The Julius Silver, Roslyn S. Silver, and Enid Silver Winslow
DIALOGUES IN ARTS AND SCIENCE

LARGE-SCALE MODELING IN NEURAL SCIENCE

David W. McLaughlin
Silver Professor of Mathematics and Neural Science

NEW YORK UNIVERSITY

Large-Scale Computational Modeling in Neural Science

Abstract

In this essay, I describe the general roles that large-scale computational modeling can play in modern neural science, and the ways it can contribute to our understanding of the human brain and its function. The essay begins by emphasizing that neural science is fundamentally an experimental science, and that its historical advances have always followed significant technological advances in experimental methods. Advances in computers and information technology of the last quarter of the 20th century constitute technological advances of similar importance. So next, I summarize some of the roles in neural science of large-scale computational modeling. As an example of these role, after presenting an overview of our large-scale computational model of a single layer of the mammalian primary visual cortex, I summarize the contributions that this model has made to our understanding the mechanistic operation and function of the primary visual cortex. Finally, I summarize the cortical operating point or state of the computational model, from which it achieves results consistent with biological experiments; thus, it is likely that I am describing the cortical operating point of the mammalian primary visual cortex itself.

1. Introduction

The human brain is a marvelous, yet very complex, dynamical system. To appreciate this, just reflect on the difficult nature of its tasks and functions — sensual perception, cognition, memory, understanding, consciousness. But one becomes even more aware of the amazing nature of the human brain upon reflection about *the ease* with which it concurrently performs and synthesizes its multiple tasks.

The "cortex" is that part of the brain that receives, remembers, processes, and directs the use of sensual information. Today, we know a great deal, and yet we know very little, about the human brain and the mechanisms underlying its functions. This lack of understanding is in spite of the vast progress that has been made over the last century and a half, at a rate that has increased dramatically over the past thirty years and continues to increase today. Most still remains unknown.

Neural Science is fundamentally an experimental science, and the progress that has been made in neural science has been driven by experimental observations and breakthroughs in experimental techniques — ranging from and including

- **Staining methods** in anatomy, such as the Golgi method discovered by Camillo Golgi in 1873 and used so effectively in neural science by Ramon y Cajal, for which they were jointly awarded the Nobel Prize in Physiology or Medicine in 1906;

- **Physiological methods**, such as

    - Single electrode extra-cellular measurements that count voltage spikes of single neurons,
    - Single electrode intra-cellular measurements that measure voltages across membranes of individual cell bodies,
    - Multi-electrode extra-cellular methods that count spike from collections of neurons at different spatial locations in the cortex;

- **Optical imaging methods** provide high resolution measurements of cortical activity across regions of the cortex, including
    - *Intrinsic imaging* that does not use dyes, and thus is not toxic,
    - *Voltage sensitive dyes imaging* with dyes of limited toxicity, for higher spatial and temporal resolution;
    - *Two photon imaging methods*, with remarkable spatial-temporal resolution of individual synapses and their plasticity;
    - *Non-invasive imaging* that allows the human brain to be imaged as it functions, including
        - *fMRI*,
        - *MEG imaging*,
        - *EEG imaging*;
- And **optical genetics**.

Although an experimental science, theory and theoretical modeling can help neuroscientists understand, interpret, and design experiments. In fact, without theory to guide it, pure experimental observation would fall far short of understanding the cortical network — its functions and the biological mechanisms by which it performs these functions.

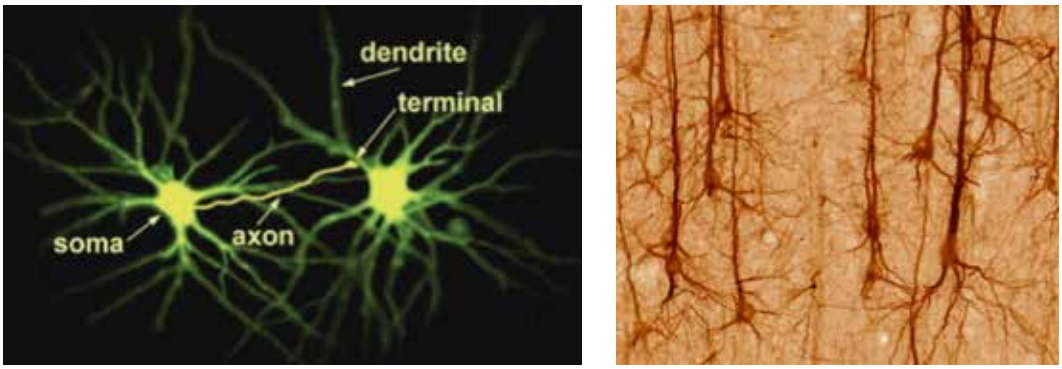

Figure 1: Neurons and their components (soma, axon, dendrite).
Left – Two photon image of two neurons; Right – staining of several neurons. Copyright held by University of California at Davis, and used for the sole purpose of this educational review web publication.

The basic components of the cortical network are neurons and their connections. Neurons are cells that are extended spatially (see Fig 1), each with a cell body (*soma*), an *axon* that projects to other neurons, and a complex dendritic arbor that receives axonal projections from other neurons. Neurons come in two categories, *excitatory* and *inhibitory*, that respectively excite/inhibit the neurons to which they project. Moreover, there are a number of distinct classes of excitatory neurons, and an even larger number of distinct classes of inhibitory neurons. These distinct classes of neurons primarily differ through neuronal anatomies, as discovered historically through Golgi staining – differences that presumably result in different biological functions and mechanisms that are just beginning to be understood. Neurons are predominantly connected through *synapses* (see Fig 2), of which there are $O(10^3)$ synaptic connections per neuron. At a synaptic connection, increases in voltage differences between the inside and outside of the presynaptic neuron (actually, its axon) induce the

release of neurotransmitters, which in turn open channels for ions to enter the postsynaptic cell and induce a current change in the postsynaptic dendrite. In this manner, chemical-electrical processes enable voltage information to be passed from the presynaptic neuron to the postsynaptic neuron. Thus, the cortical network consists in a huge number of excitatory and inhibitory neurons, connected through $O(10^3)$ more synapses than neurons. It is this cortical network that so efficiently performs the vast array of functions described above.

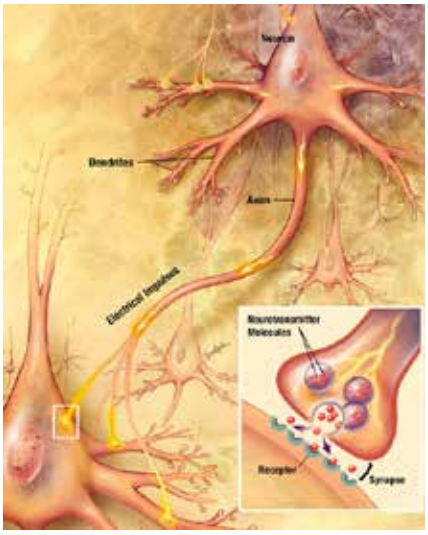

Figure 2: Schematic of a synapse between an axon of one (projecting) neuron and a dendrite of one (receiving) neuron.

In our work, we develop large-scale, biologically realistic, mathematical models of cortical regions of the mammalian brain. We emphasize that our models are biologically realistic — large scale and strongly constrained by biological measurements of the biophysics of individual neurons, their synaptic coupling, and their system properties. Most of our work models the mammalian primary visual cortex — the "front end" of the visual pathway through the visual

regions of the cortex that process visual information (See Fig 3). We then use the methods of modern applied mathematics to reduce these mathematical models, and finally study the reduced models computationally. We emphasize that we develop a *single model* (one set of parameter choices) of the primary visual cortex, and then use this single model to reproduce a large set of laboratory experiments.

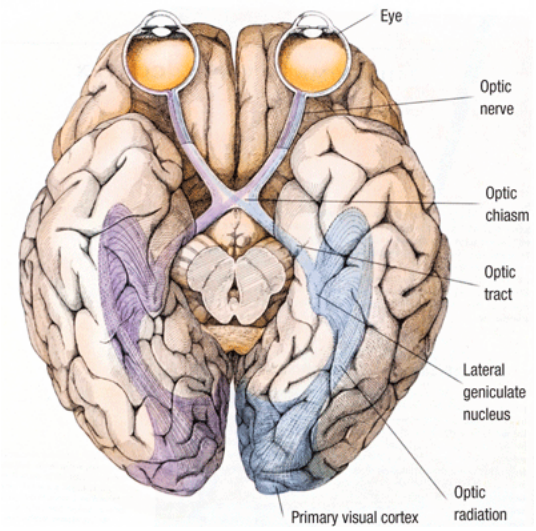

Figure 3: The visual pathway (retina→LGN→V1), from D. Hubel[31]

Without *large-scale computational models*, theorists are restricted either i) to studying one (or a few) neurons in isolation from the rest of the cortical system, or ii) to developing highly idealized abstract representations of possible cortical mechanisms whose validity for the real cortical system is either too general to inform, or impossible to validate biologically.

With *large-scale computational models*, one can check if highly idealized abstract mechanisms can be realized concretely and explicitly in biologically realistic models, and one can verify that these large-scale models perform in manners consistent with experiment. The assumptions in the construction of the

model can then be systematically changed and adjusted, to further understand exactly how the computational model works — thus identifying for the computational model the mechanisms that underlie the model's performance. Since the model is large-scale and biologically realistic, it is then reasonable to predict that these are the same mechanisms that underlie the real biological cortex.

In the remainder of this essay, we will:

1. Summarize our large-scale model of the primary visual cortex;
2. Briefly describe some of the biological properties of the cortical system, and the mechanisms that underlie these properties — including
    a. The "high-conductance" high gain operating point of the cortex;
    b. The sparsity of neuronal connections and its relation to the stability of the high-gain operating point of the visual cortex;
    c. Properties of neurons within network, including
        i. The linear/nonlinear characteristics of simple/complex cells;
        ii. The selectivity of individual neurons in V1 to the orientation of edges in the visual stimulation;
    d. Ordered versus disordered maps of cellular preferences across the cortical layer;
    e. Space-time patterns of cortical activity, including
        i. Spontaneous Cortical Activity
        ii. Cortical Activity of certain optical illusions.

3. In the conclusion, we will briefly summarize the properties of the operating point of the model cortex.

2. A Large-Scale Model of a Layer of the Primary Visual Cortex (V1)

We have chosen to model a part of the mammalian (rodent, cat, or monkey) cortex called the *primary visual cortex* (or area V1), which is the first cortical region along the visual pathway (Fig 3) to receive and process visual information from the retina. The primary visual cortex, as all cortical regions, is a layered structure of several cm^2 in lateral area[1], which is located near the rear of the skull (see Fig 4). This location at the rear of the skull makes the primary visual cortex accessible to physiological measurements with electrodes, as well as accessible to imaging procedures; thus, there is a wealth of experimental information about the primary visual cortex. We chose to study the visual cortex because of the importance of vision to mammalian perception, and the primary visual cortex because of the generally accepted belief that it is the first region along the visual pathway in which the mammalian brain processes visual information (see Fig 3)[2]. Initially [Refs(ix, i, ii)] we modeled one of the "entry layers" of the primary visual cortex, layer 4Cα. We chose to study layer 4Cα because, as an entry layer, it receives visual information directly from the retina through the LGN, with little feedback from other layers in V1 and from higher cortical regions; thus, feedback in 4Cα is largely restricted to "recurrent feedback" from and between neurons within layer 4Cα itself. Later [Refs(iii,iv)], as we refined and improved our model, we extended it to an effective single layer model which could apply to a superficial layer (such as layer 2/3) as well as an input layer (such as 4Cα).

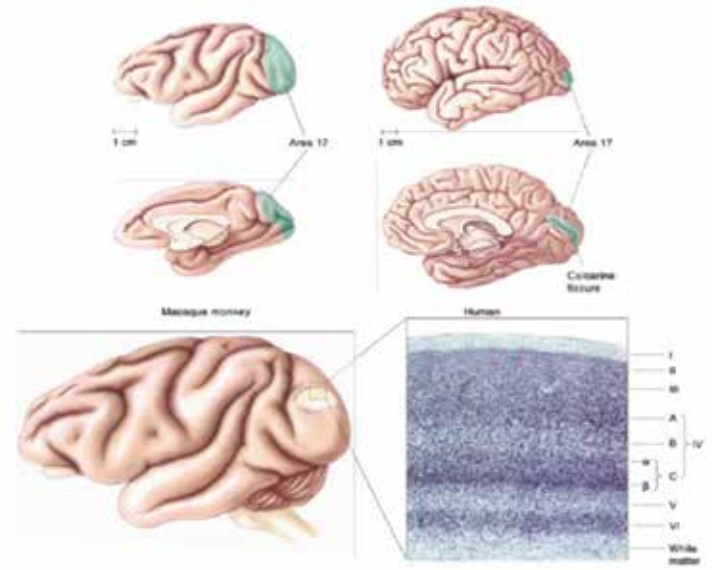

Figure 4: The primary visual cortex (V1 or Area 17).

Upper left – Macaque; Upper right – Human; Lower panels – staining showing the layered structure of V1 for macaque. The numbering – I is superficial and VI is interior. Reproduced from Neuroscience: Exploring the Brain, 4th edition, Mark F. Bear, Barry W. Connors, and Michael A. Paradiso, Wolters Kluwer (2016). The copyright is held by Wolters Kluwer, and used here with permission for the sole purpose of this educational review.

Neurons are the fundamental components of cortical networks. Neurons are spatially extended cells (Fig 1), with a cell body and both axonal and dendritic arbors that emanate from the cell body and along which the cell sends information (along axons) and receives information (along dendrites). Pairs of neurons are connected at "*synapses*," which are located at the intersection of the axon of the presynaptic (or sending) neuron with the dendrite of the postsynaptic (or receiving) neuron (Fig 2). Information is carried by intense voltage pulses, called "spikes", which are very short in temporal duration [O(3-5 ms)]. Spikes are generated at the cell body, propagate "outward" along its axon, and cause

voltage information to be passed from the presynaptic axon to the postsynaptic dendrite through a combination of chemical and electrical processes. Actually, information is carried by a "spike train", which is a temporal sequence of spikes generated at the cell body. Information is believed to be carried either by the "spike rate" of the train (the number of spikes generated by the cell body per second), or statistically by the inter-spike times of the spikes in the train. Because the spikes are so narrow, the spike train can be viewed as a binary sequence that describes the presence or absence of a spike as a function of time.[3]

As we are interested in the behavior of a cortical system containing a very large number of neurons, we opt to represent each neuron in a highly idealized fashion — as an "integrate and fire" point neuron. That is, we idealize each spatially extended neuron as a point in space, at which the voltage is governed by a version of Ohm's law [voltage = current x resistance]:

$$C\frac{d}{dt}V=-g_L(V-E_L)-g_E(t)[V-E_E)-g_I(t)[V-E_I],0<t<\tau_1,\tau_1+t_R<t<\tau_2,\text{etc}$$

$$V(\tau 1) = V_T;$$

$$V(\tau_1 + t_R) = V_R$$

Here C denotes the capacitance of the cell's membrane, $V(t)$ denotes the voltage across the cell's membrane, and the g's denotes different conductances (inverse resistances); g_L can be thought of as the membrane conductance, and $g_E(t)$ and $g_I(t)$ denote the temporal profiles of the excitatory and inhibitory[4] synaptic conductances that impinge upon this neuron from other neurons in

the cortical layer or from other regions along the visual pathway; the parameters E_L, E_E, and E_I are the values of "reversal potentials" (see below); the parameter VT denotes the "firing threshold voltage" and V_R denotes the "reset potential", where $V_T > V_R$; $\{\tau_j, j=1,2,3,...\}$ denote the spike times; and t_R denotes a short delay (spike width plus an additional delay call the "refractory period") after which the voltage is reset to V_R and begins evolve again.

The integrate and fire representation does not capture the temporal profile of the voltage spike itself; rather, it takes advantage of the following "approximate facts": i) the voltage spike is very narrow in time; ii) the voltage spike has approximately the same temporal profile each time the neuron spikes; iii) there is an approximate "spiking threshold" V_T which, when the neuron's voltage reaches $V(t=\tau) = V_T$, the neuron fires a spike; and iv) shortly (t_R = a few ms) after which the temporal profile of the voltage spike resets the voltage $V(\tau+t_R) = V_R$ ($<V_T$). Thus, $V(t=\tau) = V_T$ defines the "spike time τ", after which the integrate and fire voltage V is "re-initialized" at V_R and integrated forward in time until the next spike time. The increasing sequence of spike times (τ_1, τ_2, τ_3, τ_4, ...) label the times at which a spike occurs in the neuron's binary spike train.

The largest error in the integrate and fire representation of the neuronal voltage $V(t)$ is not the omission of the temporal profile of the voltage spike; rather, it is the point neuron approximation which eliminates the dendritic structure. Since the synapses occur at different locations along the dendritic tree, these locations affect the temporal dynamics of the postsynaptic potential $V(t)$ in

significant manners that cannot be captured by the point neuron idealization. To capture these effects, a spatially extended model of the neuron is needed, and is missing from the simple integrate and fire representation that we currently use in our model.

There are a huge number of neurons in each cortical layer. For example, in input layer 4Cα of V1, there are O(16,000) neurons per mm^2: moreover, these neurons come in a wide variety of types that differ in size, shape, and bio-physical characteristics. Some neurons, when they fire, excite the neurons to which they project to fire more rapidly, while others inhibit the firing of the neurons to which they project. The classification into "*excitatory*" or "*inhibitory*" is perhaps the most important classification of neuronal type, although there are many distinct classes of excitatory and of inhibitory neurons.

In the model, connections between cortical neurons in the layer of V1 are modeled through the conductance temporal profiles as pictured in Fig 5. When a voltage spike from a presynaptic excitatory neuron arrives at a post synaptic neuron, it activates two types of receptors, called AMPA and NMDA, which generate conductance temporal profiles on two distinct time scales – fast O(5ms) for AMPA and slow O(80ms) for NMDA. Similarly, when a voltage spike from a presynaptic inhibitory neuron arrives at a post synaptic neuron, it activates two types of receptors, called $GABA_a$ and $GABA_b$, which generate conductance temporal profiles on two time scales – fast O(5ms) for GABAa and somewhat slower for $GABA_b$. The presence of these distinct time scales influences fundamentally the model's response properties.

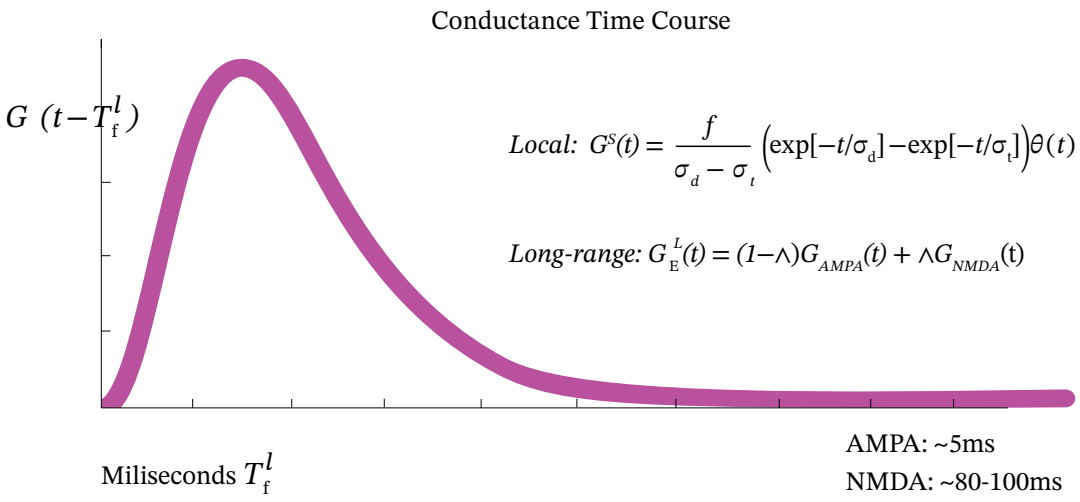

Figure 5: Excitatory Conductance Time Courses: Rapid – AMPA; Slow – NDMA

Our large-scale computational model represents a local patch of a single layer of V1, of lateral area ranging from 1mm^2 to 25mm^2. Each neuron in this patch is labeled by its type ("E" for excitatory and "I" for inhibitory), and by a two dimensional index "j" that labels the neuron's location within the patch. Thus, the model is a coupled system of 16,000/mm^2 integrate and fire point neurons, 75% of which are excitatory and 25% inhibitory. These equations are of the form

$$C\frac{d}{dt}V_\sigma^j = -g_L(V_\sigma^j - E_L) - g_{\sigma E}^j(t)[V_\sigma^j - E_E] - g_{\sigma I}^j(t)[V_\sigma^j - E_I]$$

where $j = (j_1, j_2)$, $\sigma = (E,I)$, and where $g_{\sigma E}^j(t)$, $g_{\sigma I}^j(t)$ denote the conductances from excitatory-E/inhibitory-I neurons respectively that impinge upon the neuron at cortical position j of excitatory/inhibitory type σ. The "*leak conductance*" g_L is a constant. The excitatory conductance $g_{\sigma E}^j(t)$ has the form

$$g^{j}_{\sigma E}(\mathrm{t}) = f^{j}_{lgn}(t) + f^{noise}_{E}(t) + S^{S}_{\sigma E}\sum_{k} K^{\sigma E}_{j\text{-}k}\sum_{l} G_{AMPA}(t - t^{E,k}_{l}) + S^{L}_{\sigma E}\sum_{k} L^{\sigma E}_{j\text{-}k}\sum_{l} G_{NMDA}(t - t^{E,k}_{l}),$$

and the inhibitory conductance $g^{j}_{\sigma I}(\mathrm{t})$ has the form

$$g^{j}_{\sigma I}(t) = f^{noise}_{I}(t) + S_{\sigma I}\sum_{k} M^{\sigma I}_{j\text{-}k}\sum_{l}[G_{\mathrm{GABA\text{-}A}}(t - t^{I,k}_{l}) + G_{GABA\text{-}B}(t - t^{I,k}_{l})].$$

Here $f^{noise}_{E}(\mathrm{t})$ and $f^{noise}_{I}(t)$ represent noise terms from excitation and inhibition, respectively, each of which is modeled by a Poisson spiking process; $S^{S}_{\sigma E}$, $S^{L}_{\sigma E}$ and $S_{\sigma I}$ denote parameters that set the cortical-cortical coupling strengths of short range excitation, long range excitation, and inhibition, respectively. $K^{\sigma E}_{j\text{-}k}$, $L^{\sigma E}_{j\text{-}k}$ and $M^{\sigma I}_{j\text{-}k}$ represent the pattern of spatial coupling from short range excitatory neurons, long range excitatory neurons, and from inhibitory neurons, respectively. The times $t^{E,k}_{l}$ and $t^{I,k}_{l}$ represent the l^{th} spike time of the k^{th} excitatory/inhibitory neuron, respectively. The terms of the form $\mathrm{G(t\text{-}t_{sp})}$ represent the prescribed temporal profiles of the conductances for times after the spike time $\mathrm{t_{sp}}$. There are two excitatory conductance temporal profiles {fast [O(3ms)] AMPA and slow[O(80ms)] NMDA} and two inhibitory temporal profiles {fast [O(3-5ms)] GABA-A and slow [O(7ms)] GABA-B —

$$G_{\sigma}(t) = \frac{1}{6\tau_{\sigma}}\left(\frac{t}{\tau_{\sigma}}\right)^{3}\exp(-t/\tau_{\sigma})\,\theta(t)$$

$$G_{NMDA}(t) = \frac{1}{\tau_{1} - \tau_{2}}\left[\exp(-t/\tau_{1}) - \exp(-t/\tau_{2})\right]\theta(t)$$

Here $\theta(t)$ denotes the Heaviside function ($\theta(t) = 1$ for $t > 0$, $= 0$ otherwise);

$\tau_1 = 80$ ms, $\tau_2 = 2$ ms, $\sigma =$ AMPA, $GABA_a$, or $GABA_b$, $\tau_{AMPA} = 1$ ms, $\tau_{GABA\text{-}a} = 1.67$ ms, and $\tau_{GABA\text{-}b} = 7$ms.

In the remainder of this section, we describe in more detail the spatial and temporal profiles of the cortical-cortical conductances and the LGN drive. For more detailed information, we refer the reader to the published material [Refs [5],[6],[7],[8]].

The conductances describe the coupling between other neurons in the layer of V1 and from neurons in the LGN that transmit the visual information from the retina (and thus from the outside world). Consider the visual pathway, as depicted in Fig 3. It connects "visual space" where the visual stimuli occur, with "cortical space" where the neurons in the layer reside — a pathway through the retina and LGN, to V1. Each cortical neuron is driven by the summed spike trains of a selected small number of excitatory cells in the LGN, each one of which is described by a modulated Poisson process, whose modulated parameter is defined by a space-time linear filter convolved with the visual stimuli *I(x,s)*,

$$R_k^{\pm}(t) = H\{R_B \pm \iint G(t-s)\mathrm{A}(|x_k - x|)I(\mathrm{x},\mathrm{s})\mathrm{d}^2\mathrm{x}\mathrm{d}\mathrm{s}\}\ .$$

Here $H(t) = 1$, $t > 0$; $= 0$ otherwise; and R_B (= 20 spikes per sec) denotes the spontaneous firing rate of individual LGN neurons in the absence of visual stimulation. $G(t\text{-}s) = G_{AMPA}(t\text{-}s)$. $A(|x_k - x|)$ is a circularly symmetric difference of Gaussians, centered in visual space on the receptive field location x_k of this LGN neuron (that is, centered in visual space on that spatial location from which

the LGN neuron receives visual information). The half-width of the Gaussian filter captures the small spatial extent of the receptive field of the LGN neuron, which is circularly symmetric. The spatial scales are set so that $A(|x_k - x|)$ is positive near x_k and negative far from x_k. The two possible signs capture the "on-off" or "off-on" property of LGN cells; namely, whether the cell responds to an increase in stimulus contrast from background at the center of the receptive field, and to a decrease in stimulus contrast in the circularly symmetric surround (an "on LGN cell"); or vice- versa (an "off LGN cell"). The temporal profile *G(t-s)* in eq. 0.4 was derived from experimental data on LGN neurons [Gielen & ,van Gisbargen & Vendrik, Biol. Cybern. 40, 157-170 (1981); Benardete & Kaplan, Benardete's Ph D Dissertation, Rockefeller Univ, New York (1994)].

LGN cells, with their circularly symmetric receptive fields, respond to sharp edges in the visual patterns in the same way, irrespective of the orientation of these edges. This is in contrast to neurons in V1, which have receptive fields which are considerably larger than the receptive fields of the LGN neurons, and which respond preferentially to particular orientations of edges within the visual stimuli. That individual V1 neurons respond preferentially to specific orientations of stimuli, which differ from neuron to neuron, and thus enable V1 neurons to serve of "orientation detectors of edges in visual scenes", was discovered in 1959 by David Hubel and Torsten Wiesel[9] [10] (for which they were awarded the Nobel prize in 1981). Later it was established that each layer of V1 was tiled by an organized pattern of orientation preference, with each

"tile" of size O(0.5 mm x 0.5mm), and that these organized patterns of orientation preference extend as "orientation columns" throughout the depth of the multiple cortical layers which comprise V1. (See Fig 6 for the beautiful intrinsic optical imagining observation of "orientation preference map" that tiles a superficial layer of V1).

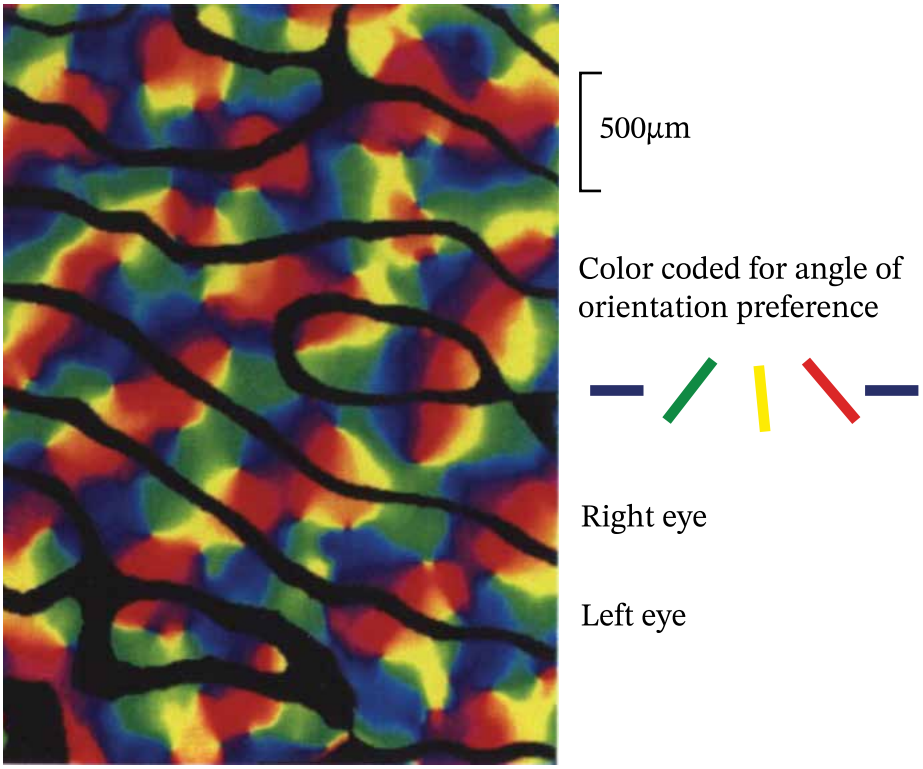

Figure 6: Optical Image of the Map of Orientation Preference for a Superficial Cortical Layer (2/3) of Macaque.

Color coding denotes the angle of orientation preference. The black curves denote the borders of ocular dominance "columns", across which input from the retina switches dominance from left to right eye. Note the "orientation pinwheel" singularities, and the scale of the individual "tiles", each containing one singularity – 500 μm x 500 μm. (From G. Blasdel, Ref [17, 18].) Copyright held by J. Neuroscience, and used for the sole purpose of this educational review web publication.

The orientation preference of individual cells in the entry layer is placed into the model as follows: A small number of LGN cells project to each cortical cell; and the receptive field of each cortical cell is a superposition of the receptive fields of each LGN neuron in the collection projecting to it, with the receptive field centers of the LGN cells displaced slightly from the receptive field center of the cortical cell, forming a structured pattern for the receptive field of the cortical cell (see Fig 7). The receptive field of a cortical cell is not circularly symmetric (as each of its LGN components), but rather has a spatial pattern imposed on it by the displacements together with the selection of "on" cells and "off" cells which comprise the collection (again, see Fig 7). Thus, the receptive field of each cortical cell is spatially larger than that of its LGN components; and it has a "sub-field" orientation structure with provides each V1 cell with its orientation preference.

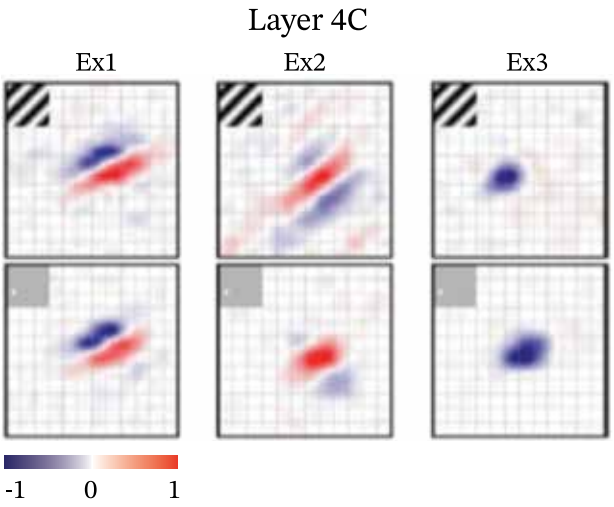

Figure 7: Receptive Field Structure for Input Layer (4C) in Macaque.

Receptive fields for three layer 4C neurons. The visual stimulation for the upper row — random gratings; for the bottom row — sparse random flashed dots. (From Ref[32], reprinted from PNAS which holds the copyright.)

Turning to the map of orientation preference that is an ordered tiling of the superficial layer, with each tile of O(500μm x 500μm): (See Fig 6) The orientation preference of each cell within a tile smoothly changes as one cycles the "pinwheel singularity" at center of each tile. These "orientation pinwheels" tile the layer, and their orientation preferences extend throughout the depth of V1 to form "orientation columns." We build this ordered pinwheel tiling into the model by mapping the orientation preference of each cortical cell to its spatial location within the layer.

In this manner, the orientation preference of each cortical neuron and the ordered pinwheel map are hard wired into the model. Each entry layer cortical neuron receives a spike train from a collection of LGN neurons. The size, location, and orientation structure of the receptive field of each entry layer neuron is set by the individual locations of the receptive field centers and their designated "on" or "off" type of the LGN neurons in the collection of LGN neurons which project to the given cortical cell. And this orientation preference is set for each cell, respecting the ordered map of orientation preference that tiles the cortical layer.

In addition to orientation preference, V1 neurons also have spatial and temporal phase preferences, wavelength preferences, etc. These are not ordered within the pinwheel tiling, but experimentally observed to be random. They are put into the model through the receptive field construction as described above for orientation, but with random (rather than ordered) spatial maps.

Now we return to the description of the excitatory cortical-cortical conductances, that is to the spatial-temporal conductance profiles generated by the spiking of presynaptic excitatory cortical neurons. These come in two types — (i) temporally fast AMPA and (ii) slow NMDA. (The precise formulas are given above.)

The conductance time courses for fast inhibition G_{GABA-A} and slow inhibition G_{GABA-B} have the same form as G_{AMPA} , with time constant $\tau = 1.67$ms and 7ms, respectively.

We turn now to the spatial patterns of the cortical-cortical coupling, as represented by $K^{\sigma E}_{j\text{-}k}$, $L^{\sigma E}_{j\text{-}k}$ and $M^{\sigma I}_{j\text{-}k}$ for short range excitatory neurons, long range NMDA excitatory neurons, and inhibitory neurons, respectively. The "short range" kernels $K^{\sigma E}_{j\text{-}k}$ and $M^{\sigma I}_{j\text{-}k}$ are taken to be isotropic Gaussians, with length scales of 200 μm for excitation and 100 - 200 μm for inhibition, respectively. These length scales are consistent with anatomical studies. The length scales of the "long range" kernels $L^{\sigma E}_{j\text{-}k}$ are much longer, O(2000 μm), and more importantly are not spatially isotropic.

Rather, these long distance connections preferentially couple neurons of similar orientation preference.[11] See Fig 8, which shows this non-isotropic connectivity for tree shrew.

All kernels are normalized to unity; hence, the coupling strengths are solely described by the parameters S_{EE}, S_{IE}, S_{EI}, and S_{II}.

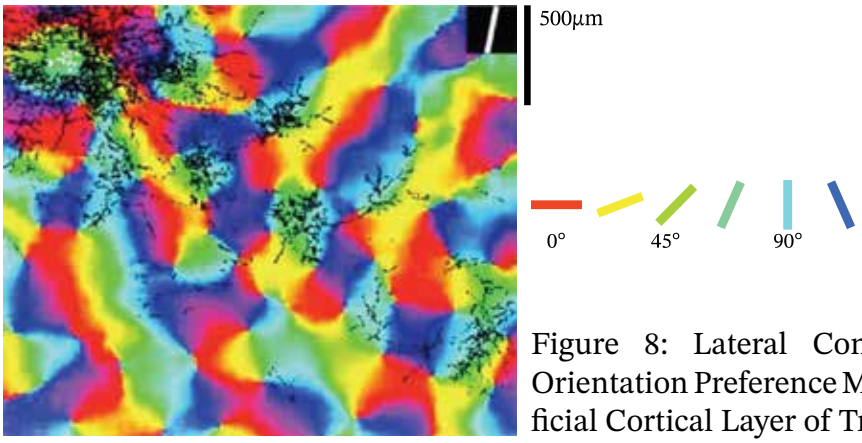

Figure 8: Lateral Connections and Orientation Preference Map for a Superficial Cortical Layer of Tree Shrew.

Connections, as displayed by a staining procedure, over the optical image of orientation preference map. The injection site for the dye is denoted in white; the neuronal connections to this injection site denoted in black. Note the injection site is in a "blue-green" region of orientation preference; the local (~500 μm) connections are rather isotropic, connecting to "all" preference angles; the long-range connection selectively connect to neurons of "like-orientation" ("blue-green" preference). From Ref [30]. Copyright held by J. Neuroscience, and used for the sole purpose of this educational review web publication.

The spatial couplings from the excitatory neurons have two important features in addition to the Gaussian fall-off with spatial distance. The first is that the spatial excitatory coupling is very sparse, with only 10-15% of the spatial connections active. In the model, which neurons are connected to each other is determined randomly from the outset, thus "hard-wiring" a sparse coupling architecture into the model. And second, we assume that those cells with fewer LGN afferents have more of their excitatory synapses taken up by cortical-cortical excitatory connections. Thus, each neuron's $S_{\sigma E}$ is inversely proportional to the its N_{LGN}, the number of its LGN afferents. A linear scaling of the coupling strengths insures this assumption. Finally, the "mean coupling strengths" are modified by a small random spread of values, to ensure qualitatively the diversity that is observed in both the "driven" and "spontaneous" firing rates.

3. Summary of Results from the Large Scale Model of a Layer of V1

In this section, we give brief overviews of the understanding of cortical mechanisms that have resulted from our large-scale computational model.

1. High Gain and Sparse Connectivity: In the neuroscience community, it is generally accepted that the primary visual cortex must operate in a regime of "high gain" if it is to be sufficiently sensitive to small changes in the visual stimuli. As described in Ref[6], in order for our large-scale computational model of V1 to operate reasonably in a high gain regime, the spatial connections between excitatory neurons must be sparse, and/or the synaptic connections must frequently experience synaptic failure. Otherwise, as the cortical network approaches a sufficiently high gain regime, the network activity becomes far too intense — entering if you like an "epileptic regime." We find that either sparse connectivity (probabilistically hardwired into the network initially) or synaptic failure (probabilistically and dynamically determined at each time a spike approaches a synapse) will stabilize the network in the high gain regime. For example, if we use only sparse connectivity, we find that the network performs well with only 10-15% of the neurons coupled to each other. The high gain regime, or operating point, itself is one of steep gain just below a region of steeper gain which supports simultaneously two stable operating points, one of low gain and a second of very high gain. The model performs best at a gain rate just below this "bi-stable domain." We emphasize that, within the model cortex, stability of this high gain region seems to demand sparse connectivity (or high synaptic failure).

2. High Gain Regime and Large Conductances: High gain is achieved through large excitatory and inhibitory conductances, and therefore large total conductance. This high gain regime holds for both the excitatory and inhibitory neurons. These high conductances result in effective relaxation time scales that are very short (compared, for example, to the "leakage" time scale. The equations take the form

$$C\frac{d}{dt}V_\sigma^j = -g_L(V_\sigma^j - E_L) - g_{\sigma E}^j(t)[V_\sigma^j - E_E] - g_{\sigma I}^j(t)[V_\sigma^j - E_I]$$

$$= -g_{\sigma T}^j(t)[V_\sigma^j - V_{\sigma,eff}^j(t)]$$

where

$$g_{\sigma T}^j(t) = g_L + g_{\sigma E}^j(t) + g_{\sigma I}^j(t),$$

$$V_{\sigma,eff}^j(t) = \frac{g_L E_L + g_{\sigma E}^j(t)E_E + g_{\sigma I}^j(t)E_I}{g_L + g_{\sigma E}^j(t) + g_{\sigma I}^j(t)}.$$

Thus, for large total conductance $g_{\sigma T}^j(t) >> g_L$, the voltage $V_\sigma^j(t)$ quickly relaxes to the effective reversal potential

$$V_\sigma^j(t) \to V_{\sigma,eff}^j(t)$$

This approximation is valid because the computational cortex operates in a high conductance regime. This rapid relaxation of the voltage to the effective reversal potential $V_{\sigma,eff}^j(t)$ is extremely useful for understanding the mechanisms of the computational cortex, by providing an analytical representation of

the voltage $V_{\sigma}^{j}(t)$ explicitly in terms of the conductances[12]. This approximation unveils the mechanisms that underlie the network responses of the computational model cortex to stimulation, and thus provides insight into the possible mechanisms that underlie the biological cortex's response properties.

3. Properties of individual cells *within network*: Individual cells, within network, respond preferentially to properties of the stimuli. For example, cells respond preferentially to specific angles of orientation (of edges) in the stimuli; that is, a given cell will fire spikes strongly to a given angle of orientation of the stimuli, and when that angle is changed, it fires less — firing very little to angles nearly 90 degrees from the optimal angle. Moreover, individual cells' selectivity to the preferred angle (as measured by the half-width of the cells firing rate as a function of the angle of stimulation) differs from cell to cell — from "sharply tuned" cells with narrow half-widths that are good "orientation detectors" to cells that are "broadly tuned" with large half-widths that are poor "orientation detectors." This orientation specificity of individual neurons extends to specificity for other stimulus properties such as phase (in time and space) and spatial and temporal wavelength.

In addition, individual cells within network respond to stimuli with varying degrees of linearity. Some cells respond to sinusoidal stimulation linearly as a sinusoid (at the same temporal and spatial frequency as the stimulus), while others respond very nonlinearly, with distorted wave forms at multiple spatial and temporal harmonics. In the literature, cells that respond linearly are called

"simple cells", while cells that respond nonlinearly are termed "complex cells."

These individual cellular response properties do not vary with time, or with repetitive stimulation. That is, to a first approximation, these specific individual neuronal response properties seem to be fixed in time.

Moreover, there is no evidence that these individual cellular response properties are the result of bio-physical properties of the individual cells. Indeed, in isolation from each other, the cells appear to be bio-physically identical. Thus, the distinct response properties of individual cells within network must result from network properties such as the coupling or "wiring" architecture of the neurons within the network, together with their coupling strengths.

Computational models can identify coupling architectures that will achieve the observed response properties of individual cells within network. One important issue is the extent to which observed response properties are caused by the "feed-forward" nature of the network, which require "feed-back" from recurrent connections within the cortical layer, and which depend upon "feed-back" connections between the cortical layer and other cortical layers or other "higher" regions of the cortex. Our single input layer model can be used to investigate (in the model) the relative importance of "feed-forward" and "recurrent within layer" connections.

Another important issue is the degree to which "preferential coupling" of neurons is required to produce the observed response properties of individual

neurons. The idea for "preferential coupling" dates back to Donald Hebb, who in 1949 formulated the principle that "neurons that fire together, wire together."[13] The idea is that, as the synaptic coupling architecture is set up during the development stage of the young mammal, the synaptic coupling strength is increased between neurons which fire together; thus, the coupling strength between two neurons with similar stimulus preference, such as similar orientation preference, will strengthen and become much stronger than the coupling strengths between neurons with different stimulus (orientation) preferences. Thus, Hebbs' principle, if correct, would imply that in the mature mammal, neurons with similar preferences are more strongly coupled than neurons with different preferences. Our model of an input layer of V1 for macaque monkey show that preferential coupling is not necessary to achieve the observed cellular responses; or more accurately, local (within 500 μm) coupling architecture that is not "preferential", but rather isotropic, decreasing in strength isotropically with distance, is sufficient to produce the observed cellular responses.

As an example, consider orientation. LGN cells have circularly symmetric receptive fields and thus no orientation preference. However, as described in Section 2, a small collection of LGN neurons projecting collectively to one V1 neuron can impart a preferred orientation to that cortical neuron. In our model, we assume each cortical neuron is driven by O(15-20) LGN neurons. By choosing patterns for each collection of LGN circularly symmetric receptive fields, cortical cells in the input layer 4Cα have orientation preference, orienta-

tion selectivity, and receptive fields with sub-field structure — all in agreement with experiment (See Fig 7). Thus, we find in the model that the orientation preference of the V1 neuron is set by the feed- forward LGN drive, and that the feed-forward orientation selectivity is modified, but only modestly, by the recurrent feedback connections within layer 4Cα. Thus, the model is primarily feed-forward with regard to the orientation tuning of cortical cells, with the local circuitry isotropic without preferential microcircuits[14] (See also Ref [6].)

However, the original model had strong, well focused for orientation, LGN drive — with O(15-20) LGN neurons projecting to each cortical input layer neuron. Recent detailed analysis[15] of the existing anatomical data has shown that, for macaque monkey, the LGN drive is much weaker and less focused for orientation — with only O(5) LGN neurons projecting to each cortical input layer neuron. In that work, Chariker, Shapley and Young have constructed a large scale computational model with realistic weak cortical drive. In their model, the orientation preference is again set by the feed-forward LGN drive; however, the orientation selectivity is modified (usually enhanced) significantly by the recurrent feed-back connections.

Turning to the linear/nonlinear properties of cortical neurons, we consider simple/complex cells within the computational model. First, consider complex cells, with their nonlinear relationship to the visual stimulus. While there are several realizations of these nonlinear response properties, perhaps the most sensitive is *frequency doubling* — the cells respond at twice the temporal

frequency of a sinusoidal (contrast reversal, standing, grating) drive. Within our model, complex cells are achieved with two modeling assumptions: i) those cells with fewer LGN afferents (N_{LGN}) have more of their excitatory synapses taken up by cortical-cortical excitatory connections; and ii) each neuron's cortical-cortical recurrent excitatory coupling strengths are inversely proportional to the number of LGN neurons afferents (N_{LGN}). That is, the weaker the LGN drive, the stronger the recurrent drive. Within the model "pure" complex cells receive no LGN drive and thus the strongest recurrent drive. In Ref[5], we show that this model produces a reasonable proportion of complex cells with the proper nonlinear response properties such as frequency doubling, and with orientation selectivity properties which are in modest agreement with experimental observations.

Perhaps more surprising than the nonlinear behavior of complex cells is the very existence of simple cells — cells in the nonlinear cortex that behave so linearly in response to the visual stimulus. Recall these cortical cells are driven by the LGN and recurrently by other cells in the layer. Both the LGN drive and the cortical network are very nonlinear. How, acting together, do they produce responses as linear as those observed in simple cells? Although too detailed mathematically to present in this essay, in Ref [16] we show, through both mathematical analysis and computer simulation, exactly how the nonlinear cortical-cortical inhibition cancels the nonlinearity in the excitatory LGN drive, and produces in the computational model a response that is linear and in close agreement with experimental observations of the behavior of simple cells.

Moreover, Ref [6] shows that our model contains a realistic distribution of both simple and complex cells.

4. Role of ordered vs disordered preference maps: Feature preferences, such as orientation preference, tile the cortical layer as a "feature map". These maps can be ordered (as orientation maps in Macaque monkey[17,18]) or disordered (as the random "salt and pepper" map of spatial frequency preference[19]). As described in Section 2: The orientation preference of cells in each layer of V1 form an ordered tiling of the layer, with each tile of O(500μm x 500μm). (See Fig 6) The orientation preference of each cell within a tile smoothly changes as one cycles the center of the tile, at which an orientation preference "singularity" exists. These "orientation pinwheels", with their pinwheel singularities, tile the layer. How does this orientation mapping affect the response properties of cells within the layer?

Within the model Ref[6,13], we can clearly describe the effect of the ordered map of orientation preference on the orientation response properties of a neuron relative to its position within the tile. Consider first a simple cell, *located near a pinwheel singularity.* It will receive strong excitation from the LGN drive (from which it derives its orientation preference), moderate excitation from other nearby excitatory neurons within the tile (at all orientation preferences), and strong inhibition from nearby inhibitory neurons (again, at all orientation preferences). These recurrent cortical-cortical drives will sum, to average over all orientation preferences; thus, the recurrent excitation and the recurrent

inhibition reaching a neuron will have no orientation preference. Moreover, inhibition will dominate and cancel excitations away from the preferred angle set by the LGN drive, resulting in simple cells near pinwheel centers that are very selective for orientation. In contrast, complex cells near pinwheel centers will not be tuned for orientation as they receive no LGN drive and a summed recurrent drive that is independent of orientation.[20]

Far from a pinwheel singularity, all nearby cells have similar orientation preferences; hence, both simple and complex cells will be quite well tuned for orientation — a result that holds for both densely and sparsely connected networks.

Our large-scale computational model shows that all of the results described above can be obtained *without* preferential coupling architectures. In the models, whether the neurons are coupled densely or sparsely, the coupling is isotropic, falling off with distance between neurons; there is no explicit preferential coupling between neurons with similar feature preferences, and hence, no explicit implementation of Hebb's principle. On the other hand, ordered maps can provide an "effective" preferential coupling architecture. Consider the orientation map described above. Except for those near pinwheel singularities, most neurons reside in "iso-orientation" domains, and thus the coupling to nearby neurons will be between neurons with similar orientation preference. The ordered map itself provides an implicit realization of preferential coupling, at least with regard to orientation.

While the map of orientation preference is ordered for cat and Macaque, it is a disordered random "salt and pepper" map for mouse visual cortex. Careful measurements have established that, even in the absence of an ordered orientation map, neurons in the mouse visual cortex have very similar orientation selectivity and tuning properties to those of Macaque. (See ref[21,22] for recent measurements.) Moreover, optogenetic techniques on mouse provide information about the orientation preference of the cortical-cortical excitatory input to V1 neurons [23,24] that is not available for Macaque. These techniques show that the cortical-cortical excitatory input to each cell in V1 has orientation preference that is aligned with the orientation preference of the LGN input to each cell. In unpublished work[25], we develop a large-scale computational model of an input layer to mouse V1, and study which properties of individual cells within the network can be obtained without preferential coupling, and which properties require preferential coupling of micro- circuits which selectively couple neurons with similar orientation and temporal phase preferences. Oriented cortical-cortical excitatory input to an individual cell that is aligned with the orientation preference of the LGN input to that cell seems to require explicit preferential coupling — in accordance with Hebb's principle.

5. Spatial-temporal patterns of cortical activity: The discovery of voltage sensitive dyes with limited toxicity has enabled neural scientists to develop optical imaging techniques that can be used to measure directly cortical activity throughout rather wide regions of superficial layers of the cortex, with

excellent spatial and temporal resolution[26],[27]. The techniques are invasive, and can only be used on animals, with most of the experiments performed on cats. The first observations were performed on anesthetized cats, observing spontaneous activity (in the absence of visual stimulation) in superficial layers of primary visual cortex[28],[29]. The observed spatial-temporal patterns of spontaneous activity were remarkably similar to those patterns observed under stimulation by oriented gratings, except that the patterns were meta-stable. A spontaneous pattern would appear that is similar to a pattern stimulated by a grating at one angle of orientation, then disappear and be replaced by the appearance of a pattern at a different angle of orientation. An irregular sequence of patterns is thus generated, with both the *angles* of orientation and the *times* of the jumps to a new angle very irregular. What mechanism creates these seemingly random sequential patterns of spontaneous cortical activity?

The investigation of this phenomena with a large-scale computational model required extending our local model (of 1 mm2 of cortical area holding approximately 4 pinwheels) to a more global model (of ~16 mm^2 of cortical area containing 64 pinwheels). In the global model, long-range synaptic connections had to be included.

Anatomical studies[30] have shown striking differences between short-range and long-range synaptic connections, as summarized for tree shrew in Fig 8 from Ref [29]. This figure traces the local and long-range excitatory synaptic connections to those neurons in a dye injected region of small lateral area — all

superimposed on the ordered map of orientation preference. The small injection region contains neurons whose angles of orientation preference are labeled as blue-green. The tracings show that, locally, the neurons in the injection region are isotropically connected to neurons with all angles of orientation preference, but that the long- range connections preferentially couple neurons with the same "blue-green" angles as the injection region. Other studies show that these long range connections definitely contain the "slow" NMDA excitatory receptors, that act on long time scales of O(80ms).

In Ref [7] we extend the computational model to an effective model of a superficial layer of V1, on a global scale of 16 mm^2, incorporating long-range couplings that prefer similar angles of orientation preference and act on the long time scales of NMDA. The spontaneous activity that we observe in this computational model has an irregular sequence of meta-stable oriented states that is very similar to the observed spontaneous cortical activity in anesthetized cat — irregular jumping between oriented states that occurs on the observed time scales. Moreover, although too detailed to describe here, analysis of the computational model, grounded upon the large conductance relaxation to the "effective reversal potential" described above, unveils the specific mechanisms by which the model produces these meta-stable state: The irregular patterns of cortical activity are initiated by small groups of excitatory neurons firing together over short time frames, coupled to neurons with similar orientation preference a long distance away through the orientation selective long-range

connections. Thus, oriented patterns of cortical activity are generated, with the NMDA long-time scale producing the time scale of the temporal jumps in the patterns. Although the initial sequential firings are rapidly quenched due to inhibition, the long range cortical activity is maintained over O(100ms) time scales through long-range connections with NMDA receptors.

Optical illusions are frequently used by neural scientists to understand the operation of the cortex. Although illusionary with respect to perceptions of the outside world, these are actual patterns of cortical activity; thus, understanding the mechanisms that generate the cortical activity associated with optical illusions enables us to better understand the mechanisms under which the cortex operates. Most illusions involve higher order cortical activity such as attention; thus, they can't be explained solely with an input region such as the primary visual cortex, but must involve higher order regions. However, there are a few illusions that seem "pre-attentive", and thus might be explainable within the primary visual cortex. One such is the *"pre-attentive" line motion illusion*. As shown in Fig 9, a stimulus that consists of (i) a square of light flashed on for a short time, then (ii) flashed off and held off for a short period of time, and finally (iii) followed by the flash of long bar of light (of width equal to the square) is a stimulation that is perceived as either a square growing into the bar, or a moving square — although in reality there is absolutely no motion or growth in the visual stimulus.

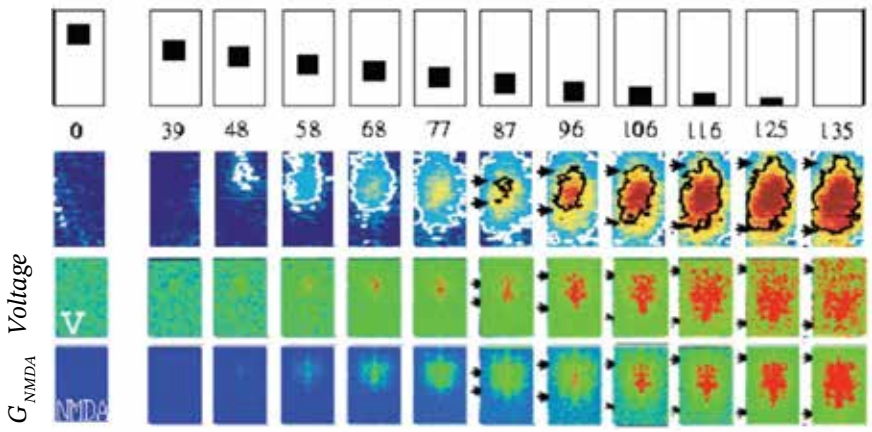

Figure 9a: Cortical Activity in a Superficial Layer of Cat V1, that is Visually Stimulated by a Moving Square. Top row: Schematic of the moving square, as a function of time labeled in ms; Second row: Voltage sensitive dye images of the cortical activity; Third row: Computer simulations of the cortical activity as represented by voltage; Bottom row: Computer simulations of the cortical activity as represented by NMDA conductance profile.

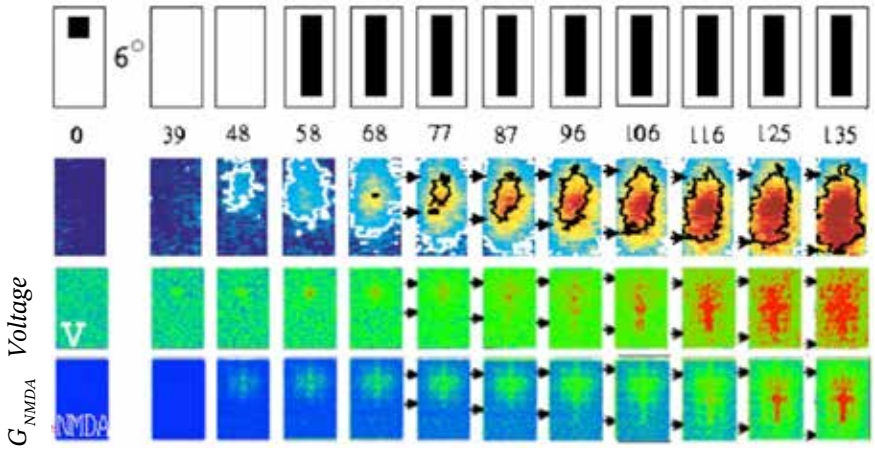

Figure 9b: Cortical Activity in a Superficial Layer of Cat V1, that is Caused by Visual Stimulation by the "Line Motion Illusion". Top row: Schematic of the visual stimulus, as a function of time labeled in ms; Second row: Voltage sensitive dye images of the cortical activity; Third row: Computer simulations of the cortical activity as represented by voltage; Bottom row: Computer simulations of the cortical activity as represented by NMDA conductance profile. Note the striking agreement between the experimental images of the cortical activity in Figures 9a and 9b; note also the good agreement between the experimentally observed cortical activity and that generated by the numerical simulations. The experimental results are from Ref [29] (reprinted with permission from Nature, which holds the copyright) and the numerical simulations from Ref[7] (reprinted from PNAS, which holds the copyright).

Experiments in Grinvald's Laboratory used voltage sensitive dye optical imagining to observe the patterns of cortical activity in anesthetized cat generated by the line-motion illusion stimulus — comparing the patterns generated by a growing bar with that of the static flashed square and bar of the line motion illusion. The observed cortical activity was identical under the two different stimulations.

In Ref[8] we stimulated a large scale model of an effective superficial layer of V1 with these line-motion illusion stimuli. As in the experiment, the space-time cortical activity was identical for the static and moving stimuli, Fig 9. Moreover, we were able to analyze and identify, using the "effective reversal potential" of the large conductance state, the mechanisms by which the computational model produced the identical cortical activity for the two distinct stimuli – identifying the crucial roles played by the long range orientation preserving connections and the slow NMDA synapses. Thus, we were able to identify the possible mechanisms by which the mammalian cortex achieves the line-motion illusion.

4. Conclusion

In this essay, we have described our large-scale computational model of a layer of mammalian primary visual cortex, and used this model to illustrate manners in which large-scale models can contribute to our understanding of the operation of mammalian cortical systems. In the construction of the large-scale model, biology provides many realistic constraints on the system — constraints

that restrict one's freedom to realize within the model abstract theoretical constructs that are proposed to be the mechanisms by which the cortex achieves its responses. It is important to emphasize that proper modeling requires that one model, with one specification of parameters, be used for every experiment. One cannot use different models, with different parameter values, to achieve agreement with different experimental observations. Further, when a large-scale model actually achieves biologically realistic results, one can identify and analyze in detail the mechanisms by which the model achieves these responses — and thus identify the likely mechanisms by which the mammalian cortex functions.

The cortical operating point, or state of cortical activity, at which our model of a layer of mammalian primary visual cortex operates may be summarized as follows: high conductance; strong inhibition; neuronal firing that is driven by temporal fluctuations in voltage (since the strong inhibition keeps mean voltages below the firing threshold); high gain; very sparse connectivity (which enables operation at high gain without instabilities); ordered patterns of cortical activity (generated by connections, often long range connections, which are selective to feature maps and operate on the long time scales of NMDA receptors); sequential, but temporally irregular, initiation of patterns of cortical activity by small groups of excitatory neurons firing together over short time frames (firings which are rapidly quenched due to inhibition, but with cortical activity that is maintained over O(100ms) time scales through long-range connections with NMDA receptors). This is the cortical operating point of the

large-scale model cortex — an operating point that can be analyzed in detail for the computational model. Thus, it is strongly suggested to be the operating point of the biological cortex.

---

1 Anatomical staining studies show that the neurons reside in distinct layers, with the density of neurons and the density of distinct neuronal types varying from layer to layer, as well as neuronal projections to and from other cortical regions varying from layer to layer.

2 Visual information enters through the retina and is transmitted along the *visual pathway* from the retina, through the lateral geniculate nucleus (LGN), to the "entry layers" of V1, and on to other cortical regions (see Fig 1). Of course, this "feed forward" pathway is a part of a complex system of feedback pathways — within layers, between layers, and between cortical regions.

3 For this binary sequence, time is discretized into small bins or windows of a few ms in duration, $t_n = n\,\Delta$, $n = 0,1,2,.....$; the presence or absence of a spike within the nth bin is denoted by a "1 or 0", and the sequence is then a binary temporal sequence of 0s and 1s. These sequences are usually depicted by "raster plots" that depict the spike time of the neuron as a function of time.

4 Neurons come in many distinct types, but perhaps the most fundamental classification is between *excitatory* and *inhibitory* neurons, that is between neurons whose voltage spikes *excite* other neurons and those whose voltage spikes *inhibit* other neurons.

5 Tao, L., Shelley, M., McLaughlin, D. & Shapley, R., "An Egalitarian Network Model for the Emergence of Simple and Complex Cells in Visual Cortex", Proc Natl Acad USA 101, 366-371 (2004).

6 Tao, L., Cai, D., McLaughlin,D., Shelley, M. & Shapley, R., "Orientation Selectivity in Visual Cortex by Fluctuation-Controlled Criticality", Proc Natl Acad USA 103, 12911-12916 (2006).

7 Cai, D., Rangan, A. & McLaughlin, D., "Architectural and Synaptic Mechanisms Underlying Coherent Spontaneous Activity in V1", Proc Natl Acad USA 102 (16), 5868–5873 (2005).

8 Cai, D., Rangan, A. & McLaughlin, D., "Modeling Spatial-temporal Cortical Activity Associated with the Line Motion Illusion in Primary Visual Cortex", Proc Natl Acad Sci USA 102 (52), 18793–18800 (2005).

9 Hubel, D. & Wiesel, T., "Receptive Fields of Single Neurones in the Cat's Striate Cortex", J Physiol 148, 574-591 (1959).

10 Hubel, D. & Wiesel, T., "Receptive fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex", J Physiol. 160, 106-154 (1962).

11 This preferential coupling is indicated by the subscript "j,k" replacing the isotropic labeling "j-k"

12 Shelley, M., McLaughlin, D., Shapley, R. & Wielaard, J., "States of High Conductance in a Large-Scale Model of the Visual Cortex", J Comp Neuroscience 13, 93-109 (2002).

13 Hebb, D., “The Organization of Behavior: A Neuropsychological Theory”, New York: Wiley and Sons (1949).

14 McLaughlin, D., Shapley, R., Shelley, M. Wielaard, J., “A Neuronal Network Model of the Macaque Primary Visual Cortex V1: Orientation Tuning and Dynamics in the Input Layer 4Cα”, Proc. Natl. Acad. Sci. USA, 97, 8087-8092 (2000).

15 Chariker, L., Shapley, R. & Young, L-S., “ Orientation Selectivity for Very Sparse LGN Inputs in a Comprehensive Model of Macaque V1 Cortex”, J Neuro Sci 36(49), 12368-12384 (2016).

16 Wielaard, J., Shelley, M., Shapley, R. & McLaughlin, D., “How Simple Cells are Made in a Nonlinear Network Model of the Visual Cortex”, J. Neurosci. 21, 5203–5211 (2001).

17 Blasdel, G., “Differential Imaging of Ocular Dominance and Orientation Selectivity in Monkey Striate Cortex”, J. Neurosci 12, 3115–3138 (1992).

18 Blasdel, G., “Orientation Selectivity, Preference, and Continuity in Monkey Striate Cortex”, J. Neurosci 12, 3139–3161 (1992).

19 Sirovich, L. & Uglesich, R., “The Organization of Orientation and Spatial Frequency in Primary Visual Cortex”, Proc. Natl. Acad. Sci. (USA) 101 (48), 16941-16946 (2004).

20 These mechanisms are based on the cortical-cortical summation over nearby cells averaging over all orientations, and thus are much reduced in networks which are very sparsely connected and thus the cortical-cortical averaging is not over all angles of preference Ref [5,6]. For very sparse networks, near pinwheel centers, simple cells are

modestly tuned (and a little better tuned than neurons far from the pinwheel center), while complex cells are broadly tuned for inhibition, but with distributed values for the angle of orientation preference relative to the preference of the impinging LGN drive.

21 Niell, C. & Stryker, M., "Highly Selective Receptive Fields in Mouse Visual Cortex", J. Neurosci 28 (30), 7520–7536 (2008).

22 Niell, C. & Stryker, M., "Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex", Neuron 65, 472-479 (2010).

23 Lien, A. & Scanziani, M., "Tuned Thalamic Excitation is Amplified by Visual Cortical Circuits", Nature Neuroscience 16 (9), 1315–1323 (2013).

24 Li, Y., Ibrahim, L., Liu, B., Zhang, L. & Huizhong, W., "Linear Transformation of Thalamocortical Input by Intracortical Excitation", Nature Neuroscience 16 (9), 1324–1330 (2013).

25 Wielaard, J., McLaughlin, D. & Shapley, R., "A Model of Mouse V1 Layer 4 – Signatures of Specificity", unpublished (2016).

26 Arieli, A., Shoham, D., Hildesheim, R., & Grinvald, A., "Coherent Spatiotemporal Patterns of Ongoing Activity Revealed by Real-time Optical Imaging Coupled with Single-unit Recordings in the Cat Visual Cortex", J. Neurophysiol. 73, 2072-2093 (1995).

27 Shoham, D., Glaser, D., Arieli, A., Kinet, T., Wijnbergen, C., Toledo, Y., Hildesheim, R., & Grinvald, "Imaging Cortical Dynamics at High Spatial and Temporal Resolution with Novel Blue Voltage Sensitive Dyes", A., Neuron 24, 791–802 (1999).

28 Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A., “Linking Spontaneous Activity of Single Cortical Neurons and the Underlying Functional Architecture”, Science 286, 1943–1946 (1999).

29 Jancke, D., Chavance, F., Naaman, S., & Grinvald, A., “Imaging Cortical Correlates of Illusion in Early Visual Cortex”, Nature 428, 423–426 (2004).

30 Bosking, W., Zhang, Y., Schofield, B., & Fitzpatrick, D., “Orientation Selectivity and the Arrangement of Horizontal Connections in Tree Shrew Striate Cortex”,J. Neurosci. 17 (6), 2112–2127 (1997).

31 Hubel, D, “Eye, Brain, and Vision”, Scientific American Library Series (22), (1987). Reprinted with permission from Scientific American Library Series which holds the copyright (1988, 1995).

32 Yeh, C., Xing, D., Williams, P. & Shapley, R., “Stimulus Ensemble and Cortical Layer Determine V1 Spatial Receptive Fields”, Proc. Natl. Acad. USA 106 (34), 14652-14657 (2009).