A population density approach that facilitates large-scale modeling of neural networks: extension to slow inhibitory synapses

Duane Q. Nykamp and Daniel Tranchina

Abstract

A previously developed method for efficiently simulating complex networks of integrate-and-fire neurons was specialized to the case in which the neurons have fast unitary postsynaptic conductances. However, inhibitory synaptic conductances are often slower than excitatory for cortical neurons, and this difference can have a profound effect on network dynamics that cannot be captured with neurons having only fast synapses. We thus extend the model to include slow inhibitory synapses. In this model, neurons are grouped into large populations of similar neurons. For each population, we calculate the evolution of a probability density function (PDF) which describes the distribution of neurons over state space. The population firing rate is given by the flux of probability across the threshold voltage for firing an action potential. In the case of fast synaptic conductances, the PDF was one-dimensional, as the state of a neuron was completely determined by its transmembrane voltage. An exact extension to slow inhibitory synapses increases the dimension of the PDF to two or three, as the state of a neuron now includes the state of its inhibitory synaptic conductance. However, by assuming that the expected value of a neuron's inhibitory conductance is independent of its voltage, we derive a reduction to a one-dimensional PDF and avoid increasing the computational complexity of the problem. We demonstrate that although this assumption is not strictly valid, the results of the reduced model are surprisingly accurate.

1 Introduction

In a previous paper (Nykamp and Tranchina, 2000), we explored a computationally efficient population density method. This method was introduced as a tool for facilitating large-scale neural network simulations by Knight and colleagues (Knight et al., 1996; Omurtag et al., 2000; Sirovich et al., 2000; Knight, 2000). Earlier works that provide a foundation for our previous and current analysis include those of Knight (1972a,b), Kuramoto (1991), Strogatz and Mirollo (1991), Abbott and van Vreeswijk (1993), Treves (1993), and Chawanya et al. (1993).

The population density approach considered in the present paper is accurate when the network is comprised of large, sparsely connected subpopulations of identical neurons each receiving a large number of synaptic inputs. We showed (Nykamp and Tranchina, 2000) that the model produces good results, at least in some test networks, even when all these conditions that guarantee accuracy are not met. We demonstrated, moreover, that the population density method can be two hundred times faster than conventional direct simulations in which the activity of thousands of individual neurons and hundreds of thousands of synapses is followed in detail.

In the population density approach, integrate-and-fire point-neurons are grouped into large populations of similar neurons. For each population, one calculates the evolution of a probability density function (PDF) which represents the distribution of neurons over all possible states. The firing rate of each population is given by the flux of probability across a distinguished voltage which is the threshold voltage for firing an action potential. The populations are coupled via stochastic synapses in which the times of postsynaptic conductance events for each neuron are governed by a modulated Poisson process whose rate is determined by the firing rates of its presynaptic populations.

In our original model (Nykamp and Tranchina, 2000), we assumed that the time course of the postsynaptic conductance waveform of both excitatory and inhibitory synapses was fast on the time scale of the typical 10–20 ms membrane time constant (Tamás et al., 1997). In this approximation, the conductance events were instantaneous (delta functions) so that the state of each neuron was completely described by its transmembrane voltage. This assumption resulted in a one-dimensional PDF for each population. The approximation of instantaneous synaptic conductance seems justified for those excitatory synaptic conductance mediated by AMPA (alpha-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid) receptors that decay with a time constant of 2–4 ms (Stern et al., 1992).

We were motivated to embellish and extend the population density methods because the instantaneous synaptic conductance approximation is not justified for some inhibitory synapses. Most inhibition in the cortex is mediated by the neurotransmitter γ-aminobutyric acid (GABA) via $\mathrm{GABA_A}$ or $\mathrm{GABA_B}$ receptors. Some $\mathrm{GABA_A}$ conductances decay with a time constant that is on the same scale as the membrane time constant (Tamás et al., 1997; Galarreta and Hestrin, 1997), and $\mathrm{GABA_B}$ synapses are an order of magnitude slower. Furthermore, slow inhibitory synapses may be necessary to achieve proper network dynamics.

With slow inhibition, the state of each neuron is described not only by its transmembrane voltage, but also by one or more variables describing the state of its inhibitory synaptic conductance. Thus, a full population density implementation of neurons with slow inhibitory synapses would require two or more dimensions for each PDF. Doubling the dimension will square the size of the discretized state space and greatly decrease the speed of the population density computations. Instead of developing numerical techniques to solve this more difficult problem efficiently, we chose to develop a reduction of the expanded population density model back to one dimension that gives fast and accurate results.

In section 2 we outline the integrate-and-fire point-neuron model with fast excitation and slow inhibition that underlies our population density formulation. In section 3 we briefly derive the full population density equations that include slow inhibition. We reduce the population densities to one dimension and derive the resulting equations in section 4. We test the validity of the reduction in section 5, and demonstrate the dramatic difference in network behavior that can be introduced by the additional of slow inhibitory synapses. We discuss the results in section 6.

2 The integrate-and-fire neuron

As in our original model (Nykamp and Tranchina, 2000), the implementation of the population density approach is based on an integrate-and-fire, single-compartment neuron (point-neuron). We describe here the changes made to introduce slow inhibition.

The evolution of a neuron's transmembrane voltage V in time t is specified by the stochastic ordinary differential equation[1]:

$$c\frac{dV}{dt} + g_r(V - \mathcal{E}_r) + G_e(t)(V - \mathcal{E}_e) + G_i(t)(V - \mathcal{E}_i) = 0, \tag{1}$$

where c is the capacitance of the membrane, g_r, $G_e(t)$, and $G_i(t)$ are the resting, excitatory and inhibitory conductances, respectively, and $\mathcal{E}_{r/e/i}$ are the corresponding equilibrium potentials. When $V(t)$ reaches a fixed threshold voltage v_{th}, the neuron is said to fire a spike, and the voltage is reset to the voltage v_{reset}.

The synaptic conductances $G_e(t)$ and $G_i(t)$ vary in response to excitatory and inhibitory input, respectively. We denote by $A^k_{e/i}$ the integral of the change in $G_{e/i}(t)$ due to a unitary excitatory/inhibitory input at time $T^k_{e/i}$. Thus, if $\hat{G}^k_{e/i}(t)$ is the response of the excitatory/inhibitory synaptic conductance to a single input at time $T^k_{e/i}$, then $A^k_{e/i} = \int \hat{G}^k_{e/i}(t)dt$. The voltage $V(t)$ is a dynamic random variable because we view the $A^k_{e/i}$ and $T^k_{e/i}$ as random quantities. We denote the average value of $A_{e/i}$ by $\mu_{A_{e/i}}$.

2.1 Excitatory synaptic input

Since the change in $G_e(t)$ due to excitatory input is assumed to be very fast, it is approximated by a delta function, $G_e(t) = \sum_k A^k_e \, \delta(t - T^k_e)$. Upon receiving an excitatory input, the voltage jumps by the amount

$$\Delta V = \Gamma^{*k}_e \left[\mathcal{E}_e - V(T^{k-}_e)\right], \tag{2}$$

where[2] $\Gamma^{*k}_e = 1 - \exp(-A^k_e/c)$. Since the size of each conductance event A^k_e is random, the Γ^{*k}_e are random numbers. We define the complementary cumulative distribution function for Γ^*_e,

$$\tilde{F}_{\Gamma^*_e}(x) = \Pr(\Gamma^*_e > x), \tag{3}$$

which is some given function of x that depends on the chosen distribution of A^k_e.

[1]In stochastic ordinary differential equations and auxiliary equations, upper case Roman letters indicate random quantities.

[2]We use the notation Γ^*_e for consistency with notation from our earlier paper.

2.2 Inhibitory synaptic input

The inhibitory synaptic input is modeled by a slowly changing inhibitory conductance $G_i(t)$. Upon receiving an inhibitory synaptic input, the inhibitory conductance increases and then decays exponentially. The rise of the inhibitory conductance can be modeled as either an instantaneous jump (which we call first order inhibition) or a smooth increase (second order inhibition).

For first order inhibition, the evolution of $G_i(t)$ is governed by

$$\tau_i \frac{dG_i}{dt} = -G_i(t) + \sum_k A_i^k \ \delta(t - T_i^k), \tag{4}$$

where τ_i is the time constant for the decay of the inhibitory conductance. The response of $G_i(t)$ to a single inhibitory synaptic input at $T = 0$ is $G_i(t) = (A_i/\tau_i)\exp(-t/\tau_i)$ for $t > 0$.

For second order inhibition, we introduce an auxiliary variable $S(t)$ and model the evolution of $G_i(t)$ and $S(t)$ with the equations $\tau_i \frac{dG_i}{dt} = -G_i(t) + S(t)$ and $\tau_s \frac{dS}{dt} = -S(t) + \sum_k A_i^k \ \delta(t - T_i^k)$, where $\tau_{s/i}$ is the time constant for the rise/decay of the inhibitory conductance ($\tau_s \leq \tau_i$). The response of $G_i(t)$ to a single inhibitory synaptic input is either a difference of exponentials or, if $\tau_i = \tau_s$, an alpha function.

Since the A_i^k are random numbers, we define the complementary cumulative distribution function for A_i,

$$\tilde{F}_{A_i}(x) = \Pr(A_i > x), \tag{5}$$

which is some given function of x.

3 The full population density model

With slow inhibitory synapses, the state of a neuron is described not only by its voltage ($V(t)$), but also by one (first order inhibition) or two (second order inhibition) variables for the state of the inhibitory conductance ($G_i(t)$ and possibly $S(t)$). Thus, the state space is now at least two-dimensional. Figure 1a illustrates the time courses of $V(t)$ and $G_i(t)$ (with first order inhibition) driven by the arrival of random excitatory and inhibitory synaptic inputs. When an excitatory event occurs, V jumps abruptly; when an inhibitory event occurs, G_i jumps abruptly. Thus, a neuron executes a "random walk" in the $V - G_i$ plane, as illustrated in figure 1b.

For simplicity, we focus on the two-dimensional case of first order inhibition. The equations for second order inhibition are similar. We first derive equations for a single population of *uncoupled* neurons with specified input rates. We denote the excitatory and inhibitory input rates at time t by $\nu_e(t)$ and $\nu_i(t)$, respectively. For the non-interacting neurons, $\nu_e(t)$ and $\nu_i(t)$ are considered to be given functions of time. In section 4.4, we outline the equations for networks of interacting neurons.

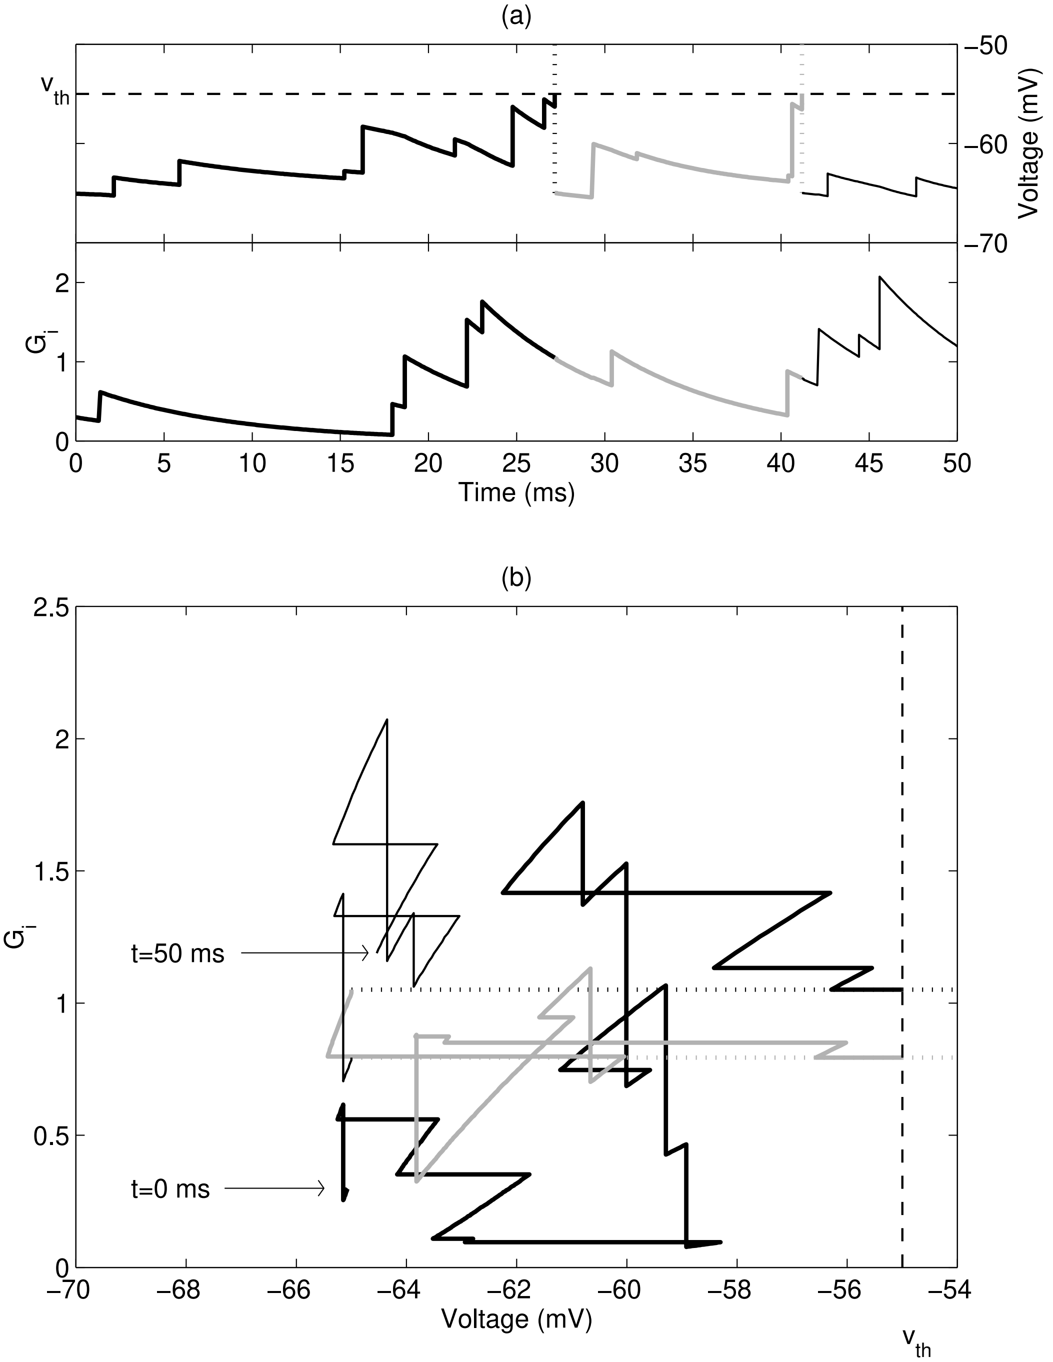

Figure 1: Illustration of the two-dimensional state space with first order inhibition. **(a)** A trace of the evolution of voltage and inhibitory conductance in response to constant input rate ($\nu_e = 400$ imp/sec, $\nu_i = 200$ imp/sec). For illustration, spikes are shown by dotted lines and the graylevel changes after each spike. **(b)** The same trace shown as a trajectory in the two-dimensional state space of voltage and inhibitory conductance.

3.1 The conservation equation

The two-dimensional state space for first order inhibition leads to the following two-dimensional PDF for a neuron

$$\rho(v, g, t)dv\, dg = \Pr\{V(t) \in (v, v + dv) \text{ and } G_i(t) \in (g, g + dg)\}, \tag{6}$$

for $v \in (\mathcal{E}_i, v_{th})$ and $g \in (0, \infty)$. For a population of neurons, the PDF can be interpreted as a population density; i.e.

$$\rho(v, g, t)dv\, dg = \text{Fraction of neurons with } V(t) \in (v, v + dv) \text{ and } G_i(t) \in (g, g + dg). \tag{7}$$

Each neuron has a refractory period of length τ_{ref} immediately after it fires a spike. During the refractory period, its voltage does not move, though its inhibitory conductance evolves as usual. When a neuron is refractory, it is not accounted for in $\rho(v, g, t)$. For this reason, $\rho(v, g, t)$ does not usually integrate to 1. Instead, $\int\int \rho(v, g, t)dv\, dg$ is the fraction of neurons that are not refractory. In appendix B, we describe the probability density function f_{ref} that describes neurons in the refractory state. However, since f_{ref} plays no role in the reduced model, we only mention the properties of f_{ref} as we need them for the derivation.

The evolution of the population density $\rho(v, g, t)$ is a described by a conservation equation, which takes into account the voltage reset after a neuron fires. Since a neuron reenters $\rho(v, g, t)$ after it becomes non-refractory, we view the voltage reset as occurring at the end of the refractory period. Thus, the evolution equation is:

$$\begin{aligned} \frac{\partial \rho}{\partial t}(v, g, t) &= -\nabla \cdot \vec{J}(v, g, t) + \delta(v - v_{reset})J_U(\tau_{ref}, g, t) \\ &= -\left(\frac{\partial J_V}{\partial v}(v, g, t) + \frac{\partial J_G}{\partial g}(v, g, t)\right) + \delta(v - v_{reset})J_U(\tau_{ref}, g, t). \end{aligned} \tag{8}$$

$\vec{J}(v, g, t) = (J_V(v, g, t), J_G(v, g, t))$ is the flux of probability at $(V(t), G_i(t)) = (v, g)$. The flux $\vec{J}$ is a vector quantity in which the J_V component gives the flux across voltage and the J_G, component gives the flux across conductance. Each component is described below. The last term in (8) is the source of probability at v_{reset} due to the reset of each neuron's voltage after it becomes non-refractory. The calculation of $J_U(\tau_{ref}, g, t)$, which is the flux of neurons becoming non-refractory, is described in appendix B.

Since a neuron fires a spike when its voltage crosses v_{th}, the population firing rate is the total flux across threshold:

$$r(t) = \int J_V(v_{th}, g, t)dg. \tag{9}$$

The population firing rate is not a temporal average, but is an average over all neurons in the population. When a neuron fires a spike, it becomes refractory for the duration of the refractory period of length τ_{ref}. Thus, the total flux of neurons becoming non-refractory at time t is equal to the firing rate at time $t - \tau_{ref}$:

$$\int J_U(\tau_{ref}, g, t)dg = r(t - \tau_{ref}). \tag{10}$$

The condition (10) is the only property of J_U that we will need for our reduced model.

3.2 The components of the flux

The flux across voltage J_V in (8) is similar to the flux across voltage from the fast synapse model (Nykamp and Tranchina, 2000). The total flux across voltage is:

$$J_V(v,g,t) = -\frac{g_r}{c}(v-\mathcal{E}_r)\rho(v,g,t) + \nu_e(t)\int_{\mathcal{E}_i}^{v} \tilde{F}_{\Gamma_e^*}\left(\frac{v-v'}{\mathcal{E}_e-v'}\right)\rho(v',g,t)dv' - \frac{g}{c}(v-\mathcal{E}_i)\rho(v,g,t) \tag{11}$$

where $\nu_e(t)$ is the rate of excitatory synaptic input at time t.

The three terms of the flux across voltage correspond to the terms on the right hand side of equation (1). The first two terms are the components of the flux due to leakage toward $\mathcal{E}_r$ and excitation toward $\mathcal{E}_e$, respectively, which are identical to the leakage and excitation fluxes in the fast synapse model. Because the slow inhibitory conductance causes the voltage to evolve smoothly toward $\mathcal{E}_i$ with velocity $-(V(t)-\mathcal{E}_i)G_i(t)/c$, the inhibition flux is the third term in equation (11).

The flux across conductance J_G consists of two components, corresponding to the two terms on the right hand side of equation (4):

$$J_G(v,g,t) = -\frac{g}{\tau_i}\rho(v,g,t) + \nu_i(t)\int_0^g \tilde{F}_{A_i}(\tau_i(g-g'))\rho(v,g',t)dg', \tag{12}$$

where $\nu_i(t)$ is the rate of inhibitory synaptic input at time t. The first term is the component of flux due to the decay of the inhibitory conductance toward 0 at velocity $-G_i(t)/\tau_i$. The second term is the flux from the jumps in the inhibitory conductance when the neuron receives inhibitory input. The derivation of the second term is analogous to the derivation of the flux across voltage due to excitation.

4 The reduced population density model

For the sake of computational efficiency, we developed a reduction of the model with slow synapses back to one dimension. This reduced model can be solved as quickly as the model with fast synapses.

To derive the reduced model, we first show that the jumps in voltage due to excitation, and thus the jumps across the firing threshold v_{th}, are independent of the inhibitory conductance. We therefore only need to calculate the distribution of neurons across voltage (and not inhibitory conductance) in order to accurately compute the population firing rate $r(t)$. We then show that the evolution of the distribution across voltage depends on the inhibitory conductance only via its first moment. The evolution of this moment can be computed by a simple ordinary differential equation if we assume it is independent of voltage.

4.1 EPSP amplitude independent of inhibitory conductance

The fact that the size of the excitatory voltage jump is independent of the inhibitory conductance is implicit in equation (2). Insight may be provided by the following argument. To calculate the response in voltage to a delta function excitatory input at time T, we integrate (1) from the time just before the input T^- to the time just after the input T^+. The terms involving g_r and $G_i(t)$ drop out because they are finite and thus cannot contribute to the integral over the infinitesimal interval (T^-, T^+). The term with $G_e(t)$ remains because across this interval $G_e(t)$ is a delta function.

Physically, the independence from $G_i(t)$ corresponds to excitatory current so fast, that it can only be balanced by capacitative current. Although the peak of the EPSP does not depend on the inhibitory conductance, the time course of the EPSP does. A large inhibitory conductance would, for example, create a shunt which would quickly bring the voltage back down after an excitatory input.

4.2 Evolution of distribution in voltage

Since excitatory jump sizes are independent of inhibitory conductance, we can calculate the population firing rate at time t if we know only the distribution of neurons across voltage. We denote the marginal probability density function in voltage by $f_V(v,t)$, which is defined by

$$f_V(v,t)dv = \Pr\{V(t) \in (v, v+dv)\} \tag{13}$$

and is related to $\rho(v,g,t)$ by [3]

$$f_V(v,t) = \int \rho(v,g,t)dg. \tag{14}$$

We derive an evolution equation for f_V by integrating (8) with respect to g, obtaining:

$$\frac{\partial f_V}{\partial t}(v,t) = -\left(\int \frac{\partial J_V}{\partial v}(v,g,t)dg + \int \frac{\partial J_G}{\partial g}(v,g,t)dg\right) + \delta(v - v_{reset})\int J_U(\tau_{ref}, g, t)dg. \tag{15}$$

Because $G_i(t)$ cannot cross 0 or infinity, the flux at those points is zero, and consequently the second term on the right hand side is zero:

$$\int \frac{\partial J_G}{\partial g}(v,g,t)dg = \lim_{g\to\infty} J_G(v,g,t) - J_G(v,0,t) = 0 \tag{16}$$

Using condition (10), we have the following conservation equation for f_V:

$$\frac{\partial f_V}{\partial t}(v,t) = -\frac{\partial \bar{J}_V}{\partial v}(v,t) + \delta(v - v_{reset})r(t - \tau_{ref}), \tag{17}$$

where $\bar{J}_V(v,t)$ is the total flux across voltage:

$$\begin{aligned}\bar{J}_V(v,t) &= \int J_V(v,g,t)dg \qquad (18)\\ &= -\frac{g_r}{c}(v - \mathcal{E}_r)f_V(v,t) + \nu_e(t)\int_{\mathcal{E}_i}^{v} \tilde{F}_{\Gamma_e^*}\left(\frac{v - v'}{\mathcal{E}_e - v'}\right) f_V(v',t)dv' - \frac{1}{c}(v - \mathcal{E}_i)\int g\rho(v,g,t)dg\end{aligned}$$

[3]Note that this enables us to rewrite the expression (9) for the population firing rate as $r(t) = \nu_e(t)\int_{\mathcal{E}_i}^{v_{th}} \tilde{F}_{\Gamma_e^*}\left(\frac{v_{th}-v'}{\mathcal{E}_e - v'}\right) f_V(v',t)dv'$ (since $\rho(v_{th},g,t) = \rho(\mathcal{E}_i,g,t) = 0$), confirming that the firing rate depends only on the distribution in voltage.

Equation (17) would be a closed equation for $f_V(v,t)$ except for its dependence on $\int g\,\rho(v,g,t)dg$, the first moment of G_i. Thus, to calculate the population firing rates, we don't need the distribution of neurons across inhibitory conductance; we only need the average value of that conductance for neurons at each voltage. This can be seen more clearly by rewriting the last term of (18) as follows.

Let $f_{G|V}(g,v,t)$ be the probability density function for G_i given V:

$$f_{G|V}(g,v,t)dg = \Pr\{G_i(t) \in (g, g+dg) | V(t) = v\}. \tag{19}$$

Then, since $\rho(v,g,t) = f_{G|V}(g,v,t)f_V(v,t)$, we have that the first moment of G_i is

$$\begin{aligned} \int g\,\rho(v,g,t)dg &= \int g f_{G|V}(g,v,t)dg\, f_V(v,t) \\ &= \mu_{G|V}(v,t)f_V(v,t), \end{aligned} \tag{20}$$

where $\mu_{G|V}(v,t)$ is the expected value of G_i given V:

$$\mu_{G|V}(v,t) = \int g\, f_{G|V}(g,v,t)dg. \tag{21}$$

Substituting (20) into (18), we have the following expression for the total flux across voltage

$$\bar{J}_V(v,t) = -\frac{g_r}{c}(v-\mathcal{E}_r)f_V(v,t) + \nu_e(t)\int_{\mathcal{E}_i}^{v} \tilde{F}_{\Gamma_e^*}\left(\frac{v-v'}{\mathcal{E}_e - v'}\right) f_V(v',t)dv' - \frac{\mu_{G|V}(v,t)}{c}(v-\mathcal{E}_i)f_V(v,t), \tag{22}$$

which, combined with (17), gives the equation for the evolution of the distribution in voltage:

$$\begin{aligned} \frac{\partial f_V}{\partial t}(v,t) &= \frac{\partial}{\partial v}\left[\frac{g_r}{c}(v-\mathcal{E}_r)f_V(v,t) - \nu_e(t)\int_{\mathcal{E}_i}^{v} \tilde{F}_{\Gamma_e^*}\left(\frac{v-v'}{\mathcal{E}_e - v'}\right) f_V(v',t)dv' \right. \\ &\quad \left. + \frac{\mu_{G|V}(v,t)}{c}(v-\mathcal{E}_i)f_V(v,t)\right] \\ &\quad + \delta(v - v_{reset})r(t-\tau_{ref}). \end{aligned} \tag{23}$$

4.3 The independent mean assumption

Equation (23) depends on the mean conductance $\mu_{G|V}(v,t)$. In appendix A, we derive an equation for the evolution of the mean conductance $\mu_{G|V}(v,t)$ (equation (34) coupled with (38) and (39)). The equation cannot be solved directly because it depends on $\mu_{G^2|V}$, the second moment of G_i. The equation for the evolution of the second moment would depend on the third, and so on for higher moments, forming a moment expansion in G_i.

To close the moment expansion, one has to make an assumption that will make the expansion terminate. We assume that $\mu_{G|V}(v,t)$ is independent of voltage, i.e. that the mean conditioned on the voltage is equal to the unconditional mean:

$$\mu_{G|V}(v,t) = \mu_G(t) \tag{24}$$

where $\mu_G(t)$ is the mean value of $G_i(t)$ averaged over all neurons in the population.

Since the equation for the evolution of G_i (4) is independent of voltage, this assumption closes the moment expansion at the first moment and enables us to derive, as shown in appendix A, a simple ordinary differential equation for the evolution of the mean voltage:

$$\frac{d\mu_G}{dt}(t) = \frac{\nu_i(t)\mu_{A_i} - \mu_G(t)}{\tau_i}. \tag{25}$$

Therefore, if we make our independence assumption (24), we can combine (25) with (23) to calculate the evolution of a single population of *uncoupled* neurons to excitatory and inhibitory input at the *given* rates $\nu_e(t)$ and $\nu_i(t)$.

4.4 Coupled populations

All the population density derivations above were for a single population of uncoupled neurons. However, we wish to apply the population density method to simulate large networks of coupled neurons. The implementation of a network of population densities requires two additional steps. A detailed description and analysis of the assumptions behind these steps has been given earlier (Nykamp and Tranchina, 2000).

The first step is to group neurons into populations of similar neurons and form a population density for each group, denoted $f_V^k(v,t)$, where $k = 1, 2, \ldots, N$ enumerates the N populations. Each population has a corresponding population firing rate, denoted $r^k(t)$. Denote the set of excitatory indices by Λ_E and the set of inhibitory indices by Λ_I, i.e., the set of excitatory/inhibitory populations is $\{f_V^k(v,t) \mid k \in \Lambda_{E/I}\}$. Denote any imposed external excitatory/inhibitory input rates to population k by $\nu_{e/i,o}^k(t)$.

The second step is to determine the connectivity between the populations and form the connectivity matrix W_{jk}, $j,k = 1, 2, \ldots, N$. W_{jk} is the average number of presynaptic neurons from population j that project to each postsynaptic neuron in population k. The synaptic input rates to each population are then determined by the firing rates of its presynaptic population as well as external firing rates by:

$$\nu_{e/i}^k(t) = \nu_{e/i,o}^k(t) + \sum_{j \in \Lambda_{E/I}} W_{jk} \int_0^\infty \alpha_{jk}(t') r^j(t-t')dt' \tag{26}$$

where $\alpha_{jk}(t')$ is the distribution of latencies of synapses from population j to population k. The choice of $\alpha_{jk}(t')$ used in our simulations is given in appendix D. Using the input rates determined by (26), one can calculate the evolution of each population using the equations for the single population.

4.5 Summary of equations

Combining (17), (22), (24), and (25), we have a partial differential-integral equation coupled with an ordinary differential equation for the evolution of a single population density with synaptic input rates $\nu_e(t)$ and $\nu_i(t)$. In a network with populations $f_V^k(v,t)$, these input rates are given by (26). The firing rates of each population are given by (9).

The following summarizes the equations of the population density approach with populations $k = 1, 2, \ldots, N$.

$$\frac{\partial f_V^k}{\partial t}(v,t) = -\frac{\partial \bar{J}_V^k}{\partial v}(v,t) + \delta(v - v_{reset}) r^k(t - \tau_{ref}) \tag{27}$$

$$\bar{J}_V^k(v,t) = -\frac{g_r}{c}(v - \mathcal{E}_r) f_V^k(v,t) + \nu_e^k(t) \int_{\mathcal{E}_i}^{v} \tilde{F}_{\Gamma_e^*}\left(\frac{v - v'}{\mathcal{E}_e - v'}\right) f_V^k(v',t)dv' - \frac{\mu_G^k(t)}{c}(v - \mathcal{E}_i) f_V^k(v,t) \tag{28}$$

$$r^k(t) = \bar{J}_V^k(v_{th}, t) \tag{29}$$

$$\frac{d\mu_G^k}{dt}(t) = \frac{\nu_i^k(t)\mu_{A_i} - \mu_V^k(t)}{\tau_i} \tag{30}$$

$$\nu_{e/i}^k(t) = \nu_{e/i,o}^k(t) + \sum_{j \in \Lambda_{E/I}} W_{jk} \int_0^\infty \alpha_{jk}(t') r^j(t - t')dt' \tag{31}$$

The boundary conditions for the partial differential-integral equations are that $f_V^k(\mathcal{E}_i, t) = f_V^k(v_{th}, t) = 0$. In general, the parameters $\mathcal{E}_{e/i/r}$, $v_{th/reset}$, g_r, c, and μ_{A_i} as well as the function $\tilde{F}_{\Gamma_e^*}$, could depend on k.

5 Tests of validity

The reduced population density model would give the same results as the full model if the expected value of the inhibitory conductance were independent of the voltage. We expect that the reduced model should give good results if the independence assumption (24) is approximately satisfied.

We performed a large battery of tests by which to gauge the accuracy of the reduced model. For each test, we compared the firing rates of the reduced population density model with the firing rates from direct simulations of groups of integrate-and-fire neurons, computed as described in appendix C. We modeled the inhibitory input accurately for the direct simulations, using the equations from section 2. Thus, the direct simulation firing rates serve as a benchmark against which we can compare the reduced model.

5.1 Testing procedure

We focus first on simulations of a single population of uncoupled neurons. For a single population of uncoupled neurons that receive specified input that is a modulated Poisson process, the population density approach without the independent mean assumption (24) would give exactly the same results as a large population of individual neurons. Thus, tests of single population results will demonstrate what errors the independent mean assumption alone introduces.

We also performed extensive tests using a simple network of one excitatory and one inhibitory population. In the network simulations, the comparison with individual neurons also tests other assumptions of the population density, such as the assumption that the combined synaptic input to each neuron can be approximated by a modulated Poisson process. We have demonstrated earlier that these assumptions are typically satisfied for sparse networks (Nykamp and Tranchina, 2000). The network tests with the reduced population density model will demonstrate the accumulated errors from all the approximations.

For both single population and two population tests, we performed many runs with random sets of parameters. We chose parameters in the physiological regime, but the parameters were chosen independently of one another. Thus, rather than focusing on parameter sets that may be relevant to a particular physiological system, we decided to stress the model with many combinations of parameters to probe the conditions under which the reduced model is valid. However, we analyzed the results of only those tests where the average firing rates were between 5 spikes/second and 300 spikes/second. Average spike rates greater than 300 spikes/second are unphysiological, and spike rates less than 5 spikes/second are too low to permit a meaningful comparison by our deviation measure, below.

We randomly varied the following parameters: the time constant for inhibitory synapses τ_i, the average unitary excitatory input magnitude μ_{A_e}, the average peak unitary inhibitory synaptic conductance μ_{A_i}/τ_i, and the connectivity of the network W_{jk}. We selected τ_i randomly from the range $\tau_i \in (2, 25)$ ms. Scaling the excitatory input by the membrane capacitance to make it dimensionless, we used the range: $\mu_{A_e}/c \in (0.001, 0.030)$. Similarly, scaling the peak unitary inhibitory conductance by the resting conductance, we used the range $\mu_{A_i}/(\tau_i g_r) \in (0.001, 0.200)$. The single population simulations represented uncoupled neurons, so we set the connectivity to zero ($W_{11} = 0$). For the two population simulations, we set all connections to the same strength ($W_{jk} = \bar{W},\ j, k = 1, 2$), and selected that strength randomly from $\bar{W} \in (5, 50)$. The other parameters were fixed for all simulations. These fixed parameters as well as the forms of the probability distributions for $A_{e/i}$ are described in appendix D.

For the purpose of illustration, we also ran some single population simulations with a larger peak unitary inhibitory conductance. We allowed $\mu_{A_i}/(\tau_i g_r)$ to range as large as 1, which is not physiological since a single inhibitory event then doubles the membrane conductance from its resting value.

To provide inputs rich in temporal structure, we constructed each of the external input rates $\nu^1_{e,o}(t)$ and $\nu^1_{i,o}(t)$ from sums of sinusoids with frequencies of 0, 1, 2, 4, 8, and 16 cycles/second. The phases and relative amplitudes of each component of the sum were chosen randomly. The steady components, which were the mean rates $\bar{\nu}^1_{e,o}$ and $\bar{\nu}^1_{i,o}$, were independently chosen from the range $\bar{\nu}^1_{e/i,o} \in (0, 2000)$ spikes/second. The amplitudes of the other 5 components were scaled by the largest common factor that would keep $\nu^1_{e/i,o}(t)$ in the range $(0, 2\bar{\nu}^1_{e/i,o})$ for all time. For the two population simulations, the second population received no external input and we set $\nu^2_{e,o}(t) = \nu^2_{i,o}(t) = 0$.

We performed similar tests using second order inhibition with larger inhibitory time constants ($\tau_i \in (50, 100)$ ms, $\tau_s \in (2, 5)$ ms) and a lower inhibitory reversal potential ($\mathcal{E}_\rangle =$

-90 mV rather than -70 mV), in order to approximate GABA$_\mathrm{B}$ synapses. Since the results were similar to those with first order inhibition, we do not focus on them here.

We also performed one test of a physiological network. We implemented a slow inhibitory synapse version of the model of one hypercolumn of visual cortex we used in an earlier paper (Nykamp and Tranchina, 2000). This model simulates the response of neurons in the primary visual cortex of a cat to an oriented flashed bar. We kept every aspect of the model unchanged except the inhibitory synapses. The original model had instantaneous inhibitory synapses. In this implementation we used first order inhibition with $\tau_i = 8$ ms and kept the integral of the unitary inhibitory conductance the same as that in the previous model.

5.2 Test results

For each run, we calculated the firing rates in response to the input using both the population density method and a direct simulation of 1000 individual integrate-and-fire point-neurons. We calculated the firing rates of the individual neurons by counting spikes in 5 ms bins, denoting the firing rate for the j^{th} bin as $\tilde{r}_j$. To compare the two firing rates, we averaged the population density firing rate over each bin, denoting the average rate for the j^{th} bin by r_j. We then calculated the following deviation measure:

$$\Delta = \frac{\sqrt{\sum_j (r_j - \tilde{r}_j)^2}}{\sqrt{\sum_j r_j^2}}, \tag{32}$$

where the sum is over all bins. The deviation measure Δ is a quantitative measure of the difference between the firing rates, but it has shortcomings, as demonstrated below. For two populations, we report the maximum Δ of the populations.

5.2.1 Single uncoupled population

We performed 10,000 single population runs and found that the reduced population density method was surprisingly accurate for all the runs. The average Δ was only 0.05 and the maximum Δ was 0.19. Figure 2a shows the results of a typical run. For this run with a Δ near average ($\Delta = 0.06$), the firing rates of the population density and direct simulations are virtually indistinguishable.

Since the population density results closely matched the direct simulation results for all 10,000 runs, we simulated a set of 500 additional runs which included unrealistically large unitary inhibitory conductance events ($\mu_{A_i}/(\tau_i g_r)$ as large as 1). In some cases, the large inhibitory conductance events led to larger deviations. The results from the run with the highest deviation measure ($\Delta = 0.43$) are shown in figure 2b. In this example, the population density firing rates are consistently lower than the direct simulation firing rates.

The derivation of the reduced population density model assumed that the mean value of the inhibitory conductance $\mu_{G|V}(v,t)$ was independent of the voltage v (24). A necessary (though not sufficient) condition for (24) is the absence of correlation between the inhibitory

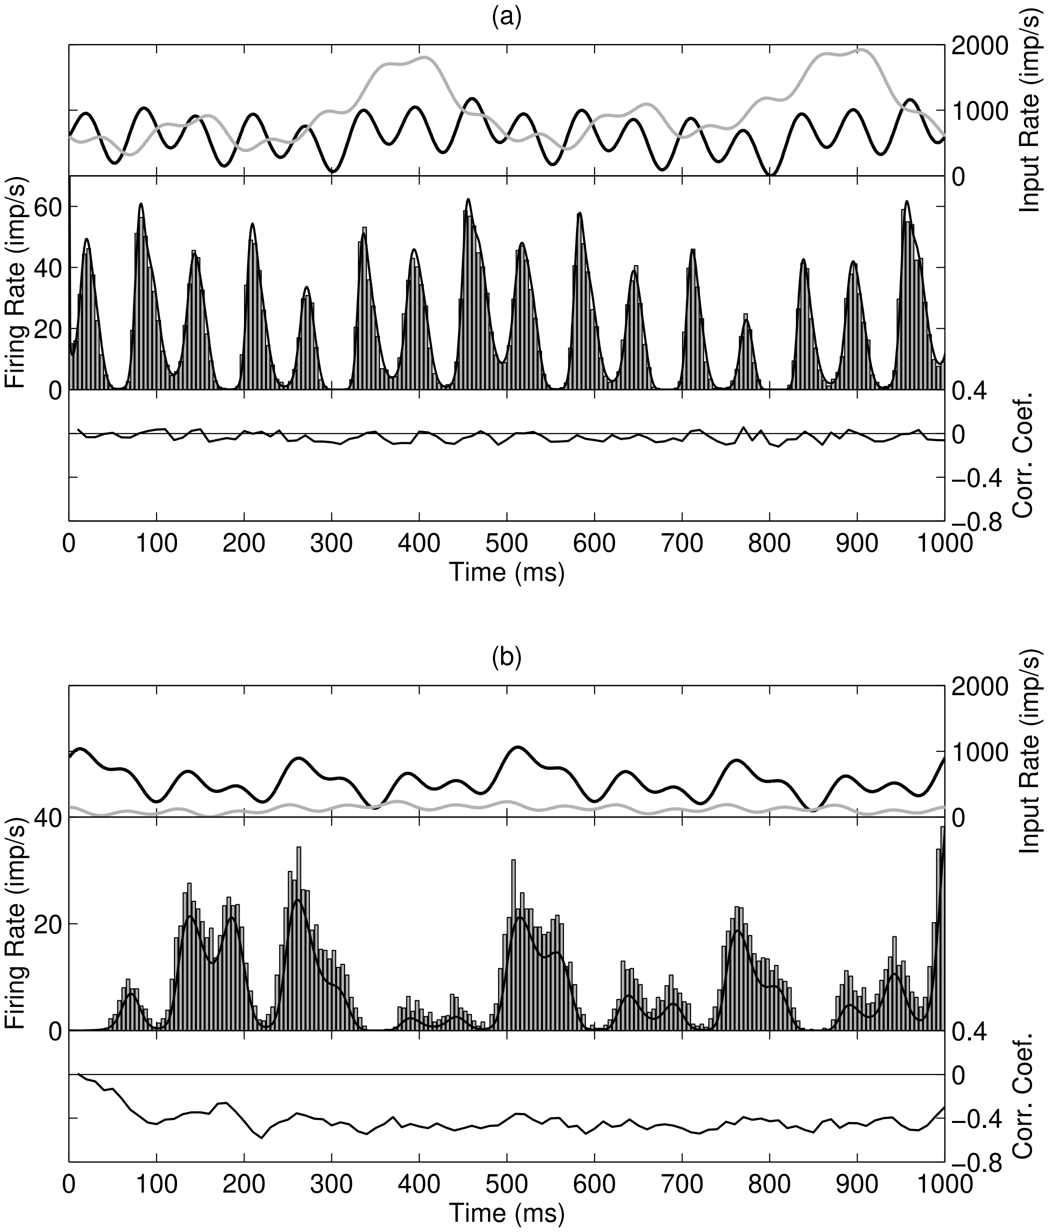

Figure 2: Example single uncoupled population results. Excitatory (black line) and inhibitory (gray line) input rates are plotted in the top panel. Population density firing rate (black line) and direct simulation firing rate (histogram) are plotted in the middle panel. The correlation coefficient between the voltage and the inhibitory conductance is shown in the bottom panel. **(a)** The results from a typical simulation. The firing rates from both simulations are nearly identical, and the correlation coefficient is small. Parameters: $\tau_i = 6.5$ ms, $\mu_{A_e}/c = 0.015$, $\mu_{A_i}/(\tau_i g_r) = 0.029$, $\bar{\nu}_e = 619$ imp/s, $\bar{\nu}_i = 961$ imp/s. **(b)** The results with the largest deviation when we allowed the unitary inhibitory conductance to be unrealistically large. The population density firing rate consistently undershoots that of the direct simulation, and the correlation coefficient is large and negative. Parameters: $\tau_i = 22.2$ ms, $\mu_{A_e}/c = 0.030$, $\mu_{A_i}/(\tau_i g_r) = 0.870$, $\bar{\nu}_e = 531$ imp/s, $\bar{\nu}_i = 118$ imp/s.

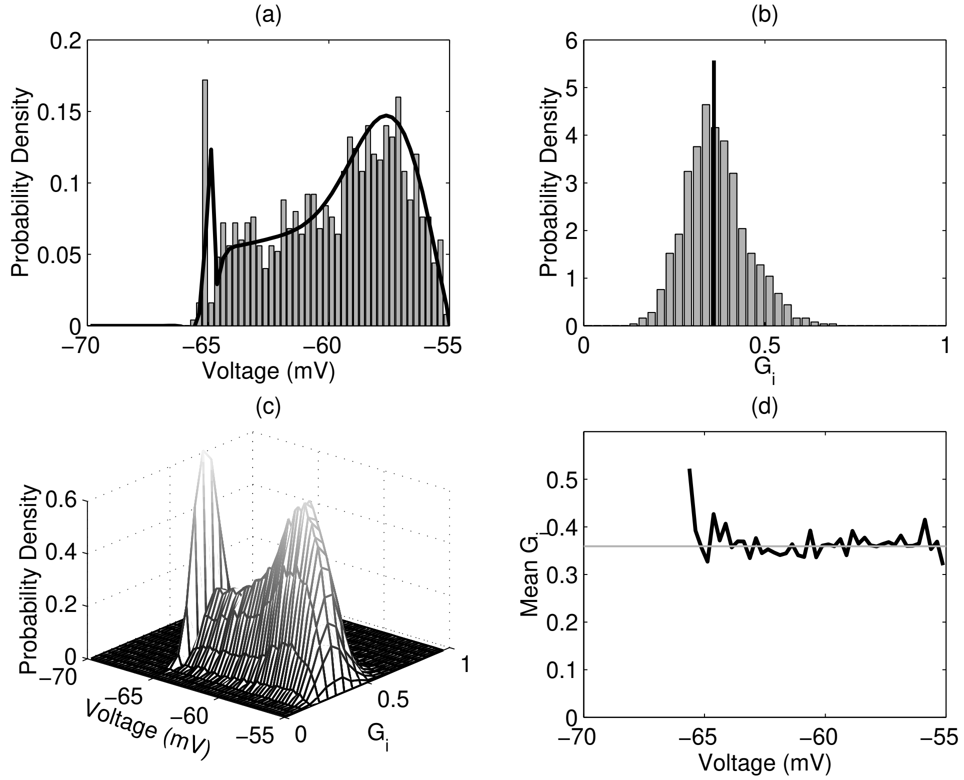

Figure 3: Snapshots from $t = 900$ ms of figure 2a. **(a)** Distribution of neurons across voltage as computed by population density (black line) and direct simulation (histogram). **(b)** Distribution of neurons across inhibitory conductance from direct simulation (histogram). Vertical black line is at the average inhibitory conductance from the population density simulation. **(c)** Distribution of neurons across voltage and inhibitory conductance from the direct simulation averaged over 100 repetitions of the input. **(d)** Conditional mean inhibitory conductance $\mu_{G|V}(v,t)$ from direct simulation (black line). Horizontal gray line is plotted at the mean inhibitory conductance $\mu_G(t)$ from the population density simulation.

conductance and the voltage. Thus the correlation coefficient calculated from the direct simulation might be a good predictor of the reduced population density's accuracy.

The correlation coefficient, which can range between -1 and 1, is plotted underneath each firing rate in figure 2. For the average results of figure 2a, the correlation coefficient is very small and usually negative. For the larger deviation in figure 2b, the inhibitory conductance is strongly negatively correlated with the voltage throughout the run. Thus, at least in this case, the correlation coefficient does explain the difference in accuracy between the two runs.

Not only do the population density simulations give accurate firing rates, they also capture the distribution of neurons across voltage (f_V) and the mean inhibitory conductance (μ_G), as shown by the snapshots in figure 3. Figure 3a compares the distribution of neurons across voltage ($f_V(v,t)$) as calculated by the population density and direct simulations. Al-

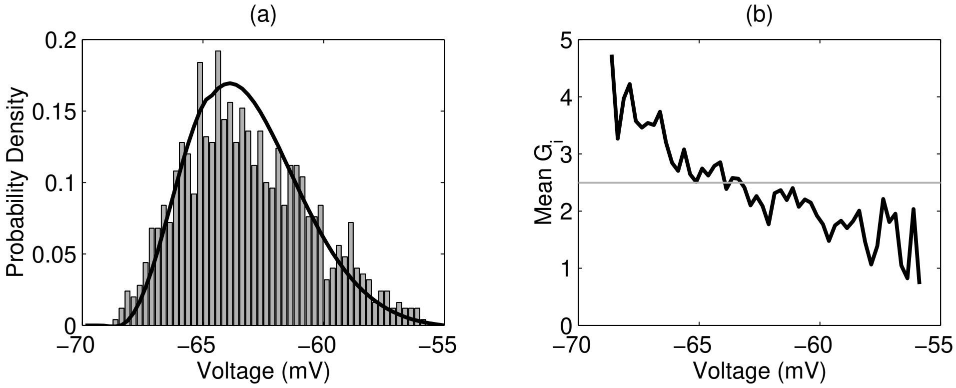

Figure 4: Snapshots from $t = 900$ ms of figure 2b. **(a)** Distribution of neurons across voltage as computed by population density (black line) and direct simulation (histogram). **(b)** Conditional mean inhibitory conductance $\mu_{G|V}(v,t)$ from direct simulation (black line). Horizontal gray line is plotted at the mean inhibitory conductance $\mu_G(t)$ from the population density simulation.

though the direct simulation histogram is still ragged with 1,000 neurons, it is very close to the $f_V(v,t)$ calculated by the population density model. Figure 3b demonstrates that the population density can calculate an accurate mean inhibitory conductance despite the wide distribution of neurons across inhibitory conductance.

For illustration, figure 3c shows the full distribution of neurons across voltage and inhibitory conductance. In order to produce a relatively smooth graph, this distribution was estimated by averaging the direct simulation over 100 repetitions of the input. The surprising result of our reduced method is that we can collapse the full distribution in figure 3c to the distribution in voltage of figure 3a coupled with the mean inhibitory conductance from figure 3b.

To analyze our critical independence assumption (24), we plot in figure 3d an estimate of the conditional mean $\mu_{G|V}$ computed by the direct simulation along with the unconditional mean value μ_G from the population density simulation. For this example, the estimate of $\mu_{G|V}$ is nearly independent of voltage, which explains why the population density results so closely match those from the direct simulation.

The snapshots in figure 4, which are from $t = 900$ ms in figure 2b, further demonstrate the breakdown of the independence assumption that can occur with unrealistically large unitary inhibitory conductances. Figure 4b shows that the mean inhibitory conductance $\mu_{G|V}(v,t)$ is not independent of voltage but decreases as the voltage increases. This dependence leads to the discrepancy in the distribution of neurons across voltage seen in figure 4a. The match between the direct simulation and the population density results might be improved if, rather than making an independence assumption (24) at the first moment of G_i, we retained the second moment from the moment expansion described at the beginning of section 4.3. Further analysis of the sources of disagreements between direct and population density computations is given in the discussion.

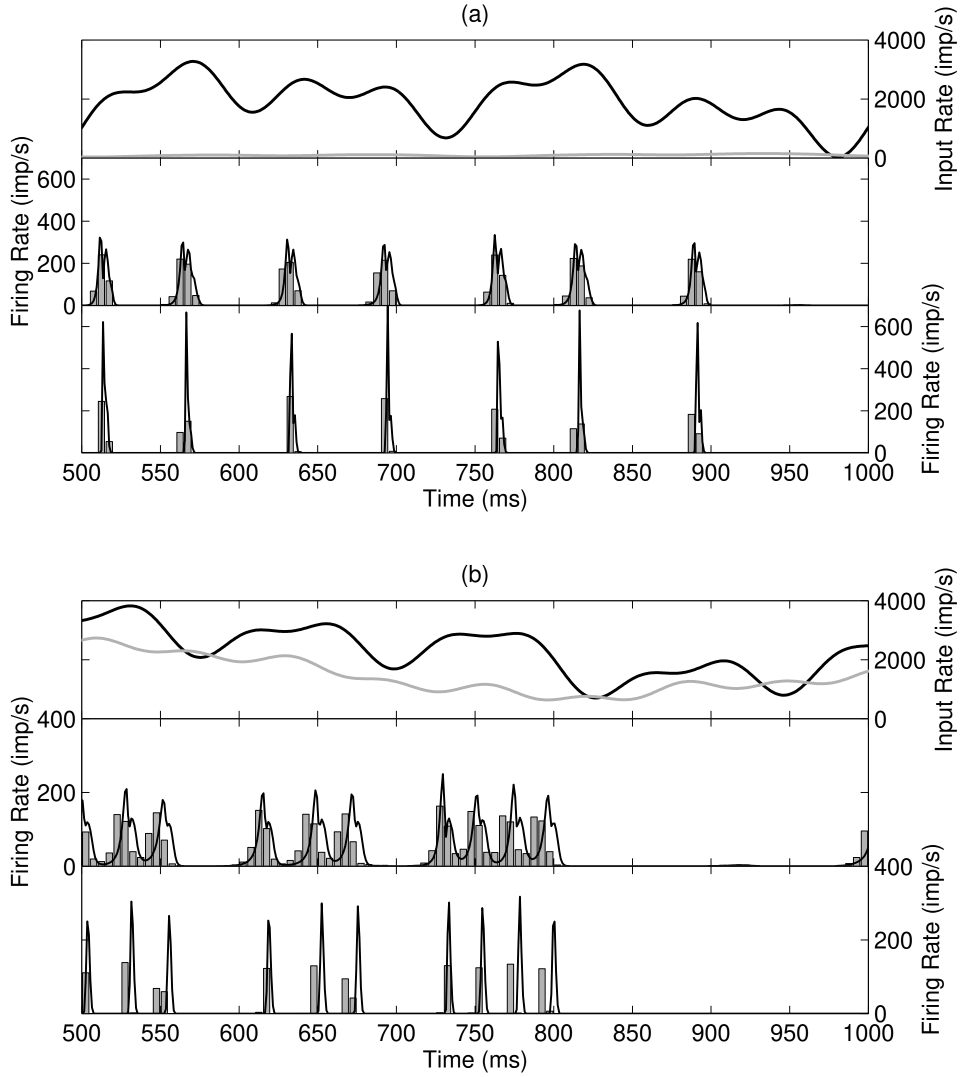

Figure 5: Examples of two population results. **(a)** The firing rates match well despite a large deviation measure ($\Delta = 0.64$). Top panel: Excitatory (black line) and inhibitory (gray line) input rates. Middle and bottom panel: Excitatory and inhibitory population firing rates, respectively. Population density (black line) and direct simulation (histogram) results are shown. Parameters: $\bar{W} = 48.05$, $\tau_i = 22.25$ ms, $\mu_{A_e}/c = 0.011$, $\mu_{A_i}/(\tau_i g_r) = 0.196$, $\bar{\nu}_e = 1853$ imp/s, $\bar{\nu}_i = 82$ imp/s. **(b)** Slight timing differences caused the large deviation ($\Delta = 0.98$). Panels as in (a). Parameters: $\bar{W} = 47.89$, $\tau_i = 2.49$ ms, $\mu_{A_e}/c = 0.005$, $\mu_{A_i}/(\tau_i g_r) = 0.137$, $\bar{\nu}_e = 1913$ imp/s, $\bar{\nu}_i = 1367$ imp/s.

5.2.2 Two population network

We simulated 300 runs of the two population network. Most (245/300) of the runs had relatively low deviation measure of Δ below 0.30. A substantial fraction (55/300) had deviations larger than 0.30. However, as demonstrated by examples below, we discovered numerous cases where Δ was high but the firing rates of the population density and direct simulations were subjectively similar. This discrepancy is largely due to the sensitivity of Δ to timing; slight differences in timing frequently lead to a large Δ. Rather than invent another deviation measure, we analyzed by eye each of the 55 simulations with $\Delta > 0.30$ to determine the source of the large Δ. We discovered that for the majority (31/55) of the simulations, the two firing rates appeared subjectively similar, as demonstrated by examples below. For only 24 of the 300 simulations, did the firing rates of the population density differ substantially from the direct simulation firing rate. However, we demonstrate evidence that these simulations represent situations in which the network has several different stable firing patterns and that neither the direct simulation nor the population density simulation can be viewed as the standard or correct result.

Figure 5a demonstrates a case in which there is good qualitative agreement between direct and population density computations despite a large deviation measure ($\Delta = 0.64$). Note that since the population firing rate is a population average, not a temporal average, the sharp peaks above 600 impulses/second in the inhibitory population firing rate do not indicate that a single neuron was firing above 600 impulses/second. Rather, these sharp peaks indicate that many neurons in the population fired a spike within a short time interval, a phenomenon we call population synchrony. For this example, the large Δ was caused by slight differences in timing of the tight population synchrony. Thus, not only does this example illustrate deficiencies in the measure Δ, it also demonstrates how well the population density model can capture the population synchrony of a strongly coupled network.

In other cases, a large Δ indicated more substantial discrepancies between the direct and population density simulations. Figure 5b shows the results of a simulation with $\Delta = 0.98$. The deviation measure is so high because the timing of the two models differs enough so that the peaks in activity sometimes do not overlap. Even so, the timings typically differ by only about 5 ms, and the firing rates are subjectively similar. In figure 6a, the population density underestimates the firing rates three times ($\Delta = 0.66$) but catches all qualitative features.

In 24 simulations, we observed periods of gross qualitative differences between the population density firing rates and the direct simulation firing rates. We discovered that in each of these cases, the network dynamics were extremely sensitive to parameters; rounding parameters to three significant digits led to large differences in firing pattern for both population density and direct simulations. In some cases, this rounding caused the firing patterns to converge and produced a low deviation measure.

An example of a simulation with a large deviation measure ($\Delta = 1.09$) is shown in figure 6b. Between $t = 250$ and 400 ms, the network appears to have a stable fixed point with both populations firing steadily at high rates. The direct simulation approaches this putative fixed point more rapidly than the population density simulation. At around $t = 400$ ms, the response of both simulations is consistent with a fixed point losing stability, as the firing

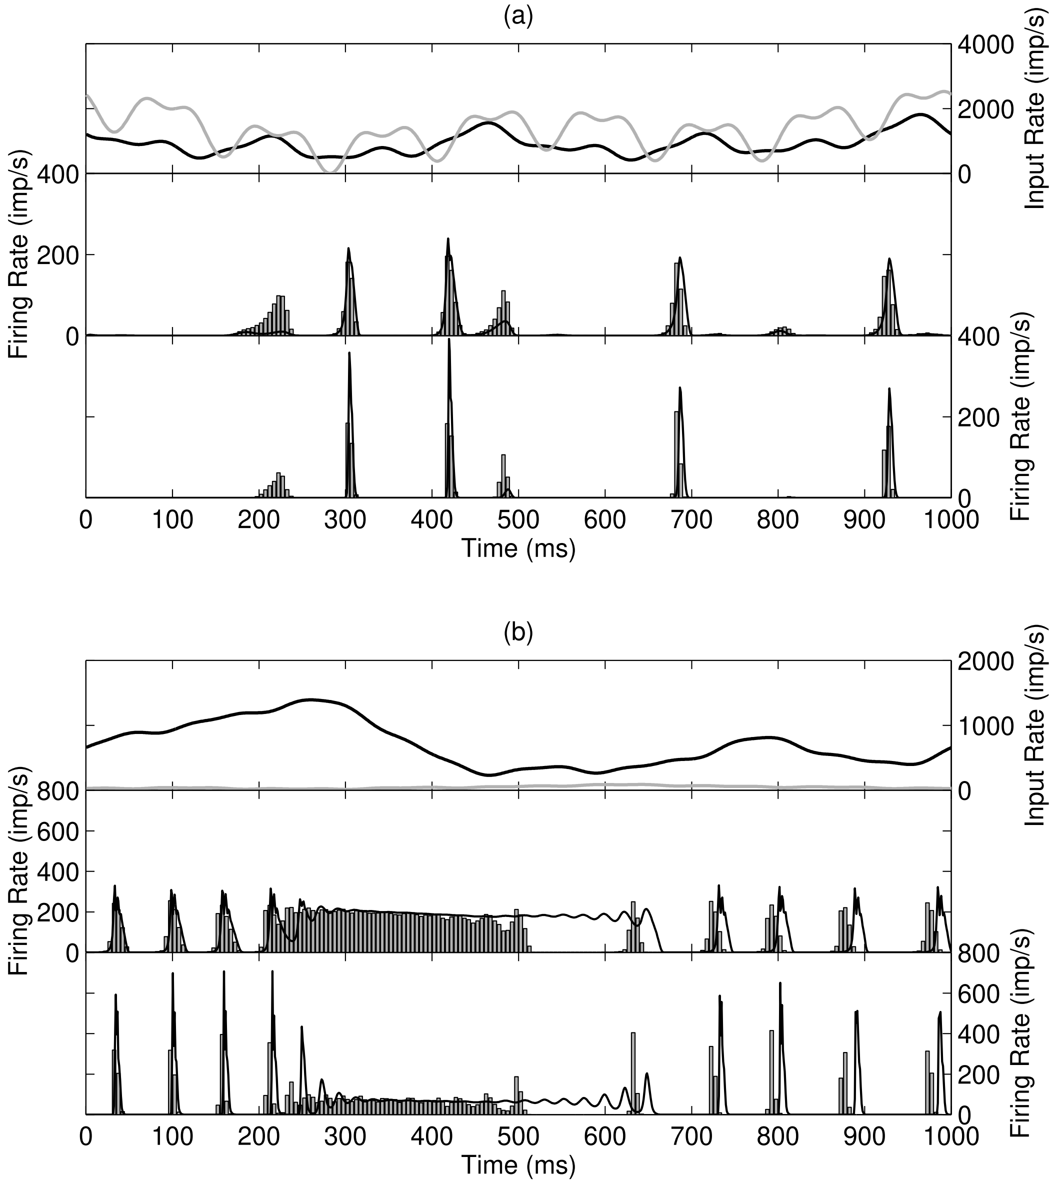

Figure 6: Further examples of two population results. Panels as in figure 5. **(a)** A simulation where the population density underestimates the firing rate. Parameters: $\bar{W} = 25.54$, $\tau_i = 11.66$ ms, $\mu_{A_e}/c = 0.018$, $\mu_{A_i}/(\tau_i g_r) = 0.172$, $\bar{\nu}_e = 908$ imp/s, $\bar{\nu}_i = 1391$ imp/s. **(b)** A simulation with large deviations between the population density and direct simulation firing rates. Parameters: $\bar{W} = 40.942$, $\tau_i = 19.992$ ms, $\mu_{A_e}/c = 0.018$, $\mu_{A_i}/(\tau_i g_r) = 0.144$, $\bar{\nu}_e = 696$ imp/s, $\bar{\nu}_i = 44$ imp/s.

rates oscillate around the fixed point with increasing amplitude. However, the oscillations of the direct simulation grow more rapidly than those of the population density simulation, and the populations in the direct simulation stop firing 150 ms before those in the population density simulation. From this point on, the population density and direct simulations behave similarly. But since the input rates for this example change slowly, the firing rates do not phase lock to the input rates, and the two simulations oscillate out of phase.

The behavior after $t = 400$ ms is reminiscent of drifting through a Hopf bifurcation and tracking the ghost of an unstable fixed point (Baer et al., 1989). Once the fixed point loses stability, the state of the system evolves away from the unstable fixed point only slowly. Any additional noise will cause the system to leave the fixed point faster (Baer et al., 1989). Thus, the additional noise in the direct simulation (e.g. from heterogeneity in the connectivity or finite size effects) might explain why the direct simulation leaves the fixed point more quickly.

Furthermore, the stability of this fixed point is extremely sensitive to parameters. Even when we leave all parameters fixed and generate a different realization of the direct simulation network, the point at which the direct simulation drops off the fixed point can differ by 100 ms. Rounding parameters to 3 significant digits changes the duration of the fixed point by hundreds of milliseconds. Moreover, when we increased $\mu_{A_i}/(\tau_i g_r)$ from 0.143 to 0.16, we no longer observed the high fixed point, and the firing rates of the two simulations matched with only small delays.

Each of the 24 simulations with high deviation measure exhibited similarly high sensitivity to parameters. Most of them appeared to have fixed points with high firing rates, giving rise to apparent bistability in the network activity; typically, these fixed points gained or lost stability differently for the two models. These observations suggest the possibility that the large deviations above reflect differences between Monte Carlo simulations and solutions of the corresponding partial differential-integral equations under bistable conditions. In these situations, neither simulation can be considered the standard; the deviation measure Δ is not a measure of the error in the population density simulation but a reflection of the sensitivity of the bistable conditions to noise. Thus, these discrepancies probably do not result from our dimension reduction approximation.

5.2.3 Model of one hypercolumn of visual cortex

In our model of one hypercolumn of visual cortex, based on that of Somers et al. (1995), excitatory and inhibitory populations are labeled by their preferred orientations. These preferred orientations are established by the input from the LGN, as demonstrated in figure 7a. The LGN synaptic input rates ($\nu^j_{e,o}(t)$) to excitatory populations in response to a flashed bar oriented at 0° are shown in the top panel. The population with a preferred orientation at 0° received the strongest input, and the population at 90° received the weakest input. Inhibitory populations received similar input. The resulting excitatory population firing rates ($r^j(t)$) are shown in the bottom panel. More details on the structure of the network are given in appendix D and in our earlier paper (Nykamp and Tranchina, 2000).

The addition of slow inhibitory synapses changes the dynamics of the network response.

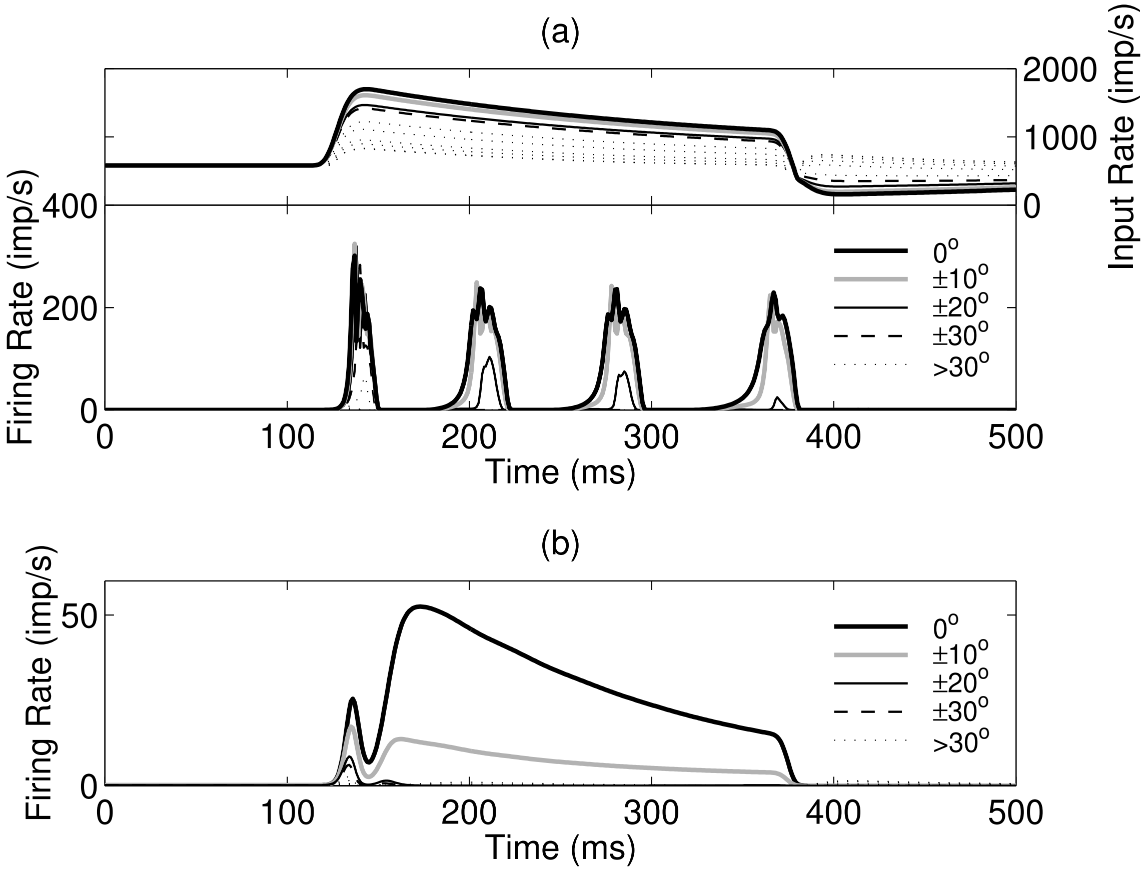

Figure 7: Results of visual cortex simulation. **(a)** Population density results with slow inhibition. Top panel: Synaptic input rates from the LGN to excitatory populations. These input rates are produced by a bar flashed from $t = 100$ ms to $t = 350$ ms, oriented at 0°; thus the input to the population at 0° is the largest. Bottom panel: The firing rates of the excitatory population densities in response to this input. **(b)** Population density results with fast inhibition. The firing rates of the excitatory population densities to the same input as in (a). The only difference between (b) and the bottom panel in (a) is that $\tau_i = 8$ ms for (a) while τ_i is essentially 0 for (b).

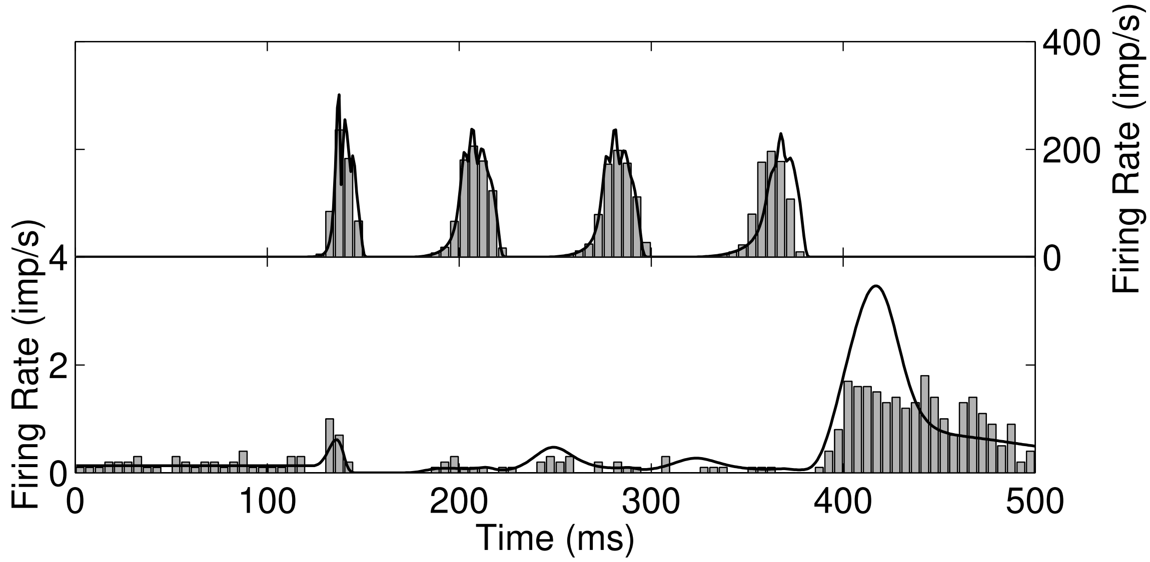

Figure 8: A comparison between population density and direct simulation firing rates for our visual cortex simulation. Population density firing rates are plotted with a black line and direct simulation firing rates are shown by the histograms. The population density firing rates are from figure 7a. **(a)** Excitatory population with a preferred orientation of 0°. **(b)** Excitatory population with a preferred orientation of 90°. Note the change in scale.

For figure 7a, we used slow inhibition with $\tau_i = 8$ ms. In figure 7b, we show the response of the same network to the same input where we have substituted fast inhibitory synapses, keeping μ_{A_i} (the average integral of a unitary inhibitory conductance) fixed. After an initial transient of activity, the firing rate for the fast inhibition model is smooth, indicating asynchronous firing. In contrast, the firing rate for the slow inhibition model (figure 7a) exhibits non-periodic cycles of population activity. Thus, this example demonstrates the importance of slow inhibitory synapses to network dynamics.

In figure 8, we plot a comparison of the population density firing rates with the firing rates of a direct simulation of the same network. As above, the direct simulation contained 1000 integrate-and-fire point-neurons per population. We averaged over two passes of the stimulus to obtain sufficient spikes for populations with low firing rates. Figure 8 compares the firing rates of excitatory populations with preferred orientations of 0° (top) and 90° (bottom). The firing rates match well. The match between the populations at 0° is excellent, though there is a slight timing difference in the fourth wave of activity. The binning of the direct simulation spikes cannot show the fine temporal structure visible in the population density firing rate. The only obvious deviation is the larger rebound in the population density than in the direct simulation at 90° after $t = 400$ ms. However, the difference is only a couple impulses per second, as the population is barely firing.

This example also serves to illustrate the speed of the population density approach. For both the population density and the direct simulations above, we used a time step of $\Delta t = 0.5$

ms. For the population density computations, we discretized the voltage state space using $\Delta v = 0.25$ mV. These simulations of a real time period of 500 ms with 36 populations took 22 seconds for the population density and 2600 seconds for the direct simulations (using a Silicon Graphics Octane computer with 1 195 MHz MIPS R10000 processor). Thus, the population density simulation was over 100 times faster than the direct simulation.

6 Discussion

Application of population density techniques to modeling networks of realistic neurons will require extensions of the theory beyond simple integrate-and-fire neurons with instantaneous postsynaptic conductances. We have taken a first step in this direction by adding a realistic time course for the inhibitory postsynaptic conductance.

The major difficulty of extending the population density approach to more complicated neurons arises from the additional dimension of the probability density function (PDF) required for each new variable describing the state of a neuron. Each additional dimension increases the computer time required to simulate the populations. Thus, dimension reduction procedures like the one employed in this paper become important in developing efficient computational methods for simulating population densities.

6.1 Analysis of dimension reduction

The dimension reduction procedure we used for this paper was inspired in part by the population density with mean field synapses by Treves (1993). Treves developed a model by approximating the value of each neuron's synaptic conductance by the mean value over the population. Although we derive our equations by making a looser assumption (24), the resulting equations are equivalent to a mean field approximation in the inhibitory conductance; the equations depend only on the mean inhibitory conductance μ_G.

6.1.1 Anticipated deficiency of reduction

Because the membrane potential of a neuron can be driven below the rest potential $\mathcal{E}_r$ only by a large inhibitory conductance, the expected value of a neuron's inhibitory conductance must depend on its voltage, contrary to our independence assumption (24). Moreover, since a larger inhibitory conductance would decrease a neuron's voltage, one would anticipate that the mean inhibitory conductance would be negatively correlated with the voltage, as was sometimes seen when we set the unitary inhibitory peak conductance to unphysiologically large values (figure 4).

The inability of our population density model to represent these correlations leads to its systematically too low firing rates when those correlations are important (e.g., figure 2b). Since the population density model cannot account for the dependence of $\mu_{G|V}(v,t)$ on voltage, it overestimates the mean inhibitory conductance for high voltages and underestimates it for low voltages (figure 4b) in the presence of large negative correlations. These voltage-dependent discrepancies in the inhibitory conductance increase the fraction of neurons at intermediate voltages in the population density model because neurons with high/low voltage receive artificially high/low inhibition. The fraction of neurons at both high and low voltages is thus lower than in the direct simulation, resulting in the differences in the distribution of neurons seen in figure 4a. Since fewer neurons are near threshold and ready to fire, the firing rate for the population density model is lower than the firing rate of the direct simulation.

6.1.2 Analysis of results

The surprising discovery of much better accuracy than might be anticipated must be explained by decorrelating actions in the dynamics of the neurons. By analyzing the 500 single population runs that included large inhibition, we obtained insight into both the cause of the deterioration of the independence assumption (24) and the decorrelating actions that help preserve the independence.

A summary of the 500 runs is plotted in figure 9. We see a confirmation in figure 9a that unphysiologically large peak unitary inhibitory conductances are required for a large deviation measure Δ. The range of peak conductances used for the first 10,000 runs is shaded, the upper end of which is already larger than what is measured experimentally (Tamás et al., 1997; Galarreta and Hestrin, 1997). The dependence of the error on the correlation between inhibitory conductance and voltage shown in figure 9b suggests that, at least for the single uncoupled population, the error in the population density results is indeed a result of a violation of the independence assumption (24).

Since many runs with large peak unitary inhibitory conductance have small Δ (figure 9a), peak conductance size alone is not a good prediction of the accuracy of the population density method. We searched for other quantities that were more strongly related to Δ and would help us better understand the source of the error. We discovered that Δ is strongly related to the average firing rate (figure 9c) and less reliably related to the relative strength of inhibition over excitation (figure 9d).

Both the voltage reset after firing a spike and the voltage jumps due to excitatory input are decorrelating actions that help preserve the independence assumption (24). The voltage reset after a spike is particularly effective because a neuron that likely has a low inhibitory conductance is moved to a low voltage, thus descreasing the negative correlation between the voltage and the inihbitory conductance. As a result, a high firing rate should reduce the correlation and thus the deviation measure, creating the relationship between firing rate and the deviation measure seen in figure 9c.

The voltage jumps due to excitatory synaptic input also reduce the correlation between voltage and inhibitory conductance, as the jumps are similar to diffusion with a diffusion coefficient that contains the factor $\nu_e(t)\mu^2_{A_e}$ (Nykamp and Tranchina, 2000). The inhibitory effects that introduce correlation are proportional to $\mu_G(t)$, whose quasi-steady-state value is $\nu_i(t)\mu_{A_i}$. We see a moderately strong relationship (figure 9d) between the deviation measure

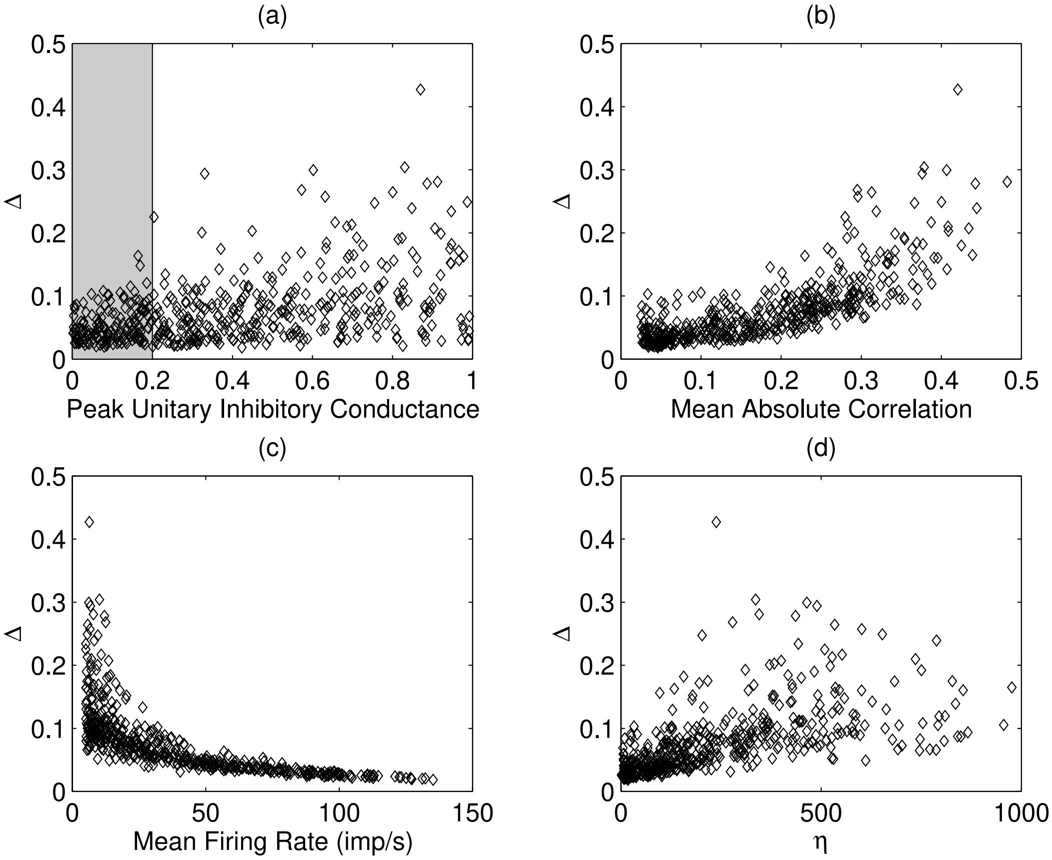

Figure 9: A summary of the 500 single population runs that included large inhibition. **(a)** A scatter plot of deviation measure (Δ) vs. peak unitary inhibitory conductance. The shaded region marks the range of more realistic conductance sizes used in the first 10,000 single population simulations. **(b)** Error measure vs. the average (over time and neurons) of the absolute value of the correlation coefficient between the voltage and inhibitory conductance. **(c)** Error measure vs. the average firing rate. **(d)** Error measure vs. η as defined by (33). For clarity, a point with $\eta = 1650$ and $\Delta = 0.1$ was omitted.

and the ratio of these quantities

$$\eta = \frac{\bar{\nu}_i \mu_{A_i} \tau_m}{\bar{\nu}_e \mu_{A_e}^2}, \tag{33}$$

where we have multiplied by $\tau_m = c/g_r$ to make the ratio dimensionless. For the vast majority of simulations, Δ increases with η, though for $\eta > 200$, there are a few simulations with higher than expected deviation measures.

The dynamics of these two decorrelating actions can be seen in individual simulations, as shown in figure 10. In figure 10a, the correlation coefficient mimics the features of the firing rate. Throughout the stimulus in figure 10b, including a period when the population is not firing, the correlation coefficient reflects the excitatory input rate. When the excitatory input rate is nearly zero around $t = 650$ ms, the correlation coefficient is largest in magnitude.

These two decorrelating actions are crucial for the accuracy of our population density model. Presumably, the reason they work so well is that they help the model most during periods when the neurons are firing. Thus, for example, the periods of highest correlation in figure 10 occur harmlessly when the population is virtually silent.

We performed the same analysis for the two population runs. However, as expected, the results were not as clean since Δ is not a satisfactory deviation measure when it is large. Moreover, for the 24 simulations with large subjective deviations, we would not expect to observe a strong link between the decorrelating actions and the deviations since these deviations are likely not a result of our independence assumption (24). Nonetheless, similar relationships as in the single population case were observed (not shown). We also observed that the probability of a large deviation increases as the coupling strength $\bar{W}$ increases. This observation is consistent with our hypothesis that network bistability may underlie some of the observed large deviations, as strong coupling facilitates network bistability.

6.2 Non-instantaneous synapses and network dynamics

Slow inhibition may be important for proper network dynamics. For example, based on results from rate models such as the Wilson and Cowan equations (Wilson and Cowan, 1972), one expects that the relation $\tau_i > \tau_e$, where τ_e is the time constant of excitatory synapses, may be important for the emergence of oscillatory behavior. Clearly, this relation cannot be retained when inhibition is approximated as instantaneous, and thus models with fast inhibition may miss oscillatory behavior that is present with slow inhibition.

Oscillatory behavior was exactly what emerged in our visual cortex simulations with the slow inhibition. The non-periodic oscillations of activity observed with $\tau_i = 8$ ms (figure 7a) were not present with fast inhibition (figure 7b). Since the excitation was fast (effectively $\tau_e = 0$ ms), the relation $\tau_i > \tau_e$ would remain true even for relatively small τ_i. Indeed, we observed oscillations even with τ_i as low as 2 ms (not shown). The oscillations began to disappear with $\tau_i = 1$ ms, and the firing rate approached the fast inhibition results (figure 7b) as τ_i was made even smaller. Similar results were obtained with second order inhibition, so these results were not peculiar to first order inhibition.

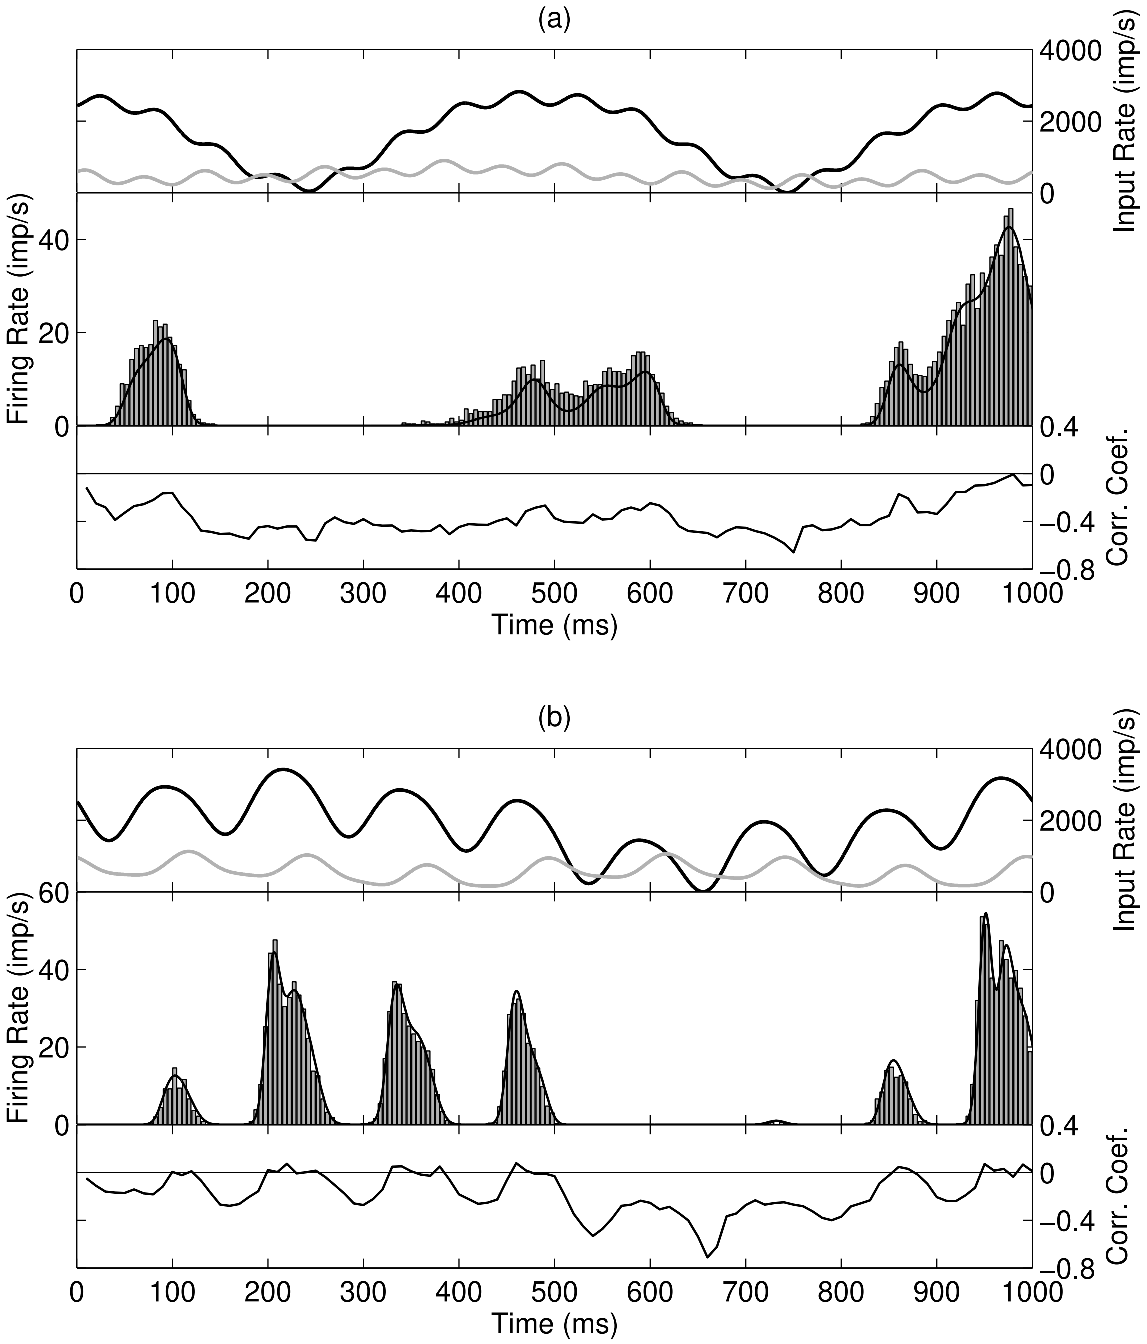

Figure 10: Further examples of results with a single population. Panels as in figure 2. **(a)** An example where the correlation coefficient is seen to mimic the firing rate. Parameters: $\tau_i = 20.60$ ms, $\mu_{A_e}/c = 0.010$, $\mu_{A_i}/(\tau_i g_r) = 0.193$, $\bar{\nu}_e = 1603$ imp/s, $\bar{\nu}_i = 449$ imp/s. **(b)** An example where the correlation coefficient closely follows the excitatory input rate $\nu_e(t)$ (black line in top panel). This is an example simulation with second order inhibition. Parameters: $\tau_i = 79.25$ ms, $\tau_s = 4.03$ ms, $\mu_{A_e}/c = 0.005$, peak unitary inhibitory conductance$/g_r = 0.003$, $\bar{\nu}_e = 1843$ imp/s, $\bar{\nu}_i = 563$ imp/s.

In networks where synapses are mainly excitatory, the time course of excitatory synapses may play a role in network dynamics such as synchronization (Abbott and van Vreeswijk, 1993; Gerstner, 2000). Furthermore, slower NMDA (N-methyl-D-aspartate) excitatory synapses cannot be accurately modeled as fast synapses. Thus, extending the model to include slow excitation would be a natural next step in the development of the population density approach.

6.3 Computational speed

We have demonstrated that the population density approach can be roughly 100 times faster than direct simulations. The speed advantage of the population density approach for slow inhibition is even greater than the advantage for fast synapses reported earlier (Nykamp and Tranchina, 2000). The population density simulations can be sped up even further using a diffusion approximation (Nykamp and Tranchina, 2000).

The computational speed of the population density simulations with slow inhibition depends on the reduction of the two- or three-dimensional PDF to one dimension. Further embellishments on the individual neuron model, such as additional compartments, ionic channels, or synaptic features, would typically increase the dimension of the PDF describing the state space of the neuron. Using dimension reduction procedures such as the one employed here or the principal component analysis method of Knight and colleagues (Knight, 2000), one may be able to reduce the dimension of these PDFs, allowing speedy computation of the resulting population density simulations. Thus, the population density approach may be a useful tool for efficiently simulating large networks of not only simple integrate-and-fire neurons but also more complicated, realistic neurons.

Acknowledgments

We gratefully acknowledge numerous helpful discussions with David McLaughlin, Robert Shapley, Charles Peskin and John Rinzel. This material is based on work supported under a National Science Foundation Graduate Fellowship.

Appendices

A Derivation of mean conductance evolution

In this appendix, we complete the derivation of equations for the reduced model in section 4 by deriving the evolution equation for the conditional mean conductance $\mu_{G|V}(v,t)$ and unconditional mean conductance $\mu_G(t)$.

A.1 The evolution of the conditional mean conductance

We derive an equation for the evolution of the conditional mean conductance $\mu_{G|V}(v,t)$ by multiplying (8) by g and integrating with respect to g. Using (20), we obtain:

$$\begin{aligned}\frac{\partial}{\partial t}[\mu_{G|V}(v,t)f_V(v,t)] &= -\frac{\partial}{\partial v}\int gJ_V(v,g,t)dg + \int J_G(v,g,t)dg \\ &\quad + \delta(v-v_{reset})\int gJ_U(\tau_{ref},g,t)dg. \end{aligned} \tag{34}$$

We used integration by parts to obtain the second term on the right hand side. The boundary terms were zero because $J_G(v,g,t)=0$ for[4] $g=0,\infty$.

Using (12) and (21), the integral of J_G becomes:

$$\int J_G(v,g,t)dg = -\frac{1}{\tau_i}\mu_{G|V}(v,t)f_V(v,t) + \nu_i(t)\int dg\int_0^g dg'\tilde{F}_{A_i}(\tau_i(g-g'))\rho(v,g',t). \tag{35}$$

We change the order of integration in the integral from the last term of (35):

$$\begin{aligned}\int_0^\infty dg\int_0^g dg'\tilde{F}_{A_i}(\tau_i(g-g'))\rho(v,g',t) &= \int_0^\infty\left[\int_{g'}^\infty \tilde{F}_{A_i}(\tau_i(g-g'))dg\right]\rho(v,g',t)dg' \\ &= \left[\int_0^\infty \frac{1}{\tau_i}\tilde{F}_{A_i}(x)dx\right]\left[\int_0^\infty \rho(v,g',t)dg'\right] \end{aligned} \tag{36}$$

where $x=\tau_i(g-g')$. Since $\tilde{F}_{A_i}(x)=\int_x^\infty f_{A_i}(y)dy$, where f_{A_i} is the probability density function for A_i, we have

$$\begin{aligned}\int_0^\infty \tilde{F}_{A_i}(x)dx &= \int_0^\infty\left[\int_x^\infty f_{A_i}(y)dy\right]dx \\ &= \int_0^\infty\left[\int_0^y dx\right]f_{A_i}(y)dy \\ &= \int_0^\infty y f_{A_i}(y)dy \\ &= \mu_{A_i} \end{aligned} \tag{37}$$

where μ_{A_i} is the average value of A_i. Combining equations (35) – (37), the second term from the right had side of (34) becomes:

$$\int J_G(v,g,t)dg = \frac{\nu_i(t)\mu_{A_i} - \mu_{G|V}(v,t)}{\tau_i}f_V(v,t). \tag{38}$$

Using (11), the integral from first term on the right hand side of (34) can be written:

$$\begin{aligned}\int gJ_V(v,g,t)dg &= -\frac{g_r}{c}(v-\mathcal{E}_r)\mu_{G|V}(v,t)f_V(v,t) \\ &\quad + \nu_e(t)\int_{\mathcal{E}_i}^v \tilde{F}_{\Gamma_e^*}\left(\frac{v-v'}{\mathcal{E}_e-v'}\right)\mu_{G|V}(v',t)f_V(v',t)dv' \\ &\quad - \frac{1}{c}(v-\mathcal{E}_i)\mu_{G^2|V}(v,t)f_V(v,t). \end{aligned} \tag{39}$$

[4] We implicitly assume J_G approaches zero faster than $1/g$ for large g.

In (39), $\mu_{G^2|V}(v,t)$ is defined similarly to $\mu_{G|V}(v,t)$, and we used the identity

$$\int g^2 \rho(v,g,t)dg = \mu_{G^2|V}(v,t) f_V(v,t). \tag{40}$$

In a similar manner, equations for $\mu_{G^k|V} f_V(v,t)$ can be derived for all k, in which $\mu_{G^k|V} f_V(v,t)$ depends on $\mu_{G^{k+1}|V} f_V(v,t)$.

A.2 The evolution of the unconditional mean conductance

With the independent mean assumption (24), we simply need to derive an equation for the unconditional mean $\mu_G(t)$. Since the evolution of $G_i(t)$ is independent of the voltage (see equation (4)), we can immediately write down the probability density function of G_i

$$f_G(g,t)dg = \Pr\{G_i(t) \in (g, g+dg)\} \tag{41}$$

and its evolution equation

$$\frac{\partial f_G}{\partial t}(g,t) = -\frac{\partial \bar{J}_G}{\partial g}(g,t). \tag{42}$$

The total flux across conductance $\bar{J}_G$ is identical in form to J_G (12):

$$\bar{J}_G(g,t) = -\frac{g}{\tau_i} f_G(g,t) + \nu_i(t) \int_0^g \tilde{F}_{A_i}(\tau_i(g-g')) f_G(g',t) dg'. \tag{43}$$

The unconditional mean $\mu_G(t)$ can then be written in terms of f_G:

$$\mu_G(t) = \int g\, f_G(g,t) dg. \tag{44}$$

The evolution equation for $\mu_G(t)$ is obtained in the same way as that of $\mu_{G|V}(v,t)$ above, so the intermediate steps are omitted. We multiply (42) by g and integrate with respect to g to obtain:

$$\begin{aligned}
\frac{d\mu_G}{dt}(t) &= -\int \frac{\partial \bar{J}_G}{\partial g}(g,t) dg \\
&= \int \bar{J}_G(g,t) dg \\
&= -\frac{1}{\tau_i}\mu_G(t) + \nu_i(t) \int dg \int_0^g dg' \tilde{F}_{A_i}(\tau_i(g-g')) f_G(g',t) \\
&= \frac{\nu_i(t)\mu_{A_i} - \mu_G(t)}{\tau_i},
\end{aligned} \tag{45}$$

which is equation (25).

B A refractory state probability density function

In general, a neuron becomes refractory for a period after firing. During the refractory period, a neuron's voltage does not move, though its synaptic conductances evolve as usual.

The presence of a refractory period adds complications to the full population density model. These complications stem from the fact that refractory neurons are not accounted for in $\rho(v, g, t)$. Since refractory neurons evolve according to different rules, we track them by a separate probability density function f_{ref}:

$$f_{ref}(u, g, t)du\,dg = \Pr\{U(t) \in (u, u + du) \text{ and } G_i(t) \in (g, g + dg)\}, \tag{46}$$

where $U(t)$ is the time since the neuron fired. Since f_{ref} accounts for only refractory neurons and a neuron becomes non-refractory when $U(t) = \tau_{ref}$, the function f_{ref} is defined only for $u \in (0, \tau_{ref})$. Since a neuron is either refractory or accounted for in $\rho(v, g, t)$, we have the conservation condition

$$\int\int \rho(v, g, t)dv\,dg + \int\int f_{ref}(u, g, t)du\,dg = 1. \tag{47}$$

The evolution equation of $f_{ref}(u, g, t)$ is:

$$\frac{\partial f_{ref}}{\partial t}(u, g, t) = -\left(\frac{\partial J_U}{\partial u}(u, g, t) + \frac{\partial J_{G'}}{\partial g}(u, g, t)\right), \tag{48}$$

where we denote the the flux across U by J_U and flux across inhibitory conductance for a refractory neuron by $J_{G'}$. The flux across U is simply

$$J_U(u, g, t) = \frac{dU}{dt} f_{ref}(u, g, t) = f_{ref}(u, g, t) \tag{49}$$

because $\frac{dU}{dt} = 1$ (the time since firing increases identically with time). Consequently, the boundary condition at $u = 0$ is $f_{ref}(0, g, t) = J_V(v_{th}, g, t)$ since neurons that fire immediately become refractory.

The flux across conductance for refractory neurons, $J_{G'}(u, g, t)$, is identical in form to the flux across conductance for non-refractory neurons (12):

$$J_{G'}(u, g, t) = -\frac{g}{\tau_i} f_{ref}(u, g, t) + \nu_i(t) \int_0^g \tilde{F}_{A_i}(\tau_i(g - g')) f_{ref}(u, g', t)dg'. \tag{50}$$

Since $J_U(\tau_{ref}, g, t)$ is the flux of neurons becoming non-refractory, it is the source of probability at v_{reset} in (8).

The refractory probability density function f_{ref} plays no role in the reduced model because we no longer need to track the inhibitory conductance of refractory neurons. Nonetheless, since $f_V(v, t)$ doesn't integrate to 1, we still have the modified conservation condition:

$$\int f_V(v, t)dv + \int_{t-\tau_{ref}}^{t} r(t')dt' = 1, \tag{51}$$

which replaces (47).

C Individual neuron simulation procedure

Numerical methods for the individual neuron simulations were similar to those of appendix C of Nykamp and Tranchina (2000). Although the evolution of the voltage and inhibitory conductance were no longer computed analytically between input events, the simulation was still based on an event-driven simulation. The event-driven procedure was used because the excitatory synaptic events were still delta functions. A time stepping algorithm where an excitatory input did not coincide with a time step would have produced significant error. Furthermore, since a neuron could cross threshold only at the time of an excitatory synaptic input, the event-driven simulation generalizes easily to the slow inhibition case.

In this approach, a neuron's voltage is calculated only when it receives synaptic input. When a neuron receives synaptic input, its voltage and inhibitory synaptic conductance are evolved up to the time of the input. Then either the voltage or the inhibitory conductance are jumped up, depending if the input were excitatory or inhibitory, respectively. We used a 2nd order Runge-Kutta method to evolve the voltage and conductance up to the synaptic input time. For each event, the step size of the Runge-Kutta method is chosen so that the synaptic input time coincides exactly with a time step. After each excitatory input voltage jump, the voltage is compared with threshold to determine if the neuron spiked. If it spiked, the voltage is reset to v_{reset}.

This hybrid method between a time stepping algorithm and an event-driven simulation is an efficient way to simulate the hybrid neurons with both delta function and slower synaptic conductances.

D Parameters

We used a gamma distribution for the distribution of unitary synaptic conductance sizes $A_{e/i}$:

$$f_{A_{e/i}}(x) = \frac{\exp(-x/a_{e/i})}{a_{e/i}(n_{e/i}-1)!} \left(\frac{x}{a_{e/i}}\right)^{n_{e/i}-1}, \tag{52}$$

where $f_{A_{e/i}}(x)$ is the probability density function of $A_{e/i}$.

For excitatory synapses, since $\Gamma_e^* = 1 - \exp(-A_e/c)$, the complementary cumulative distribution function (3) for Γ_e^* is

$$\tilde{F}_{\Gamma_e^*}(x) = e^{-u(x)} \sum_{l=0}^{n_e-1} \frac{u(x)^l}{l!}, \tag{53}$$

where $u(x) = -(c/a_e)\log(1-x)$.

We used a coefficient of variation of 0.5 for all A's. Thus, once $\mu_{A_{e/i}}$, the average value of $A_{e/i}$, was chosen, the values of $a_{e/i}$ and $n_{e/i}$ were determined. For the single and two population simulations, $\mu_{A_{e/i}}$ was chosen randomly. For the model of a hypercolumn in visual cortex, we used the following: $\mu_{A_e} = 0.008/c$, $\mu_{A_i} = 0.027/c$ for excitatory neurons, and $\mu_{A_e} = 0.020/c$, $\mu_{A_i} = 0.066/c$ for inhibitory neurons. These values corresponded to postsynaptic potentials that were half those used in the model by Somers et al. (1995).

For both excitatory and inhibitory neurons, we set $\mathcal{E}_i = -70$ mV, $\mathcal{E}_r = v_{reset} = -65$ mV, $v_{th} = -55$ mV, and $\mathcal{E}_e = 0$ mV. For excitatory neurons, $c/g_r = \tau_m = 20$ ms and $\tau_{ref} = 3$ ms, and for inhibitory neurons, $c/g_r = \tau_m = 10$ ms and $\tau_{ref} = 1$ ms. We used the parameters for excitatory neurons for single population simulations.

For the distribution of synaptic latencies $\alpha(t)$, we used:

$$\alpha(t) = \begin{cases} \bar{\alpha}\dfrac{\exp(-t/\tau_\alpha)}{\tau_\alpha(n_\alpha - 1)!}\left(\dfrac{t}{\tau_\alpha}\right)^{n_\alpha - 1} & \text{for } 0 \leq t \leq 7.5 \text{ ms} \\ 0 & \text{otherwise} \end{cases} \tag{54}$$

where $n_\alpha = 9$, $\tau_\alpha = 1/3$ ms, and $\bar{\alpha}$ is a constant so that $\int \alpha(t)dt = 1$. We used the same α for all interactions.

For the model of hypercolumn in visual cortex, the connectivity W between two populations depended on their difference in preferred orientation. The connectivity between a presynaptic population of type j and a postsynaptic population of type k, for $j, k \in \{E, I\}$, whose preferred orientation differed by $l\Delta\theta$ was

$$W_{jkl} = \begin{cases} \bar{W}_{jk} C_j \exp(-|l\Delta\theta|^2/(2\sigma_j^2)) & \text{for } |l\Delta\theta| < 60° \\ 0 & \text{otherwise} \end{cases} \tag{55}$$

where C_j was a constant such that $\sum_l W_{jkl} = \bar{W}_{jk}$. Thus $\bar{W}_{jk}$, for $j, k \in \{E, I\}$, is the average total number of synapses from all populations of type j onto each neuron of type k. We used $\bar{W}_{EE} = 72$, $\bar{W}_{EI} = 112$, $\bar{W}_{IE} = 48$, and $\bar{W}_{II} = 32$, which were double the numbers from Somers et al. (1995). We let $\sigma_I = 60°$ and $\sigma_E = 7.5°$.

References

Abbott LF, van Vreeswijk C (1993) Asynchronous states in networks of pulse-coupled oscillators. *Phys. Rev. E*, 48:1483–1490.

Baer SM, Erneux T, Rinzel J (1989) The slow passage through a Hopf bifurcation: delay, memory effects, and resonance. *SIAM J. Appl. Math.*, 49: 55–71.

Chawanya T, Aoyagi A, Nishikawa I, Okuda K, Kuramoto Y (1993) A model for feature linking via collective oscillations in the primary visual cortex. *Biol. Cybern.*, 68:483–490.

Galarreta M, Hestrin S (1997) Properties of $GABA_A$ receptors underlying inhibitory synaptic currents in neocortical pyramidal neurons. *J. Neurosci.*, 17:7220–7227.

Gerstner W (2000) Population dynamics of spiking neurons: fast transients, asynchronous states, and locking. *Neural Computation*, 12:43–89.

Knight BW (1972a) Dynamics of encoding in a population of neurons. *J. Gen. Physiol.* 59:734–766.

Knight BW (1972b) The relationship between the firing rate of a single neuron and the level of activity in a population of neurons. Experimental evidence for resonant enhancement in the population response. *J. Gen. Physiol.* 59:767–778.

Knight BW (2000) Dynamics of encoding in neuron populations: some general mathematical features. *Neural Computation*, to appear.

Knight BW, Manin D, Sirovich L (1996) Dynamical models of interacting neuron populations. In EC Gerf, ed. Symposium on Robotics and Cybernetics: Computational Engineering in Systems Applications. Cite Scientifique, Lille, France.

Kuramoto Y (1991) Collective synchronization of pulse-coupled oscillators and excitable units. *Physica D*, 50:15–30.

Nykamp D, Tranchina D (2000) A population density approach that facilitates large-scale modeling of neural networks: analysis and an application to orientation tuning. *J. Comp. Neurosci.*, 8:19–50.

Omurtag A, Knight BW, Sirovich L (2000) On the simulation of large populations of neurons. *J. Comp. Neurosci.*, 8:51–63.

Sirovich L, Knight BW, Omurtag A (2000) Dynamics of neuronal populations: the equilibrium solution. *SIAM*, to appear.

Stern P, Edwards FA, Sakmann B (1992) Fast and slow components of unitary EPSCs on stellate cells elicited by focal stimulation in slices of rat visual cortex. *J. Physiol. (Lond)*, 449:247–78.

Strogatz SH, Mirollo RE (1991) Stability of incoherence in a population of coupled oscillators. *J. Stat. Phys.*, 63:613–635.

Tamás G, Buhl EH, Somogyi P (1997) Fast IPSPs elicited via multiple synaptic release sites by different types of GABAergic neurone in the cat visual cortex. *J. Phys.*, 500:715–738.

Treves A (1993) Mean-field analysis of neuronal spike dynamics. *Network*, 4:259–284.

Wilson HR, Cowan JD (1972) Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical J.*, 12:1–24