhibitory conductance. The strong inhibitory regime requires very high levels of inhibitory conductance that are well out of the range that has been observed from intracellular experimental measurements. However, even the feedforward NMDA parameter regime requires levels of inhibitory conductance that are somewhat larger than what has been observed.

One possible method for satisfying the set of constraints described in this chapter with lower total levels of inhibitory conductance is by including Gaba-B-mediated inhibitory conductances. Gaba-B receptors are linked to K^+ channels that have a more hyperpolarized reversal potential than Gaba-A Cl^- receptors, and therefore require less conductance to achieve the same inhibitory current. Furthermore, the slow time constant of Gaba-B-mediated IPSCs, as compared to Gaba-A-mediated IPSCs, should be more effective at demodulating higher temporal frequencies, and therefore demand less mean inhibitory current to achieve comparable shifts in temporal frequency tuning. We have performed some preliminary tests of the model with Gaba-B included, and have demonstrated that many constraints can be met with lower mean inhibitory conductances. However, the presence of Gaba-B also suppressed the mean firing rates of the excitatory cells below the experimentally determined levels. More thorough parameter searches that include Gaba-B need to be done to test whether or not all of the experimental constraints can be satisfied simultaneously with weaker mean inhibitory conductance.

Chapter 4

Inhibitory Cell Orientation and Temporal Frequency Tuning

4.1 Introduction

The response properties of a population of layer 4 inhibitory cells to gratings of different orientations and contrasts was a major prediction of our study on the contrast invariance of excitatory cell orientation tuning in the first chapter. These inhibitory cells were predicted to be tuned for orientation, but to respond in a contrast-dependent manner to gratings at the null orientation. The latter property was required by the model in order to counteract the excess mean thalamocortical excitatory input onto cortical excitatory cells at the null orientation. In light of the previous chapters' results on the dependence of the model on different parameters, this prediction has become more subtle. In the second chapter, we demonstrated that synaptic depression can filter out a significant portion of the mean input, particularly at lower temporal frequencies. As a result, the inhibitory cells will not respond nearly as strongly to the null orientation at low temporal frequencies. Furthermore, their mean spike rate is not well tuned to orientation without the intracortical excitatory-to-inhibitory connections, which do not otherwise significantly affect the behavior of the model, as discussed in the previous chapter. These connections also partially impart the temporal frequency tuning of the excitatory cells onto the inhibitory cells. The mean response of inhibitory cells otherwise closely follows the temporal tuning of the LGN inputs, regardless of whether NMDA is present in their thalamocortical connections.

4.2 Orientation Tuning of Inhibitory Cells

The inhibitory cells in the model have a sufficiently high firing response to the mean LGN input that they tend not to rectify, even in response to an optimal stimulus that evokes a strong modulation of the membrane potential. As a result, the mean firing rate of the inhibitory cells closely follows

the mean level of total LGN input, largely ignoring the magnitude of the modulation of the total LGN input. In particular, the mean spike rates of inhibitory cells will not be tuned to orientation, just as the mean LGN input is not orientation tuned (figure 4.1). At low contrast (10%), where the mean LGN input does not increase significantly above background, the inhibitory cells respond at their background firing rate at all orientations. But, without synaptic depression, they increase their response with increasing levels of contrast at all orientations, and at all temporal frequencies that the LGN responds to, again, following the mean level of LGN input.

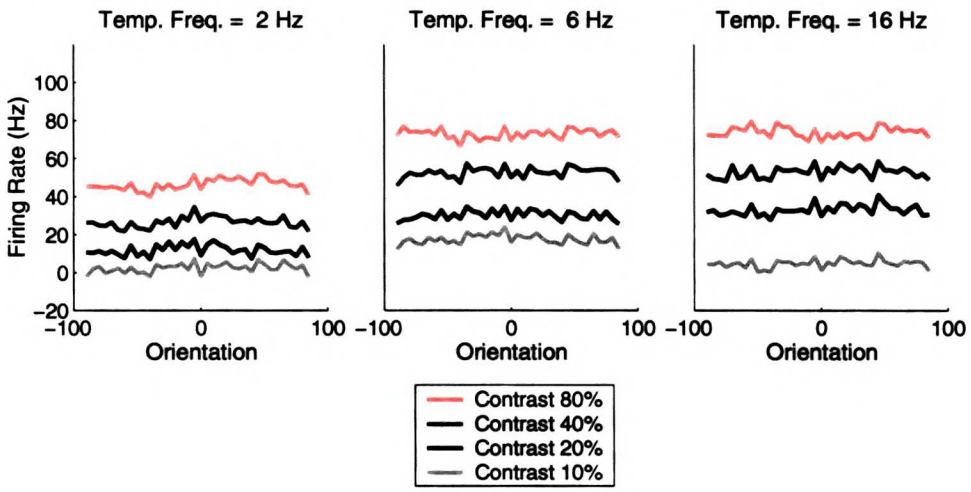

Figure 4.1:

Orientation tuning of mean firing rates of model inhibitory cells, without synaptic depression and with no E⇒I connections. 90% NMDA onto excitatory cells. Plotted is stimulus-induced mean firing rate, that is, mean firing rate averaged over cells in each orientation bin, less the background mean firing rate averaged over cells in each bin. The different colors represent different contrasts (80, 40, 20 and 10% as defined in the key), and the different plots are at three temporal frequencies (2, 6 and 16 Hz from left to right). Default parameters: without depression and with amplifier: thalamocortical excitation = 5 nA msec, inhibition = 3.5 nA msec, intracortical excitation = 4 nA msec.

When feedback connections from excitatory cells to inhibitory cells are included in the model (the E⇒I connections described in the previous chapter), the inhibitory cells inherit sensitivity to the orientation of the stimulus from the cortical excitatory cells (figure 4.2). They still respond at the null orientation in a contrast dependent manner, following the mean LGN input. This was the response pattern that was described in the first chapter, in which the default parameter conditions included the E⇒I connections. As we have described in the preceeding chapter, there are very few other features of the model that are sensitive to the presence or absence of these connections.

Synaptic depression reduces the mean level of LGN input, particularly at low temporal frequencies, as shown in the second chapter in figure 2.28. As a result, including synaptic depression

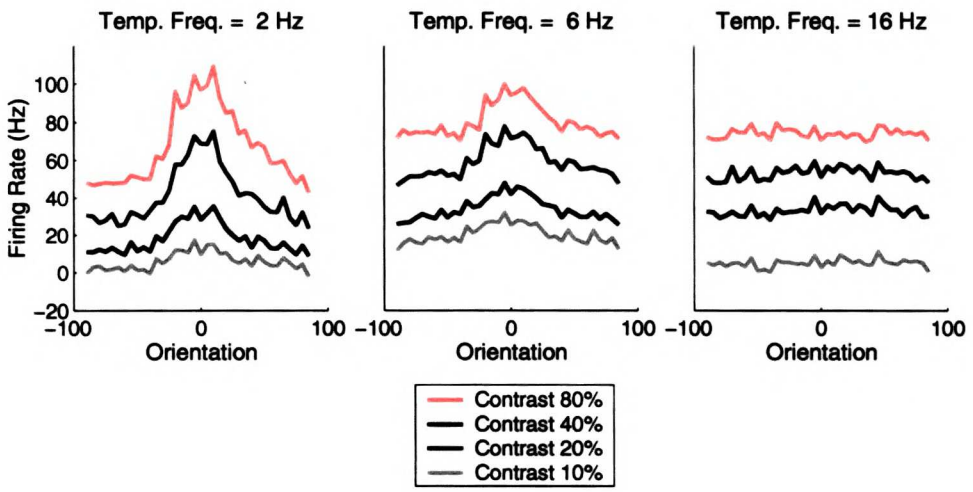

Figure 4.2:

Orientation tuning of mean firing rate of model inhibitory neurons, without synaptic depression and with E⇒I connections. 90% NMDA onto excitatory cells. Plotted is stimulus-induced mean firing rate, that is, mean firing rate averaged over cells in each orientation bin, less the background mean firing rate averaged over cells in each bin. The different colors represent different contrasts (80, 40, 20 and 10% as defined in the key), and the different plots are at three temporal frequencies (2, 6 and 16 Hz from left to right). Default parameters: without depression and with amplifier: thalamocortical excitation = 5 nA msec, inhibition = 3.5 nA msec, intracortical excitation = 4 nA msec.

reduces the response of the inhibitory cells, such that at low frequencies they respond at close to background levels, even at high contrast (figure 4.3). At 2Hz, there is still some firing above background at 80% contrast, at all orientations, but it is significantly reduced from the firing rates without depression. Without the E⇒I connections, the inhibitory cells are not sensitive to orientation, as was the case without depression. However, with the E⇒I connections, the inhibitory cells again inherit the orientation tuning of the excitatory cells. In fact, the orientation tuning is stronger with depression, particularly at lower frequencies where much of the untuned mean input is suppressed.

4.3 Temporal Frequency Tuning of Inhibitory Cells

A similar transformation occurs upon including the E⇒I connections when considering the temporal frequency tuning of the inhibitory cells. Without the E⇒I connections, the inhibitory cells follow the temporal tuning of the LGN inputs almost perfectly without synaptic depression (figure 4.5, left). With synaptic depression, the match is very close, except at the lowest temporal frequencies, where the suppression of the DC is stronger than at higher frequencies (figure 4.6, left).

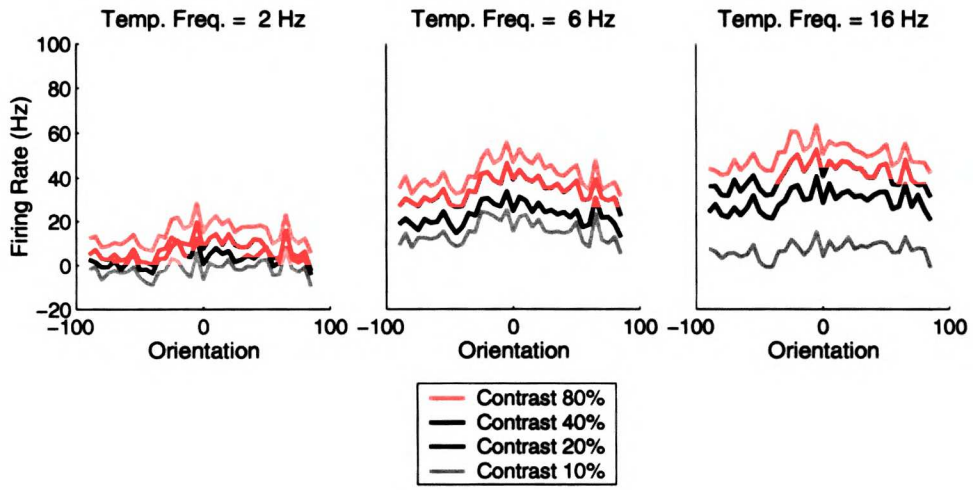

Figure 4.3:

Orientation tuning of mean firing rate of model inhibitory neurons, with synaptic depression and with no E⇒I connections. 90% NMDA onto excitatory cells. Plotted is stimulus-induced mean firing rate, that is, mean firing rate averaged over cells in each orientation bin, less the background mean firing rate averaged over cells in each bin. The different colors represent different contrasts (80, 40, 20 and 10% as defined in the key), and the different plots are at three temporal frequencies (2, 6 and 16 Hz from left to right). Default parameters: with depression and with amplifier: thalamocortical excitation = 20 nA msec, inhibition = 5 nA msec, intracortical excitation = 7 nA msec.

Including NMDA does not affect the mean level of input, and therefore the tuning of the inhibitory cells is not dependent on the presence of NMDA in the thalamocortical connections. When the E⇒I connections are included, the inhibitory cells inherit the temporal frequency tuning of the excitatory cells to varying degrees, both without synaptic depression (figure 4.5, right) and with synaptic depression (figure 4.6, right). Including thalamocortical NMDA onto both excitatory and inhibitory cells causes the greatest shift of the inhibitory cell tuning towards lower frequencies, but the inhibitory cells are predicted to never completely suppress their responses to higher frequencies, unlike excitatory cells.

4.4 Discussion

We have shown that certain parameters of the model effect the orientation and temporal frequency tuning of the mean firing rate of the model inhibitory cells. Synaptic depression suppresses much of the response of the inhibitory cells to null oriented gratings, at low temporal frequencies. Furthermore, it is the presence of the intracortical excitatory-to-inhibitory connections which partially impart the orientation and temporal frequency tuning of the excitatory cells of the model onto their

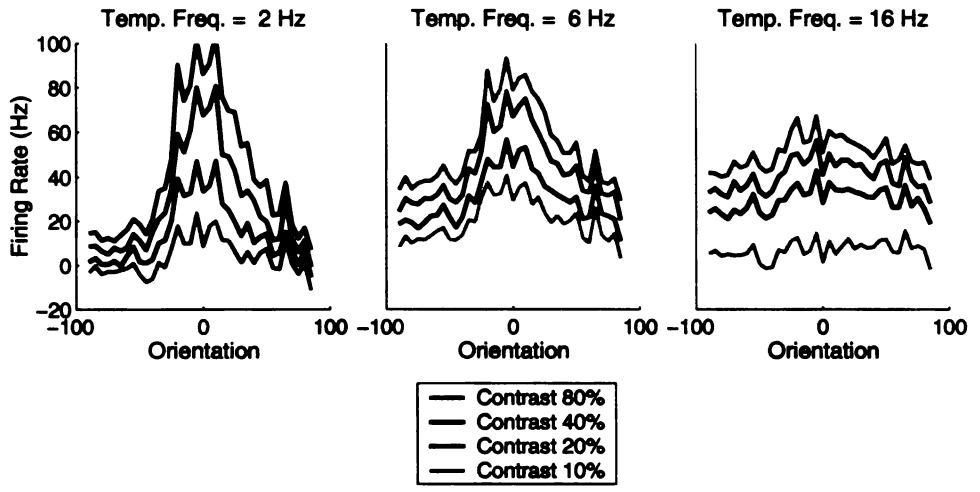

Figure 4.4:

Orientation tuning of mean firing rate of model inhibitory neurons, with synaptic depression and with E⇒I connections. 90% NMDA onto excitatory cells. Plotted is stimulus-induced mean firing rate, that is, mean firing rate averaged over cells in each orientation bin, less the background mean firing rate averaged over cells in each bin. The different colors represent different contrasts (80, 40, 20 and 10% as defined in the key), and the different plots are at three temporal frequencies (2, 6 and 16 Hz from left to right). Default parameters: with depression and with amplifier: thalamocortical excitation = 20 nA msec, inhibition = 5 nA msec, intracortical excitation = 7 nA msec.

neighboring inhibitory cells. Without these E⇒I connections, the mean response of inhibitory cells closely follows the temporal tuning the LGN inputs, and the lack of orientation tuning of the mean LGN input.

It is important to note that these predictions are all for the mean firing rate of the inhibitory cells. Many experimental measures involve the modulating, or F1, spiking response of cortical cells. As long as the inhibitory cells are not rectifying, these two measures should lead to significantly different tuning curves. The modulating response of the inhibitory cells should be more similar to the excitatory cells, both in orientation and in temporal tuning. We have yet to characterize these differences, to further describe our predictions about inhibitory cells responses. Furthermore, we need to characterize the parameter regimes that should lead to inhibitory cell rectification, which would also change the predictions described in this chapter. It is possible that with different levels of thalamocortical and/or background non-specific input onto inhibitory cells, it would be possible to find parameter sets that would preserve reasonable background firing rates, but drive inhibitory cells strongly enough, with the stimulus presentation, to rectify.

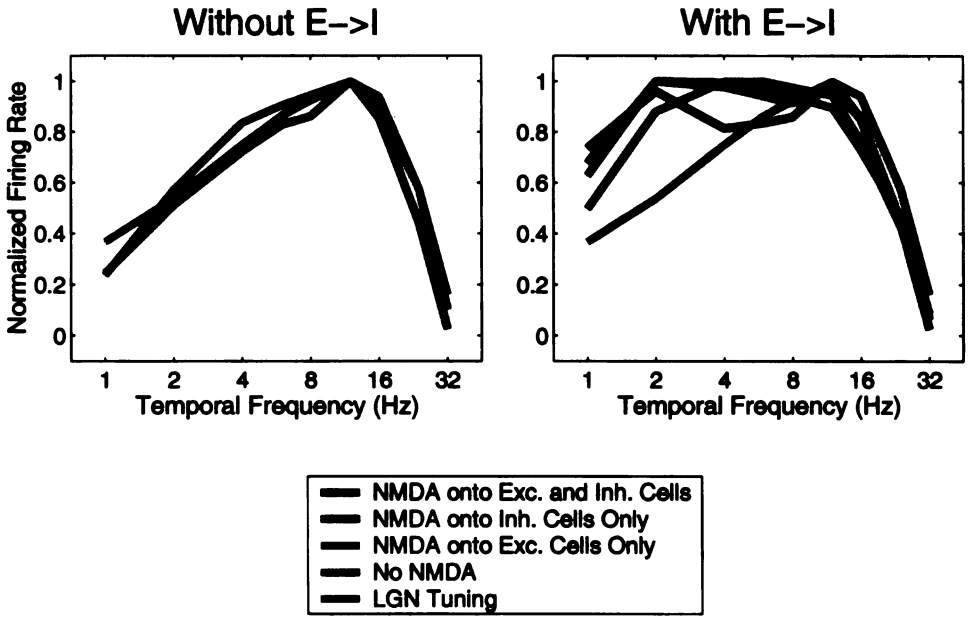

Figure 4.5:

Temporal frequency tuning of the mean firing rate of inhibitory cells at their preferred orientation, with no depression and with (right) and without (left) E⇒I connections, compared to the tuning of the mean LGN firing rate calculated from Sclar [1987]. We show four cases: no NMDA, 90% TC NMDA⇒E and 0% TC NMDA⇒I, 90% TC NMDA⇒E and 0% TC NMDA⇒I, and 90% TC NMDA⇒E and 90% TC NMDA⇒I. Model curves: at each temporal frequency, mean firing rates of all cells preferring the stimulus orientation were averaged together; the resulting population tuning curve was normalized to have a peak of 1. LGN data curve: the temporal frequency tuning curve of mean firing rate, calculated from the F1 amplitudes reported in Sclar [1987] as described in methods of chapter 2, was normalized to have a peak of 1. Default parameters: without depression and with amplifier: thalamocortical excitation = 5 nA msec, inhibition = 3.5 nA msec, intracortical excitation = 4 nA msec;

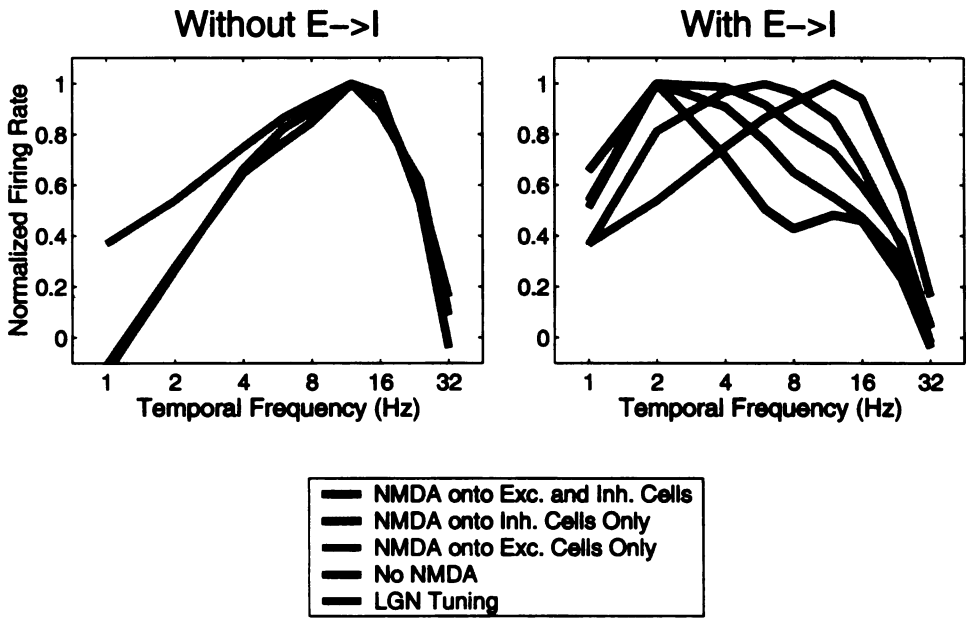

Figure 4.6:

Temporal frequency tuning of the mean firing rate of inhibitory cells at their preferred orientation, with depression and with (right) and without (left) E⇒I connections, compared to the tuning of the mean LGN firing rate calculated from Sclar [1987]. We show four cases: no NMDA, 90% TC NMDA⇒E and 0% TC NMDA⇒I, 90% TC NMDA⇒E and 0% TC NMDA⇒I, and 90% TC NMDA⇒E and 90% TC NMDA⇒I. Model curves: at each temporal frequency, mean firing rates of all cells preferring the stimulus orientation were averaged together; the resulting population tuning curve was normalized to have a peak of 1. LGN data curve: the temporal frequency tuning curve of mean firing rate, calculated from the F1 amplitudes reported in Sclar [1987] as described in methods, was normalized to have a peak of 1. Default parameters: with depression and with amplifier: thalamocortical excitation = 20 nA msec, inhibition = 5 nA msec, intracortical excitation = 7 nA msec;

Bibliography

L. F. Abbott, J. A. Varela, K. Sen, and S. B. Nelson. Synaptic depression and cortical gain control. *Science*, 275:220–224, 1997.

P. Adorján, J. S. Lund, and K. Obermayer. A model for generation and sharpening of orientation selectivity with the same cortical microcircuits in macaque striate cortex. *Soc. Neurosci. Abstr.*, 23:2061, 1997.

D. G. Albrecht. Visual cortex neurons in monkey and cat: Effect of contrast on the spatial and temporal phase transfer functions. *Vis. Neurosci.*, 12:1191–1210, 1995.

D. G. Albrecht and W. S. Geisler. Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Vis. Neurosci.*, 7:531–546, 1991.

J. D. Allison, J. F. Kabara, R. K. Snider, V. A. Casagrande, and A. B. Bonds. GABA-B-receptor-mediated inhibition reduces the orientation selectivity of the sustained response of striate cortical neurons in cats. *Vis. Neurosci.*, 13:559–566, 1996.

J. D. Allison and K. A. C. Martin. Contrast adaptation produced by null direction and cross orientation stimulation of neurons in cat visual cortex. *Soc. Neurosci. Abstr.*, 23:454, 1997.

J. M. Alonso, W. M Usrey, and R. C. Reid. Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature*, 383:815–819, 1996.

J. Anderson, M. Carandini, and D. Ferster. Orientation tuning of input conductance in cat primary visual cortex. *J. Neurophysiol.*, (submitted), 2000.

M. C. Angulo, J. Rossier, and E. Audinat. Postsynaptic glutamate receptors and integrative properties of fast-spiking interneurons in the rat neocortex. *J. Neurophysiol.*, 82:1295–1302, 1999.

M. Armstrong-James, E. Welker, and C. A. Callahan. The contribution of NMDA and non-NMDA receptors to fast and slow transmission of sensory information in the rat SI barrel cortex. *J. Neurosci.*, 13:2149–2160, 1993.

R. Azouz, C. M. Gray, L. G. Nowak, and D. A. McCormick. Physiological properties of inhibitory interneurons in cat striate cortex. *Cerebral Cortex*, 7:534–545, 1997.

C. L. Baker. Spatial- and temporal-frequency selectivity as a basis for velocity preference in cat striate cortex neurons. *Vis. Neurosci.*, 4:101–113, 1990.

M. F. Bear, A. Kleinschmidt, Q. Gu, and W. Singer. Disruption of experience-dependent synaptic modifications in striate cortex by infusion of an NMDA receptor antagonist. *J. Neurosci.*, 10: 909–925, 1990.

R. Ben-Yishai, R. L. Bar-Or, and H. Sompolinsky. Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA*, 92:3844–3848, 1995.

N. J. Berman, R. J. Douglas, K. A. C. Martin, and D. Whitteridge. Mechanisms of inhibition in cat visual cortex. *J. Physiol.*, 440:697–722, 1991.

O. Bernander, R. Douglas, K. Martin, and C. Koch. Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proc. Nat. Acad. Sci., USA*, 88:11569–11573, 1991.

P. O. Bishop, W. Kozak, and G. J. Vakkur. Some quantitative aspects of the cat's eye: Axis and plane of reference, visual field co-ordinates and optics. *J. Physiol.*, 163:466–502, 1962.

S. Bisti, G. Carmignoto, L. Galli, and L. Maffei. Spatial-frequency characteristics of neurones of area 18 in the cat: Dependence on the velocity of the visual stimulus. *J. Physiol.*, 359:259–268, 1985.

A. B. Bonds. Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Vis. Neurosci.*, 2:41–55, 1989.

L. J. Borg-Graham, C. Monier, and Y. Frégnac. Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, 393:369–373, 1998.

J. C. Brumberg, D. Goldreich, H. T. Kyriazi, and D. J. Simons. Adjacent whisker input to a layer IV barrel does not arise from adjacent barrels. *Soc. Neurosci. Abstr.*, 23:2344, 1997.

J. C. Brumberg, D. J. Pinto, and D. J. Simons. Spatial gradients and inhibitory summation in the rat whisker barrel system. *J. Neurophysiol.*, 76:130–140, 1996.

J. Bullier and G. H. Henry. Laminar distribution of first-order neurons and afferent terminals in cat striate cortex. *J. Neurophysiol.*, 42:1271–1281, 1979.

P. C. Bush and N. J. Priebe. GABAergic inhibitory control of the transient and sustained components of orientation selectivity in a model microcolumn in layer 4 of cat visual cortex. *Neural Comput.*, 10:(in press), 1998.

D. Cai, G. C. DeAngelis, and R. D. Freeman. Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. *J. Neurophysiol.*, 78:1045–1061, 1997.

M. Carandini and D. Ferster. A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science*, 276:949–952, 1997.

M. Carandini and D. Ferster. Membrane potential and firing rate in cat primary visual cortex. *J. Neurosci.*, 20:470–484, 2000.

M. Carandini and D. J. Heeger. Summation and division by neurons in visual cortex. *Science*, 264: 1333–1336, 1994.

M. Carandini, D. J. Heeger, and J. A. Movshon. Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.*, 17:8621–8644, 1997.

M. Carandini, D. J. Heeger, and J. A. Movshon. Linearity and gain control in V1 simple cells. In E. G. Jones and A. Peters, editors, *Cerebral Cortex, Vol. 13*. Plenum, New York, 1998.

M. Carandini and D. Ringach. Predictions of a recurrent model of orientation selectivity. *Vision Res.*, 37:3061–3071, 1997.

G. Carmignoto and S. Vicini. Activity-dependent decrease in NMDA receptor responses during development of the visual cortex. *Science*, 258:1007–1011, 1992.

F. S. Chance, S. B Nelson, and L. F. Abbott. Effects of synaptic depression of temporal nonlinearities in responses of model simple cells. *Soc. Neurosci. Abstr.*, 23:1266, 1997.

F. S. Chance, S. B Nelson, and L. F. Abbott. Synaptic depression and the temporal response characteristics of V1 cells. *J. Neurosci.*, 18:4785–4799, 1998.

H. Cheng, Y. M Chino, E. L. Smith, J. Hamamoto, et al. Transfer characteristics of lateral geniculate nucleus X neurons in the cat: Effects of spatial frequency and contrast. *J. Neurophysiol.*, 74:2548–2557, 1995.

S. Chung and D. Ferster. The size of thalamic input to simple cells of the cat visual cortex. *Soc. Neurosci. Abstr.*, 23:2059, 1997.

S. Chung and D. Ferster. Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. *Neuron*, 20:1177–89, 1998.

M. C. Crair and R. C. Malenka. A critical period for long-term potentiation at thalamocortical synapses. *Nature*, 375, 1995.

O. Creutzfeldt, F. C. Hellweg, and C. Schreiner. Thalamocortical transformation of responses to complex auditory stimuli. *Experimental Brain Research*, 39:87–104, 1980.

Y. Dan, J. J. Atick, and R. C. Reid. Efficient coding of natural scenes in the lateral geniculate nucleus: Experimental test of a computational theory. *J. Neurosci.*, 16:3351–3362, 1996.

A. F. Dean, D. J. Tolhurst, and N. S. Walker. Non-linear temporal summation by simple cells in cat striate cortex demonstrated by failure of superposition. *Exp. Brain Res.*, 45:456–458, 1982.

G. C. DeAngelis, G. M. Ghose, I. Ohzawa, and R. D. Freeman. Spatiotemporal receptive field structure and phase relationships between adjacent simple cells in the cat's striate cortex. *Soc. Neurosci. Abstr.*, 18:10, 1992.

G. C. DeAngelis, I. Ohzawa, and R. D. Freeman. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *J. Neurophysiol.*, 69:1091–1117, 1993a.

G. C. DeAngelis, I. Ohzawa, and R. D. Freeman. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II. Linearity of temporal and spatial summation. *J. Neurophysiol.*, 69:1118–1135, 1993b.

A. M. Derrington and A. F. Fuchs. Spatial and temporal properties of X and Y cells in the cat lateral geniculate nucleus. *J. Physiol.*, 293:347–364, 1979.

R. J. Douglas, C. Koch, M. Mahowald, K. A. Martin, and H. H. Suarez. Recurrent excitation in neocortical circuits. *Science*, 269:981–985, 1995.

R. J. Douglas, K. A. C. Martin, and D. Whitteridge. Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature*, 332:642–644, 1988.

R. J. Douglas, K. A. C. Martin, and D. Whitteridge. A canonical microcircuit for neocortex. *Neural Comput.*, 1:480–488, 1989.

R. J. Douglas, K. A. C. Martin, and D. Whitteridge. An intracellular analysis of the visual responses of neurones in cat visual cortex. *J. Physiol.*, 440:659–696, 1991.

J. J. Eggermont. Differential effects of age on click-rate and amplitude modulation-frequency coding in primary auditory cortex of the cat. *J. Physiol.*, 293:347–364, 1993.

D. Feldmeyer, V. Egger, J. Lübke, and B. Sakmann. Reliable synaptic connections between pairs of excitatory layer 4 neurones within a sibgle "barrel" of developing rat somatosensory cortex. *J. Physiol.*, 521:169–190, 1999.

D. G. Ferrington and M. Rowe. Differential contributions to coding of cutaneous vibratory inforemation by cortical somatosensory areas I and II. *J. Neurophysiol.*, 43:310–331, 1985.

D. Ferster. Orientation selectivity of synaptic potentials in neurons of cat primary visual cortex. *J. Neurosci.*, 6:1284–1301, 1986.

D. Ferster. Origin of orientation-selective EPSPs in simple cells of cat visual cortex. *J. Neurosci.*, 7:1780–1791, 1987.

D. Ferster. Spatially opponent excitation and inhibition in simple cells of the cat visual cortex. *J. Neurosci.*, 8:1172–1180, 1988.

D. Ferster. X- and Y-mediated current sources in areas 17 and 18 of cat visual cortex. *Vis. Neurosci.*, 4:135–145, 1990a.

D. Ferster. X- and Y-mediated synaptic potentials in neurons of areas 17 and 18 of cat visual cortex. *Vis. Neurosci.*, 4:115–133, 1990b.

D. Ferster, S. Chung, and H. Wheat. Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature*, 380:249–252, 1996.

D. Ferster and S. Lindström. An intracellular analysis of geniculo-cortical connectivity in area 17 of the cat. *J. Physiol.*, 342:181–215, 1983.

I. A. Fleidervish, A. M. Binshtok, and M. J. Guntick. Functionall distinct NMDA receptors mediate horizontal connectivity within layer 4. *Neuron*, 21:1055–1065, 1998.

K. H. Foster, J. P. Gaska, and M. Nagler. Spatial and temporal frequency selectivity of neurones in visualcortical areas V1 and V2 of the macaque monkey. *J. Physiol.*, 365:331–363, 1985.

K. Fox, H. Sato, and N. Daw. The location and function of NMDA receptors in cat and kitten visual cortex. *J. Neurosci.*, 9:2443–2454, 1989.

K. Fox, H. Sato, and N. Daw. The effect of varying stimulus intensity on NMDA-receptor activity in cat visual cortex. *J. Neurophysiol.*, 64:1413–1428, 1990.

R. D. Freeman, G. C. DeAngelis, G. M. Ghose, and I. Ohzawa. Clustering of response properties of neurons in the visual cortex. *Soc. Neurosci. Abstr.*, 23:567, 1997.

E. Gary-Bobo, J. Przybyslawski, and P. Saillour. Experience-dependent maturation of the spatial and temporal characteristics of the cell receptive fields in the kitten visual cortex. *Neuroscience Letters*, 189:147–150, 1995.

G. M. Ghose, G. C. DeAngelis, I. Ohzawa, and R. D. Freeman. Similarity between the spatiotemporal response profiles of nearby cells in the cat's and kitten's visual system. *Soc. Neurosci. Abstr.*, 19:424, 1993.

Z. Gil and Y. Amitai. Adult thalamocortical transmission involves both NMDA and non-NMDA receptors. *J. Neurophysiol.*, 76:2547–2554, 1996.

C. D. Gilbert. Laminar differences in receptive field properties of cells in cat primary visual cortex. *J. Physiol.*, 268:391–421, 1977.

C. D. Gilbert and T. N. Wiesel. Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature*, 280:120–125, 1979.

S. V. Girman, Y. Sauve, and R. D. Lund. Receptive field properties of single neurons in rat primary visual cortex. *J. Neurophysiol.*, 82:301–311, 1999.

V. D. Glezer, T. A. Tsherbach, V. E. Gauselman, and V. M. Bondarko. Spatio-temporal organization of receptive fields of the cat striate cortex: The receptive fields as the grating filters. *Biol. Cybern.*, 43:35–49, 1982.

K. Hagihara, T. Tsumoto, H. Sato, and Y. Hata. Actions of excitatory amino acid antagonists on geniculo-cortical transmission in the cat's visual cortex. *Exp. Brain Res.*, 69:407–416, 1988.

J. Hamamoto, H. Cheng, K. Yoshida, E. L. III Smith, and Y. M. Chino. Transfer characteristics of lateral geniculate nucleus X-neurons in the cat: effects of temporal frequency. *Exp. Brain Res.*, 98:191–199, 1994.

P. Hammond and C. J. Pomfrett. Influence of spatial frequency on tuning and bias for orientation and direction in the cat's striate cortex. *Vision Res.*, 30:359–369, 1990.

M. J. Hawken, C. Blakemore, and J. W. Morley. Development of contrast sensitivity and temporal-frequency selectivity in primate later geniculate nucleus. *Exp. Brain Res.*, 114:86–98, 1997.

M. J. Hawken, R. M. Shapley, and D. H. Grosof. Temporal-frequency selectivity in monkey visual cortex. *Vis. Neurosci.*, 13:477–492, 1996.

D. J. Heeger. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.*, 9:181–198, 1992.

P. Heggelund and K. Albus. Orientation selectivity of single cells in striate cortex of cat: The shape of orientation tuning curves. *Vision Res.*, 18:1067–1071, 1978.

F. C. Hellweg, W. Schultz, and O. D. Creutzfeldt. Extracellular and intracellular recordings from cat's cortical whisker projection area: thalamocortical repsonse transformation. *J. Neurophysiol.*, 40:463–479, 1977.

J. A. Hirsch, J. M. Alonso, and R. C. Reid. Two-dimensional spatial maps of excitation and inhibition in the simple receptive field. *Soc. Neurosci. Abstr.*, 21:21, 1995a.

J. A. Hirsch, J. M. Alonso, and R. C. Reid. Visually evoked calcium action potentials in cat striate cortex. *Nature*, 378:612–616, 1995b.

J. A. Hirsch, J.-M. Alonso, R. C. Reid, and L.M. Martinez. Synaptic integration in striate cortical simple cells. *J. Neurosci.*, 18:9517–9528, 1998.

R. A. Holub and M. Morton-Gibson. Response of visual cortical neurons of the cat to moving sinusoidal gratings: Response-contrast functions and spatiotemporal interactions. *J. Neurophysiol.*, 46:1244–1259, 1981.

D. H. Hubel and T. N. Wiesel. Integrative action in the cat's lateral geniculate body. *J. Physiol.*, 155:385–398, 1961.

D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.*, 160:106–154, 1962.

Mark Hübener, Doron Shoham, Amiram Grinvald, and Tobias Bohnoeffer. Spatial relationships among three columnar systems in cat area 17. *J. Neurosci.*, 17:9285–9307, 1997.

A. L. Humphrey, M. Sur, D. J. Uhlrich, and S. M. Sherman. Projection patterns of individual X- and Y-cell axons from the lateral geniculate nucleus to cortical area 17 in the cat. *J. Comp. Neurol.*, 233:159–189, 1985a.

A. L. Humphrey, M. Sur, D. J. Uhlrich, and S. M. Sherman. Termination patterns of individual X- and Y-cell axons in the visual cortex of the cat: Projections to area 18, to the 17/18 border region, and to both areas 17 and 18. *J. Comp. Neurol.*, 233:190–212, 1985b.

H. Ikeda and M. J. Wright. Spatial and temporal properties of sustained and transient neurones in area 17 of the cat's visual cortex. *Exp. Brain Res.*, 22:363–383, 1975.

C. E. Jahr and C. F. Stevens. Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *J. Neurosci.*, 10:3178–3182, 1990.

J. P. Jones and L. A. Palmer. The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J. Neurophysiol.*, 58:1187–1211, 1987.

J. P. Jones, A. Stepnoski, and L. A. Palmer. The two-dimensional spectral structure of simple receptive fields in cat striate cortex. *J. Neurophysiol.*, 59:1212–1232, 1987.

E. Kaplan, K. Purpura, and R. M. Shapley. Contrast affects the transmission of visual information trough the mammalian lateral geniculate nucleus. *J. Physiol.*, 391:267–288, 1987.

T. Kasamatsu, K. Imamura, N. Mataga, E. Hartveit, U. Heggelund, and P. Heggelund. Roles of N-methyl-D-aspartate receptors in ocular dominance plasticity in developing visual cortex: re-evaluation. *Neuroscience*, 82:687–700, 1998.

A. S. Kayser, N. J. Priebe, and K. D. Miller. Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. *J. Neurosci.*, (submitted), 1999.

C. Koch, R. Douglas, and U. Wehmeier. Visibility of synaptically induced conductance changes: theory and simulations of anatomically characterized cortical pyramidal cells. *J. Neurosci.*, 10: 1728–1744, 1990.

Y. Komatsu. $GABA_B$ receptors, monoamine receptors, and postsynaptic inositol trisphosphate-induced Ca^{2+} release are involved in the induction of long-term potentiation at visual cortical inhibitory synapses. *J. Neurosci.*, 16:6342–6352, 1996.

A. Krukowski, N. J. Priebe, and K. D. Miller. A model of simple-cell orientation tuning: Feed-forward tuning and correlation-based intracortical connectivity. *Soc. Neurosci. Abstr.*, 22:642, 1996.

L. J. Larson-Prior, P. S. Ulinksi, and N. T. Slater. Excitatory amino acid receptor-mediated transmission in geniculocortical and intracortical pathways within visual cortex. *J. Neurophysiol.*, 66:293–306, 1991.

B. B. Lee, A. Elepfandt, and V. Virsu. Phase of responses to moving sinusoidal gratings in cells of cat retina and lateral geniculate nucleus. *J. Neurophysiol.*, 45:807–817, 1981a.

B. B. Lee, A. Elepfandt, and V. Virsu. Phase of responses to sinusoidal gratings of simple cells in cat striate cortex. *J. Neurophysiol.*, 45:818–828, 1981b.

R. A. Lester, J. D. Clements, G. L. Westbrook, and C. E. Jahr. Channel kinetics determine the time course of NMDA receptor-mediated synaptic currents. *Nature*, 346:565–567, 1990.

M. W. Levine and J. B. Troy. The variability of the maintained discharge of cat dorsal lateral geniculate cells. *J. Physiol.*, 375:339–359, 1986.

D. S. F. Ling and L. S. Benardo. Recruitment of $GABA_A$ inhibition in rat neocortex is limited and not NMDA dependent. *J. Neurophysiol.*, 74:2329–2335, 1995.

R. A. Linsenmeier, L. J. Frishman, H. G. Jakiela, and C. Enroth-Cugell. Receptive field properties of X and Y cells in the cat retina derived from contrast sensitivity measurements. *Vision Res.*, 22:1173–1183, 1982.

R. Maex and G. A. Orban. A model circuit for cortical temporal low-pass filtering. *Neural Comput.*, 4:932–945, 1992.

N. J. Mangini and A. L. Pearlman. Laminar distribution of receptive field properties in the primary visual cortex of the mouse. *Journal of Comparative Neurology*, 80:203–222, 1980.

H. Markram, J. Lübke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275:213–215, 1997.

H. Markram and M. Tsodyks. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382:807–810, 1996.

K. A. C. Martin. From single cells to simple circuits in the cerebral cortex. *J. Exp. Physiol.*, 73: 637–702, 1988.

S. L. McCabe. Cortical synaptic depression and auditory perception. *preprint*, pages 111–116, 1999.

D. A. McCormick, B. W. Connors, J. W. Lighthall, and D. A. Prince. Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J. Neurophysiol.*, 54: 782–805, 1985.

J. McLean and L. A. Palmer. Organization of simple cell responses in the three-dimensional (3-D) frequency domain. *Visual Neuroscience*, 11:295–306, 1994.

B. W. Mel. Synaptic integration in an excitable dendritic tree. *J. Neurophysiol.*, 70:1086–1101, 1993.

M. Merleau-Ponty. *Phenomenology of perception.* Humanities Press, New York, 1962.

K. D. Miller. A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. *J. Neurosci.*, 14:409–441, 1994.

K. D. Miller, B. Chapman, and M. P. Stryker. Responses of cells in cat visual cortex depend on NMDA receptors. *Proc. Natl. Acad. Sci. USA*, 86:5183–5187, 1989.

C. Monier, L. Borg-Graham, and Y. Fregnac. The dominant increase in synaptic conductance evoked by visual input in cat area 17 is due to shunting inhibition. *Soc. Neurosci. Abstr.*, 23: 2363, 1997.

M. C. Morrone, M. Di Stefano, and D. C. Burr. Spatial and temporal properties of neurons of the lateral suprasylvian cortex of the cat. *J. Neurophysiol.*, 56:969–986, 1986.

V. B. Mountcastle, W. H. Talbot, H. Sakata, and J. Hyvarinen. Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. neuronal periodicity and frequency discrimination. *J. Neurophysiol.*, 32:452–484, 1969.

J. A. Movshon, I. D. Thompson, and D. J. Tolhurst. Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat visual cortex. *J. Physiol.*, 283:101–120, 1978a.

J. A. Movshon, I. D. Thompson, and D. J. Tolhurst. Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J. Physiol.*, 283:53–77, 1978b.

S. Nelson, L. Toth, B. Sheth, and M. Sur. Orientation selectivity of cortical neurons during intracellular blockade of inhibition. *Science*, 265:774–777, 1994.

L. G. Nowak, M. V. Sanchez-Vives, and D. A. McCormick. Influence of low and high frequency inputs on spike timing in visual cortical neurons. *Cerebral Cortex*, 7:487–501, 1997.

L. P. O'Keefe, J. B. Levitt, D. C. Kiper, R. M. Shapley, and J. A. Movshon. Functional organization of owl monkey lateral geniculate nucleus and visual cortex. *J. Neurophysiol.*, 80:594–609, 1998.

G. A. Orban. Quantitative electrophysiology of visual cortical neurones. pages 173–222, 1991.

G. A. Orban, K.-P. Hoffmann, and J. Duysnes. Velocity selectivity in the cat visual system. I. responses of LGN cells to moving bar stimuli: A comparison with cortical areas 17 and 18. *J. Neurophysiol.*, 54:1026–1049, 1985.

G. A. Orban, H. Kennedy, and J. Bullier. Velocity sensitivity and direction selectivity of neurons in areas V1 and V2 of the monkey: Influence of eccentricity. *J. Neurophysiol.*, 56:462–480, 1986.

G. A. Orban, H. Kennedy, and H. Maes. Response to movement of neurons in areas 17 and 18 of the cat: Velocity sensitivity. *J. Neurophysiol.*, 45:1043–1058, 1981.

L. A. Palmer and T. L. Davis. Receptive-field structure in cat striate cortex. *J. Neurophysiol.*, 46: 260–276, 1981.

L. Peichl and H. Wassle. Size, scatter and coverage of ganglion cell receptive field centres in the cat retina. *J. Physiol.*, 291:117–141, 1979.

M. Perouansky and Y. Yaari. Kinetic properties of NMDA receptor-mediated synaptic currents in rat hippocampal pyramidal cells *versus* interneurones. *J. Physiol.*, 465:223–244, 1993.

A. Peters and E. Yilmaz. Neuronal organization in area-17 of cat visual cortex. *Cerebral Cortex*, 3:49–68, 1993.

D. J. Pinto, J. C. Brumberg, D. J. Simons, and G. B. Ermentrout. A quantitative population model of whisker barrels: Re-examining the Wilson-Cowan equations. *J. Comput. Neurosci.*, 3: 247–264, 1996.

N. J. Priebe, A. S. Kayser, A. E. Krukowski, and K. D. Miller. A model of simple cell orientation tuning: The role of synaptic depression. *Soc. Neurosci. Abstr.*, 23:2061, 1997.

R. C. Reid and J. M. Alonso. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378:281–284, 1995.

R. C. Reid, J. D. Victor, and R. M. Shapley. Broadband temporal stimuli decrease the integration time of neurons in cat striate cortex. *Vis. Neurosci.*, 9:39–45, 1992.

E. B. Roberts, A. Meredith, and A. S. Ramoa. Suppression of NMDA receptor function using antisense DNA blocks ocular dominance plasticity while preserving visual responses. *J. Neurophysiol.*, 80:1021–1032, 1998.

T. E. Salt, C. L. Meier, N. Seno, T. Krucker, and P. L. Herrling. Thalamocortical and corticocortical excitatory postsynaptic potentials mediated by excitatory amino acid receptors in the cat motor cortex *in vivo*. *Neuroscience*, 64:433–442, 1995.

H. Sato, N. Katsuyama, H. Tamura, Y. Hata, and T. Tsumoto. Mechanisms underlying orientation slectivity of neurons in the primary visual cortex of the macaque. *J. Physiol.*, 494:757–771, 1996.

A. B. Saul and A. L. Humphrey. Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *J. Neurophysiol.*, 64:206–224, 1990.

A. B. Saul and A. L. Humphrey. Evidence of input from lagged cells in the lateral geniculate nucleus to simple cells in cortical area 17 of the cat. *J. Neurophysiol.*, 68:1190–1208, 1992.

C. Schreiner. Coding of amplitude modulation in auditory thalamus and cerebral cortex. *preprint*, 1999.

G. Sclar. Expression of "retinal" contrast gain control by neurons of the cat's lateral geniculate nucleus. *Exp. Brain Res.*, 66:589–596, 1987.

G. Sclar and R. D. Freeman. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp. Brain Res.*, 46:457–461, 1982.

Michael N. Shadlen and William T. Newsome. Noise, neural codes and cortical organization. *Curr. Op. Neurobiology*, 4:569–579, 1994.

R. M. Shapley and J. D. Victor. The effect of contrast on the transfer properties of cat retinal ganglion cells. *J. Physiol. (London)*, 285:275–298, 1978.

H. Sherk. Location and connections of visual cortical areas in the cat's suprasylvian sulcus. *J. Comp. Neurol.*, 274:1–31, 1986.

S. M. Sherman. Functional organization of the W-, X-, and Y-cell pathways in the cat: A review and hypothesis. *Progr. Psychobiol. Physiol. Psychol.*, 11:233–314, 1985.

D. Shoham, M. Hübener, S. Shulze, A. Grinvald, and T. Bonhoeffer. Spatio-temporal frequency domains and their relation to cytochrome oxidase staining in cat visual cortex. *Nature*, 385: 529–533, 1997.

A. M. Sillito. The contribution of inhibitory mechanisms to the receptive field properties of neurones in the striate cortex of the cat. *J. Physiol.*, 250:305–329, 1975.

D. J. Simons. Response properties of vibrissa units in rat S1 somatosensory neocortex. *J. Neurophysiol.*, 41:798–820, 1978.

D. J. Simons. Temporal and spatial integration in the rat S1 vibrissa cortex. *J. Neurophysiol.*, 54: 615–635, 1985.

D. J. Simons and G. E. Carvell. Thalamocortical response transformation in the rat vibrissa/barrel system. *J. Neurophysiol.*, 61:311–330, 1989.

B. C. Skottun, A. Bradley, G. Sclar, I. Ohzawa, and R. D. Freeman. The effects of contrast on visual orientation and spatial frequency discrimination: A comparison of single cells and behavior. *J. Neurophysiol.*, 57:773–786, 1987.

B. C. Skottun, R. L. De Valois, D. H. Grosof, J. A. Movshon, D. G. Albrecht, and A. B. Bonds. Classifying simple and complex cells on the basis of response modulation. *Vision Res.*, 38: 1079–86, 1991.

D. Somers, S. B. Nelson, and M. Sur. An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.*, 15:5448–5465, 1995.

J. F. Storm. Potassium currents in hippocampal pyramidal cells. *Prog. Brain Res.*, 83:161–187, 1990.

K. J. Stratford, K. Tarczy-Hornoch, K. A. Martin, N. J. Bannister, and J. J. Jack. Excitatory synaptic inputs to spiny stellate cells in cat visual cortex. *Nature*, 382:258–261, 1996.

H. Suarez, C. Koch, and R. Douglas. Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit. *J. Neurosci.*, 15:6700–6719, 1995.

H. A. Swadlow. Efferent neurons and suspected interneurons in binocular visual cortex of the awake rabbit: Receptive fields and binocular properties. *J. Neurophysiol.*, 59:1162–1187, 1988.

H. A. Swadlow. Efferent neurons and suspected interneurons in S-1 vibrissa cortex of the awake rabbit: Receptive fields and axonal properties. *J. Neurophysiol.*, 62:288–308, 1989.

H. A. Swadlow. Efferent neurons and suspected interneurons in S-1 forelimb representation of the awake rabbit: Receptive fields and axonal properties. *J. Neurophysiol.*, 63:1477–1498, 1990.

H. A. Swadlow. Efferent neurons and suspected interneurons in second somatosensory cortex of the awake rabbit: Receptive fields and axonal properties. *J. Neurophysiol.*, 66:1392–1409, 1991.

H. A. Swadlow. Efferent neurons and suspected interneurons in motor cortex of the awake rabbit: Axonal properties, sensory receptive fields, and subthreshold synaptic inputs. *J. Neurophysiol.*, 71:437–453, 1994.

K. Tanaka. Cross-correlation analysis of geniculostriate neuronal relationships in cats. *J. Neurophysiol.*, 49:1303–1318, 1983.

A. C. Tang, A. M. Bartels, and T. J. Sejnowski. Effects of cholinergic modulation on responses of neocortical neurons to fluctuating input. *Cerebral Cortex*, 7:502–509, 1997.

Y.-P. Tang, E. Shimizu, G. R. Dube, C. Rampon, G. A. Kerchner, M. Zhuo, G. Liu, and J. Z. Tsien. Genetic enhancement of learning and memory in mice. *Nature*, 401:63–69, 1999.

K. Tarczy-Hornoch. *Physiology of Synaptic Inputs to layer IV of Cat Visual Cortex.* PhD thesis, Oxford University, England, 1996.

A. M. Thomson and J. Deuchars. Temporal and spatial properties of local circuits in neocortex. *Trends Neurosci.*, 17:119–126, 1994.

A. M. Thomson, J. Deuchars, and D. C. West. Single axon excitatory postsynaptic potentials in neocortical interneurons exhibit pronounced paired pulse facilitation. *Neurosci.*, 54:347–360, 1993.

A. M. Thomson and D. C. West. Fluctuations in pyramid-pyramid excitatory postsynaptic potentials modified by presynaptic firing pattern and postsynaptic membrane potential using paired intracellular recordings in rat neocortex. *Neurosci.*, 54:329–346, 1993.

D. J. Tolhurst and A. F. Dean. The effects of contrast on the linearity of spatial summation of simple cells in the cat's striate cortex. *Exp. Brain Res.*, 79:582–588, 1990.

D. J. Tolhurst and J. A. Movshon. Spatial and temporal contrast sensitivity of striate cortical neurons. *Nature*, 257:674–675, 1975.

D. J. Tolhurst and I. D. Thompson. On the variety of spatial frequency selectivities shown by neurons in area 17 of the cat. *Proc. Royal. Soc. Lond. B*, 213:183–199, 1981.

D. J. Tolhurst, N. S. Walker, I. D. Thompson, and A. F. Dean. Non-linearities of temporal summation in neurones in area 17 of the cat. *Exp. Brain Res.*, 38:431–435, 1980.

K. Toyama, M. Kimura, and K. Tanaka. Organization of cat visual cortex as investigated by cross-correlation technique. *J. Neurophysiol.*, 46:202–214, 1981.

J. B. Troy. Spatio-temporal interaction in neurones of the cat's dorsal lateral geniculate nucleus. *J. Physiol.*, 344:419–432, 1983.

T. W. Troyer, A. E. Krukowski, N. J. Priebe, and K. D. Miller. Contrast-invariant orientation tuning in cat visual cortex: Feedforward tuning and correlation-based intracortical connectivity. *J. Neurosci.*, 18:5908–5927, 1998.

T. W. Troyer and K. D. Miller. Integrate-and-fire neurons matched to physiological f-I curves yield high input sensitivity and wide dynamic range. In J. M. Bower, editor, *Computational Neuroscience: Trends in Research 1997*, pages 197–201. Plenum Press, New York, 1997a.

T. W. Troyer and K. D. Miller. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comput.*, 9:971–983, 1997b.

M. V. Tsodyks and H. Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. USA*, 94, 1997.

T. Tsumoto, W. Eckart, and O. D. Creutzfeld. Modification of orientation sensitivity of cat visual cortex neurons by removal of GABA-mediated inhibition. *Exp. Brain Res.*, 34:351–363, 1979.

R. J. Tusa, L. A. Palmer, and A. C. Rosenquist. The retinotopic organization of area 17 (striate cortex) in the cat. *J. Comp. Neurol.*, 177:213–235, 1978.

T. R. Vidyasagar and J. A. Sigüenza. Relationship between orientation tuning and spatial frequency in neurones of cat area 17. *Exp. Brain Res.*, 57:628–631, 1985.

M. A. Webster and R. L. De Valois. Relationship between spatial-frequency and orientation tuning of striate-cortex cells. *J. Opt. Soc. Am. A.*, 2:1124–1132, 1985.

F. Wörgötter and C. Koch. A detailed model of the primary visual pathway in the cat: Comparison of afferent excitatory and intracortical inhibitory connection schemes for orientation selectivity. *J. Neurosci.*, 11:1959–1979, 1991.

R. Yuste and D. W. Tank. Dendritic integration in mammalian neurons, a century after Cajal. *Neuron*, 16:701–716, 1996.

S. W. Zucker. The computational connection in vision: Early orientation selection. *Behav. Res. Meth., Instr. and Computers*, 18:608–617, 1986.

T. J. Zumbroich and C Blakemore. Spatial and temporal selectivity in the suprasylvian visual cortex of the cat. *J. Neurosci.*, 7:482–500, 1987.

T. J. Zumbroich, D. J. Price, and C Blakemore. Development of spatial and temporal selectivity in the suprasylvian visual cortex of the cat. *J. Neurosci.*, 8:2713–2728, 1988.

arch 13, 2000

o: Ann Desai
ociety for Neuroscience
ax: (202) 462 1547

rom: Anton Krukowski
niversity of California
epartment of Physiology
ax: (415) 476 4929
hone: (415) 476 1007

ear Ms. Desai,

am a graduate student at the University of California,
an Francisco in the Biophysics Graduate Group.
am requesting permission to use material from the following
rticle that was published in the Journal of
euroscience in August of 1998, as part of my dissertation:

odd W. Troyer, Anton E. Krukowski, Nicholas J. Priebe, and Kenneth D. Miller,
"Contrast-invariant orientation tuning in cat visual cortex:
Feedforward tuning and correlation-based intracortical connectivity."
J. Neurosci. 18:5908-5927, 1998.

I would like to include material from all of the pages (5908-5927)
as part of my dissertation; in fact, the entire article is what
I had in mind.

Thank you very much.

Sincerely,

Anton Krukowski

Nancy Berry 3/13/2000

Permission is granted by the Society for Neuroscience provided the original *The Journal of Neuroscience* reference is always cited.