J Comput Neurosci
DOI 10.1007/s10827-010-0253-4

Synaptic information transfer in computer models of neocortical columns

Samuel A. Neymotin · Kimberle M. Jacobs · André A. Fenton · William W. Lytton

Received: 15 December 2009 / Revised: 24 May 2010 / Accepted: 31 May 2010

Abstract Understanding the direction and quantity of information flowing in neuronal networks is a fundamental problem in neuroscience. Brains and neuronal networks must at the same time store information about the world and react to information in the world. We sought to measure how the activity of the network alters information flow from inputs to output patterns. Using neocortical column neuronal network simulations, we demonstrated that networks with greater internal connectivity reduced input/output correlations from excitatory synapses and decreased negative correlations from inhibitory synapses, measured by Kendall's τ correlation. Both of these changes were associated with reduction in information flow, measured by normalized transfer entropy (*n*TE). Information handling by the network reflected the degree of internal connectivity. With no internal connectivity, the feedforward network *transformed* inputs through nonlinear summation and thresholding. With greater connectivity strength, the recurrent network *translated* activity and information due to contribution of activity from intrinsic network dynamics. This dynamic contribution amounts to added information drawn from that stored in the network. At still higher internal synaptic strength, the network corrupted the external information, producing a state where little external information came through. The association of increased information retrieved from the network with increased gamma power supports the notion of gamma oscillations playing a role in information processing.

Keywords Information transfer · Neuronal networks · Simulation · Modeling

Action Editor: P. Dayan

This research was supported by NIH grants MH082417 and MH057068 and by DARPA grant N66001-10-C-2008.

S. A. Neymotin (✉) · A. A. Fenton · W. W. Lytton
Biomedical Engineering,
SUNY Downstate Medical Center, 450 Clarkson Avenue,
P.O. Box 31, Brooklyn, NY 11203-2098, USA
e-mail: samn@neurosim.downstate.edu

K. M. Jacobs
Departments of Anatomy and Neurobiology,
Virginia Commonwealth University,
Richmond, VA 23284, USA
e-mail: kmjacobs@vcu.edu

A. A. Fenton · W. W. Lytton
Department of Physiology and Pharmacology,
SUNY Downstate Medical Center, 450 Clarkson Avenue,
P.O. Box 31, Brooklyn, NY 11203-2098, USA

A. A. Fenton
Center for Neural Science, New York University,
4 Washington Place, New York, NY 10003-6621, USA
e-mail: afenton@nyu.edu

A. A. Fenton
e-mail: afenton@downstate.edu

W. W. Lytton
Department of Neurology,
SUNY Downstate Medical Center,
450 Clarkson Avenue, P.O. Box 31,
Brooklyn, NY 11203-2098, USA
e-mail: billl@neurosim.downstate.edu

1 Introduction

Understanding information flow in the brain is a fundamental problem in neuroscience (Victor 2006; Rieke et al. 1999; Salinas and Sejnowski 2001; Destexhe and Contreras 2006; Vogels et al. 2005). Brains and

Published online: 17 June 2010

neuronal networks must both store information and pass information from a sensory periphery reliably (Aldworth et al. 2005; Lazar and Pnevmatikakis 2008). There is a spectrum of network involvement in signal processing. At the extreme of a pure feedforward network, neurons are uncoupled and fire independently of each other. Any information processing that occurs is then not at the network level, but rather at the level of individual neurons. Addition of recurrent network connections permits activity flow through multiple loops of connected neurons. The dynamics of the network then permit novel combinations to produce different output patterns, as implemented in an echo state network (Jaeger and Haas 2004; Buonomano and Maass 2009). From an information viewpoint, one can view the synapses in a recurrent network as the location of stored patterns produced through learning. Information that is present in the network can then be added to information coming in from the inputs so as to provide more complex output patterns.

We here investigate how the strength of synaptic weights influences information flow-through in a neocortical column simulation. We focus on neocortex because it is a functionally specialized information storage and information processing structure. Neocortex is thought to use its complex architecture to provide multiple levels of recurrent connections, producing interlaced subnetworks that may play complementary roles in information processing (Douglas et al. 1989; Hill and Tononi 2005; Izhikevich and Edelman 2008; Nelson 2002; Sporns et al. 2005). In this study, we quantified global activity transfer using Kendall's tau correlation (Kendall 1938) and information transfer using normalized transfer entropy (*n*TE) (Gourevitch and Eggermont 2007).

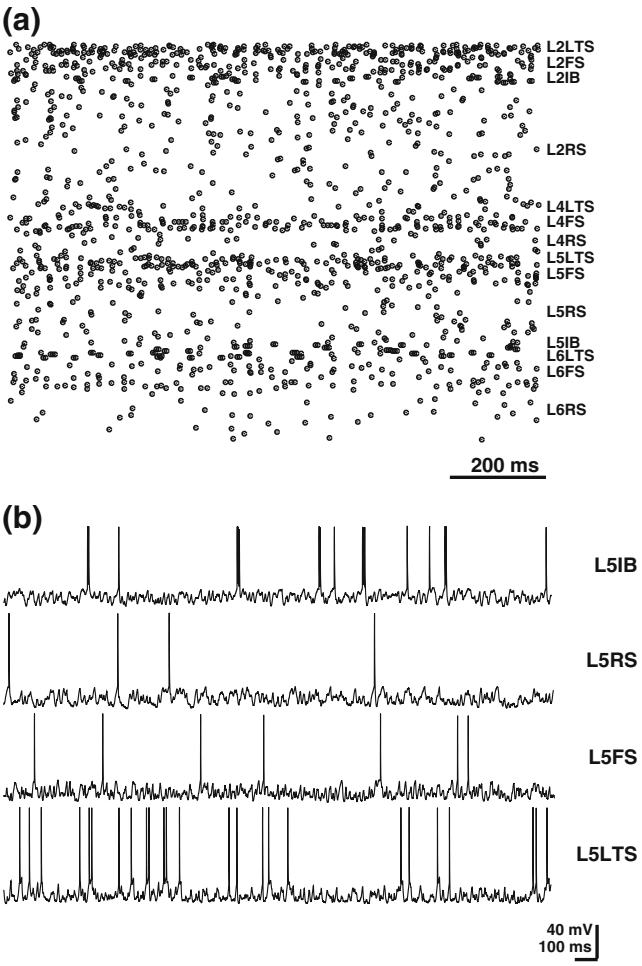

Fig. 1 (**a**) Raster plot of 1 s of simulation time. Cells are sorted and labeled by layer and type, i.e., L5RS is a layer 5 regular spiking cell. (**b**) Representative activity traces from axon of four cells taken from raster. L5IB—layer 5 intrinsically bursting excitatory cell; L5FS—layer 5 fast spiking interneuron. L5LTS—layer 5 low-threshold spiking interneuron

2 Methods and materials

2.1 Cortical column simulation

Our model of neocortical column consisted of four cell classes oriented in four layers: layer 2/3 (listed as simply 2), 4, 5 and 6. This produced only 14 different cell groups since not all cell classes are present in every level (Fig. 1). Excitatory cells were either regular-spiking (RS) or intrinsically bursting (IB). Inhibitory interneurons consisted of fast-spiking (FS), and low-threshold spiking (LTS) cells. The model consisted of 470 cells (Table 1).

Each cell was implemented as a three-compartment model: dendrite (d), soma (s), and axon (a). Each compartment had 0–3 intrinsic voltage-sensitive channels (Friesen and Friesen 1994). Dendrites had no active channels, and were the location of the excitatory AMPA and NMDA excitatory inputs and some inhibitory $GABA_A$ inputs. Other $GABA_A$ inputs synapsed onto the soma. Dendrites were 500 μm long with a diameter of 2 μm, the soma was 30 μm long with a diameter of 30 μm, and the axon was 200 μm long with a diameter of 1 μm.

Each compartment is modeled using a variant of the standard Hodgkin–Huxley parallel conductance model, using the following equations:

$$-C \cdot \dot{V} = I_{\text{pas}} + I_{\text{v}} + I_{\text{syn}} \tag{1}$$

$$I_{\text{v}} = g_{\text{i}} \cdot (V - \text{Erev}_{\text{i}}) \quad \text{for i} = K, H, A \ldots \tag{2}$$

$$g_{\text{i}} = \bar{g}_{\text{i}} \cdot m_{\text{i}} \cdot h_{\text{i}} \quad \text{for i} = H, A, \ldots \tag{3}$$

Table 1 Number of each cell type used in simulations

Type	L2RS	L2IB	L2LTS	L2FS	L4RS	L4LTS	L4FS	L5RS	L5IB	L5LTS	L5FS	L6RS	L6LTS	L6FS
Count	142	8	13	25	30	14	20	65	17	13	25	60	13	25

L4RS represents regular spiking (RS) excitatory cells in layer 4

$$\dot{x} = \frac{-x + x_\infty}{\tau_x} \quad \text{for } x = K, m, h \tag{4}$$

$$x_\infty = 1 \Big/ \left(1 + exp\left(\frac{-V + V_{1/2}}{\text{slope}}\right)\right) \tag{5}$$

In Eq. (1), I_v are voltage-sensitive channel currents, I_{syn} synaptic currents and I_{pas} passive (leak) current. One cell type, layer 2 IB, also had injected current in soma, to provide spontaneous bursting. Dendrites had no voltage sensitive channels. Excitatory synapses (NMDA, AMPA) were on dendrites; inhibitory ($GABA_A$) on dendrites and soma. Synaptic currents were parameterized as previously (Lytton 1996; Destexhe et al. 1994). The m, h particles of Eq. (3) are not exponentiated for any channels.

Each compartment had a passive current described by Eq. (2) with Erev_{pas} (in mV) = −65 (RS all compartments), −69 (IB all), −60 (LTS s), −69 (LTS d,a), −63 (FS all); $\bar{g}_{pas}$ (in nS) = 0.26, 0.29, 1.32 (RS d,s,a); 0.26, 0.07, 1.32 (IB d,s,a); 0.17, 0.20, 0.83 (LTS d,s,a); 0.29, 0.15, 1.43 (FS d,s,a). Axial resistivity, R_a , was same for all cell types: 25.13 Ω-cm for d–s, 15.71 Ω-cm for s–a. Membrane capacitance was adjusted in order to get desired time constant/input impedance relations for single compartments representing large structures: C_m (in $\mu F/\text{cm}^2$) = 5.02, 1, 25.10 (RS d,s,a), 5.02, 5.58, 25.10 (IB d,s,a), 1.97, 2.19, 9.85 (LTS d,s,a), 2.38, 3.8, 11.89 (FS d,s,a).

All cell types had I_A channels in the soma using the parameterizations of Eqs. (2)–(5) with $\bar{g}_A$ (in nS) = 5.13 (RS), 35.5 (E2B), 7.22 (E5B), 30.4 (LTS), 7 (FS). Erev_A (in mV) = −82 (RS), −116 (IB), −85 (LTS), −99 (FS). Equation (5) slopes (in mV for m,h) = 3.2, −4.9 (RS); 8.4, −5.9 (IB), 11.4, −4.9 (LTS), 9.4, −1.9 (FS); $V_{\frac{1}{2}}$ (in mV for m,h) = −24, 8 (RS), 10, 3 (IB), 54, −45 mV (LTS); 13, −14 (FS). τ_m,τ_h (in ms) = 82, 5 ms (RS); 5, 20 (IB), 2, 37 (LTS), 2, 5 (FS). IB cells had hyperpolarization-activated (anomalous rectifier) I_H channels in the soma with $\text{Erev}_H = -20$ mV, $\bar{g}_H$ = 17.68 nS. For the single activation variable m: $V_{\frac{1}{2}} = -74$ mV, slope = −10 mV, $\tau_m = 50$ ms. E2B cells had L-type Ca^{2+} channels in their soma, responsible for calcium spikes and parameterized as in Zhu et al. (1999a, b).

The axon used a hybrid rule/conductance-based spike-generator, that checked for voltage crossings above a threshold (Friesen and Friesen 1994). When axon voltage crossed threshold, the Na^+ channel activated instantly to $\bar{g}_{Na}$ and was maintained for duration dur_{Na}. Erev_{Na} (in mV) = 55 (RS), 50 (IB), 51 (LTS), 55 (FS); $\bar{g}_{Na}$ (in μS) = 0.6 (RS), 0.3 (IB), 0.6 (LTS), 0.3 (FS). dur_{Na} (in ms) = 0.9 (RS), 1.2 (IB), 0.5 (LTS), 0.6 (FS). Absolute refractory period lasted (in ms): 11 (RS), 4.7 (IB), 11 (LTS), 2.9 (FS). A relative refractory period was used for IB and LTS cells. After a spike adaptive threshold, $V_{\text{th\_adapt}}$ incremented by $V_{\text{th\_inc}}$ and then decayed to baseline (Eq. (4) where x_∞ is baseline). Baseline (in mV) = −40 (RS,IB,FS; fixed for RS,FS), −47 (LTS); $V_{\text{th\_inc}}$ (in mV) = 1.3 (IB), 3.5 (LTS); τ = 70 (IB), 89 (LTS).

A threshold-crossing-activated K^+ channel with adaptation was present in the axon. At threshold crossing, the K^+ adaptation variable, g_{kadapt}, increased by g_{kinc} and then decayed with τ_{kadapt} and kadapt_∞ (Eq. (4)). At the end of an action potential, the K^+ conductance, g_K, was set to the current value of g_{kadapt}. g_K then decayed with τ_K and $gK_\infty = 0$ (Eq. (4)). kadapt_∞ (in nS) = 91 (RS), 70 (IB), 60 (LTS), 58 (FS); τ_{kadapt} (in ms) = 69 (RS), 80 (IB), 107 (LTS), 1 (FS). τ_K (in ms) = 2.3 (RS), 1.2 (IB), 1.5 (LTS), 1.5 (FS). Erev_K (in mV) = −82 (RS), −116 (IB), −85 (LTS), −91 (FS). g_{kinc} (in nS) = 13 (RS), 2 (IB), −0.3 (LTS), 0 (FS).

Cells were randomly connected by type and layer using the connection densities shown in Table 2. Note that the connection densities are several times higher than in real neuronal networks. This was necessary to sustain activity due to the small scale of the model. Local field potential (LFP) was calculated at each time step as the sum of differences of somatic and dendritic voltages across all the excitatory cells.

A multiplicative gain factor to control the strength of all intrinsic synaptic weights (excitatory and inhibitory) ranged from 0 (no connectivity) to 1. We provided randomly-timed sub-threshold synaptic inputs independently to each synapse from 160–240 Hz for excitatory synapses, and 80–120 Hz for inhibitory synapses. The rate of inputs to each synapse was selected from the ranges above using a uniform distribution. Once the rate was selected, the timing of the inputs was determined by drawing samples from a uniform distribution to maintain the selected rate for each synapse. These external inputs were independent for each dendritic AMPA synapse, dendritic NMDA synapse, somatic $GABA_A$ synapse, and dendritic $GABA_A$ synapse. The

Table 2 Connection densities between cell types

From ⇓	To ⇒													
	L2RS	L2IB	L2LTS	L2FS	L4RS	L4LTS	L4FS	L5RS	L5IB	L5LTS	L5FS	L6RS	L6LTS	L6FS
L2RS	0.187	0.187	0.51	0.51	0.024	–	–	0.057	0.024	–	–	–	–	–
L2IB	0.187	0.187	0.51	0.51	0.024	–	–	0.057	0.024	–	–	–	–	–
L2LTS	0.35	0.35	0.09	0.09	–	–	–	0.35	0.5	–	0.53	0.25	–	0.53
L2FS	0.35	0.35	0.09	0.09	–	–	–	0.35	0.5	–	0.53	0.25	–	0.53
L4RS	0.145	–	–	–	0.243	0.51	0.43	0.116	0.122	–	–	0.032	–	–
L4LTS	–	–	–	–	0.35	0.09	0.53	–	–	–	–	–	–	–
L4FS	–	–	–	–	0.44	0.34	0.62	–	–	–	–	–	–	–
L5RS	0.022	–	–	–	0.007	–	–	0.191	0.08	0.51	0.43	0.032	–	–
L5IB	0.018	–	–	–	0.007	–	–	0.017	0.07	0.51	0.43	0.07	–	–
L5LTS	0.35	–	–	–	–	–	–	0.35	0.35	0.09	0.53	0.25	–	0.53
L5FS	–	–	–	–	–	–	–	0.44	0.44	0.34	0.62	–	–	–
L6RS	–	–	–	–	–	–	–	0.006	0.028	–	–	0.028	0.51	0.43
L6LTS	0.35	–	–	–	–	–	–	0.25	0.25	–	0.53	0.35	0.09	0.53
L6FS	–	–	–	–	–	–	–	–	–	–	–	0.44	0.34	0.62

random external inputs were utilized not only to maintain network activity but also as the source of inputs for the correlation and information measures of external inputs to spike outputs. Excitatory synapses were located on the dendrite; $GABA_A$ synapses on soma or dendrite.

We initially tuned the networks by providing lower strength external inputs to LTS, so that they started with low firing rates. We then increased their rates with larger internal gains from excitatory cells. In this way, we avoided the positive feedback latch-up seen with increased synaptic strength and increased excitatory firing rates (Lytton and Omurtag 2007). The FS cells were connected so that increasing internal synaptic strengths resulted in lower rates through the influence of the LTS cells. This was required in order to avoid over-inhibition of the excitatory cells.

The major results were taken from simulations using five different internal synaptic weight strengths, ten different sets of random external input streams and five different wiring randomizations for a total of 250 100 s simulations. Each of the ten different external input streams had the same strengths, but were different outcomes of the same random process. Additional 30 s simulations were run for various controls mentioned in results. Going from the unconnected to the connected network results in substantial rate changes in some cell groups, primarily the fast-spiking inhibitory cells. Subsequent gain increases then produced moderate increases in most groups with further decreases in the fast-spiking populations (Table 3). In some control networks, we hand-tuned the weights of external excitatory and inhibitory inputs to balance against changes in internal gain and preserve comparable rates. We also tested variable synaptic change in two networks: starting with a gain factor of 0.25, we multiplied each synaptic weight by a random value from a uniform distribution on the interval of [0.4, 3.6] to produce a final synaptic gain in range [0.1, 0.9] (rates in RPD column in Table 3). This allowed for both potentiation and depression of synapses.

Table 3 Average ± standard-error of firing rates for the different cell types in 100 s of simulation with different internal gain factors shown in columns

(RPD: Network with randomly potentiated/depressed synapses run for only 30 s)

Type ↓	0	0.25	0.5	0.75	1.0	RPD
L2LTS	1.1 ± 0.1	5.3 ± 0.3	12.2 ± 0.5	16.2 ± 0.6	18.8 ± 0.7	12.9 ± 0.6
L2FS	16.2 ± 0.6	6.3 ± 0.4	5.7 ± 0.4	4.6 ± 0.4	3.8 ± 0.4	5.9 ± 0.4
L2IB	3.0 ± 0.2	3.0 ± 0.4	5.6 ± 0.6	7.3 ± 0.7	8.5 ± 0.8	5.9 ± 0.7
L2RS	1.7 ± 0.0	1.0 ± 0.0	1.5 ± 0.1	1.7 ± 0.1	1.7 ± 0.1	1.6 ± 0.1
L4LTS	0.9 ± 0.1	1.1 ± 0.1	3.3 ± 0.4	6.0 ± 0.6	8.3 ± 0.8	3.6 ± 0.5
L4FS	16.9 ± 0.7	6.2 ± 0.5	6.2 ± 0.6	6.1 ± 0.7	6.1 ± 0.8	6.3 ± 0.6
L4RS	1.7 ± 0.1	1.0 ± 0.1	1.7 ± 0.1	2.3 ± 0.1	2.6 ± 0.2	1.8 ± 0.1
L5LTS	0.9 ± 0.1	3.5 ± 0.3	8.6 ± 0.5	12.5 ± 0.6	15.0 ± 0.7	8.9 ± 0.6
L5FS	15.2 ± 0.7	5.4 ± 0.4	4.1 ± 0.3	3.1 ± 0.3	2.2 ± 0.2	4.3 ± 0.3
L5RS	1.7 ± 0.1	0.9 ± 0.1	1.3 ± 0.1	1.6 ± 0.1	1.6 ± 0.1	1.5 ± 0.1
L5IB	8.4 ± 0.3	5.0 ± 0.4	5.1 ± 0.5	4.8 ± 0.6	4.5 ± 0.7	5.1 ± 0.6
L6LTS	0.9 ± 0.1	1.3 ± 0.1	3.1 ± 0.2	5.2 ± 0.4	7.1 ± 0.4	3.2 ± 0.3
L6FS	16.1 ± 0.6	4.7 ± 0.3	2.8 ± 0.3	1.6 ± 0.2	1.0 ± 0.2	2.8 ± 0.3
L6RS	1.7 ± 0.1	0.6 ± 0.0	0.7 ± 0.1	0.7 ± 0.1	0.7 ± 0.1	0.7 ± 0.1

Simulations were run under NEURON using Crank–Nicholson integration with dt = 0.025 ms (Carnevale and Hines 2006; Hines and Carnevale 2001). Organization of simulations and analysis of output data was aided by Neural Query System (Lytton 2006; Lytton and Stewart 2007). Simulation and analysis software is posted on ModelDB (http://senselab.med.yale.edu/modeldb). Simulations and analysis were run under Linux on a 2.7 GHz quad-core Intel Xeon CPU. A 100 s simulation ran in approximately 80 min.

2.2 Spectral analysis

To correlate changes in network activity with activity at the macroscopic level we analyzed the local field potentials of pyramidal cells. Spectral analysis was performed using Matlab (2009b, The MathWorks. Natick, MA). Since simulations were run with a time-step of 0.025 ms this resulted in a 40 KHz sampling rate. To determine spectral power in the different frequency bands, we used the standard Matlab spectrogram function with a WINDOW size of 32768, NFFT of 131072, and NOVERLAP of 16384. To determine the power in a frequency band we used the resulting spectrogram and summed the values in the given frequency range, across time. We then applied a base 10 logarithm to the sum to allow for easier comparison of power across the different networks. All comparisons were performed on simulations run for equal amounts of time.

2.3 Kendall's τ

Other covariance measures cannot handle the preponderance of zeros encountered using short bin sizes (e.g., 100 ms), which must be counted as ties (Press et al. 2007). For this reason, Kendall's τ non-parametric rank correlation, defined as:

$$\tau = \frac{n_c - n_d}{\frac{1}{2}n(n-1)},$$

is used with these data (Kendall 1938). Kendall's τ is a normalized difference between concordant (n_c) and discordant pairs (n_d); ties are taken into account by the normalizing term, $\frac{1}{2}n(n-1)$, which represents the total number of ordered pairs in the time-series. To reduce the runtime of the calculation from $O(n^2)$ to $O(n \cdot \log(n))$, we used a fast version of Kendall's τ based on sorting procedures (Knight 1966). We note that high correlation between inputs and outputs does not imply high transfer of information: two identical time-series are fully correlated, but there is no information transfer. Information transfer requires a time delay.

2.4 Transfer entropy

Information-theoretic directed measures between two processes have previously been described (Jumarie 1990; Paluš 1996; Hlaváčková-Schindler et al. 2007). Recently, transfer entropy (Schreiber 2000) was developed to determine the direction and quantity of information transferred between two processes, $X1$ and $X2$ (from Eq. 1 in Gourevitch and Eggermont 2007):

$$TE_{X1\to X2} = H\left(X2_{\text{future}}|X2_{\text{past}}\right) - H\left(X2_{\text{future}}|X2_{\text{past}}, X1_{\text{past}}\right) \quad (6)$$

To calculate the transfer entropy from $X1$ to $X2$ we cannot simply use $H\left(X2_{\text{future}}|X1_{\text{past}}\right)$, because this ignores the possibility that both processes may be driven in tandem by a hidden process that helps determine both their present and their future. We must therefore also consider the conditional entropy of $X2$'s future given its past ($H(X2_{\text{future}}|X2_{\text{past}})$) and the conditional entropy of $X2$'s future given both its past AND the past of $X1$, $H(X2_{\text{future}}|X2_{\text{past}}, X1_{\text{past}})$. When $X1$ is transferring information to $X2$, the additional knowledge of $X1$'s past will reduce the conditional entropy of $X2$'s future values given $X2$'s past. Transfer entropy quantifies how much entropy of $H(X2_{\text{future}}|X2_{\text{past}})$ is left over once $X1_{\text{past}}$ is known. Another key to transfer entropy is that it is different in two directions: knowing $X1$'s past may tell you something about $X2$'s future, but knowing $X2$'s past does not have to tell you anything about $X1$'s future.

A problem with transfer entropy is that bias may be introduced in the estimate. Several similar approaches were proposed to deal with this problem, effective transfer entropy (Marschinski and Kantz 2002) and normalized transfer entropy (nTE) (Gourevitch and Eggermont 2007). We used nTE, because it was formulated for working with neuronal data, specifically with spike counts per time:

$$nTE_{X1\to X2} = \frac{TE_{X1\to X2} - < TE_{X1_{\text{shuffled}}\to X2} >}{H(X2_{\text{future}}|X2_{\text{past}})} \quad (7)$$

nTE removes bias from the estimate of transfer entropy by subtracting the average transfer entropy from $X1$ to $X2$ using a shuffled version of $X1$ denoted $< TE_{X1_{\text{shuffled}}\to X2} >$, over several shuffles. It then divides the estimate by the entropy of $H(X2_{\text{future}}|X2_{\text{past}})$, to get a value between 0 and 1. nTE will be 0 when $X1$ transfers no information to $X2$, and will be 1 when $X1$ transfers maximal information to $X2$. If $X1$ is a particularly poor predictor of $X2$, then the shuffled $X1$ may on average, produce a better estimate of $X2$'s future. In this case, nTE can be less than 0. However, the

number of these cases can be reduced if both processes are observed for long periods of time and many shuffles are used. For all the nTE calculations, we shuffled $X1$ 30 times.

The main equation for nTE requires the formation of two probability distributions: $P(X2_{\text{future}}|X2_{\text{past}})$ and $P(X2_{\text{future}}|X2_{\text{past}}, X1_{\text{past}})$. These distributions, in turn, require calculating the following distributions: $P(X2_{\text{future}}, X2_{\text{past}})$, $P(X2_{\text{past}})$, $P(X2_{\text{future}}, X2_{\text{past}}, X1_{\text{past}})$, and $P(X2_{\text{past}}, X1_{\text{past}})$. The time-series we are working with consist of the counts of discrete events (spikes or synaptic inputs) per non-overlapping, consecutive time-bin. As a result, the 1-D probability distribution, $P(X2_{\text{past}})$, is defined on the interval from the minimum to maximum number of events in the given time-series. Then for this 1-D distribution, the probability of occurrence of n events (spikes or inputs) in a time-bin is the number of times that particular count occurs (n_c) in the time-series, divided by the total length of the time-series (t), or $P(n) = \frac{n_c}{t}$. The joint probability distribution $P(X2_{\text{future}}, X2_{\text{past}})$ requires iterating through the time-series and counting the number of occurrences of event doubles using $X2$'s future at time t_i and $X2$'s past using time t_{i-1}. In this case, because we start at t_1 instead of t_0, we normalize the probability by $t-1$, where t is the length of the time-series. The distributions $P(X2_{\text{past}}, X1_{\text{past}})$ and $P(X2_{\text{past}})$, $P(X2_{\text{future}}, X2_{\text{past}}, X1_{\text{past}})$ are formed in a similar fashion. Once the distributions are formed, the conditional entropies are calculated using the probabilities of event counts in the distributions, using:

$$H(X_{\text{future}}|X_{\text{past}}) = -\sum p(X_{\text{future}}, X_{\text{past}}) \times \log \frac{p(X_{\text{future}}, X_{\text{past}})}{p(X_{\text{past}})} \quad (8)$$

which is the conditional entropy of X_{future} given X_{past}, and uses a log of base 2.

We note that in the calculation of nTE since only $X1$ is shuffled, all entropies that only use $X2$, need only be calculated once. Other methods for calculating the probabilities of events, include kernels, which require iterating over all data-points to obtain a single probability. Due to efficiency reasons and because event counts are most natural for this type of data, we choose to use simple event counts. Although multiple time-bins in the past can be used for calculating nTE we found that the nTE values do not change significantly with increasing number of bins, and therefore used a single time-bin from the past. However, varying the size of the bin can have a large impact on the value of nTE and even determine whether a significant value will be detected (Gourevitch and Eggermont 2007). Additionally, since we are using one time-bin from the past, only interactions that occur in consecutive time-bins will be detected. For example, if there is a unidirectionally coupled process $X_1 \rightarrow X_2$, with a characteristic delay, t_d, then using a time-bin (t_b) smaller than t_d will result in the event at X_1 occurring at time t_1 and the effect occurring in X_2 at time $t_1 + t_d > t_1 + t_b$. This will result in the effect occurring in X_2 later than the next time-interval, and will be undetectable. Similarly, if the time-bin used, t_b, is greater than the delay, the effect will also not be detected in the next time-interval. Instead, the optimal bin-size is the delay time, because this will maximize the number of interactions occurring in consecutive time intervals. This example demonstrates how nTE allows detection of characteristic delays in directional interactions. This issue is explored in detail in Gourevitch and Eggermont (2007).

2.5 Analysis of simulation output

To assess the amount of information transferred from inputs to outputs, and to measure the strength of the modulation, we measured nTE and Kendall's τ non-parametric rank correlation (Kendall 1938; Press et al. 2007) across input/output pairs where the outputs were spike trains from a neuron and individual inputs were the signals at single synapses on that cell (once for each of AMPA, NMDA, $GABA_A$ to dendrite, $GABA_A$ to soma). Before nTE and correlation calculations, input streams and output spike-trains were binned as counts per time.

The bin size for Kendall's τ correlation calculations was 100 ms, and for nTE calculations it was 10 ms in Figs. 2 and 3 and 8 ms in Fig. 5. To determine how the nTE and Kendall's τ values changed in the different simulations, we took the average nTE and Kendall's τ across all the excitatory cells at a given synapse. Averaging was used because values across excitatory cells were similar. The resulting values gave us an estimate of overall information transfer and modulation.

2.6 Internal vs external information transfer

Internal information transfer was measured between neurons within the network. To measure the nTE between neurons within the network we formed input vectors for each cell's synapses. The input vectors were defined as event counts per time, and were summed from all spikes of presynaptic cells synapsing on the postsynaptic cell at the given synapse. The events occurred after the presynaptic cell spiked with the addition of the synaptic transmission delay. A multi-unit type of measure was used because the interactions

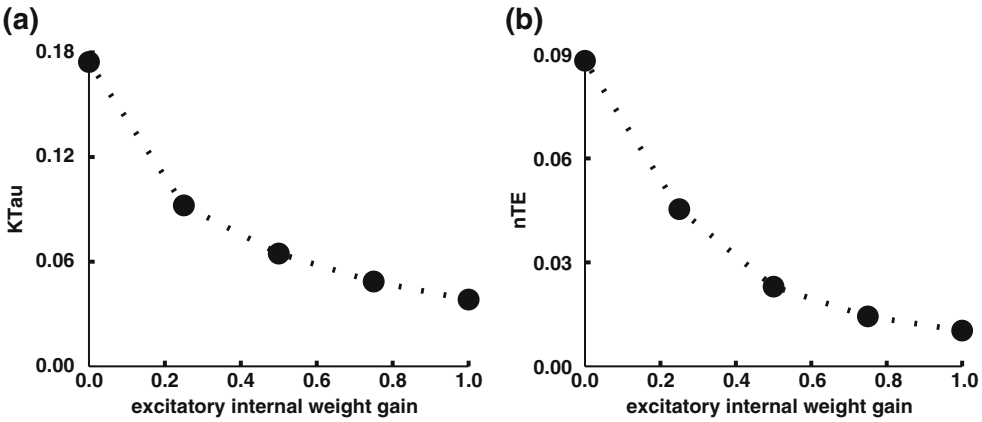

Fig. 2 Input/output relation at AMPA synapses. (**a**) Average Kendall's τ correlations between external inputs and output spikes at AMPA synapses of excitatory cells. (**b**) Average nTE from external inputs to output spikes at AMPA synapses of excitatory cells. Standard error bars within symbol size in (**a**, **b**)

between single cells was relatively weak. External information transfer was measured similarly, but the input vectors were formed from the external inputs driving each synapse independently (and not modeled explicitly as cells).

The counts in an 8 ms bin ranged from 0–9 interneuronal synaptic events for AMPA, NMDA and $GABA_A$ synapses, and from 0–8 and 0–11 external synaptic inputs for $GABA_A$ and AMPA/NMDA synapses respectively. Each 8 ms bin, on average had less than one event.

3 Results

3.1 Information transfer in neocortical column simulations

We ran over 250 30–300 s simulations, using five different random wirings and ten different input streams. The results shown were consistent across these random variations.

Internal synaptic weight strength was varied and activity flow (correlation) measured from external excitatory (AMPA) inputs onto excitatory cells, whose spiking provided the outputs. Our initial network was feedforward with no internal (recurrent) connectivity (gain 0). In this network, firing times for each neuron showed higher correlation with inputs (Fig 2(a)). With increasing internal weight gain, external input/output correlations gradually declined, suggesting that network dynamics interfered with direct activity flow. The network with the strongest internal weight gain showed little correlation, hence little direct flow from external inputs to outputs.

We next turned to nTE to quantify this activity flow alteration in information terms (Fig. 2(b)). In the uncoupled network, nTE demonstrated relatively high information transfer from external inputs to outputs: the neurons had no influence on each other, so each cell's external inputs were good predictors of its outputs. Note that information transfer at this level is still small (nTE is normalized to 1.0). This means that an individual input had relatively little influence on cell spiking. This is to be expected because five–ten

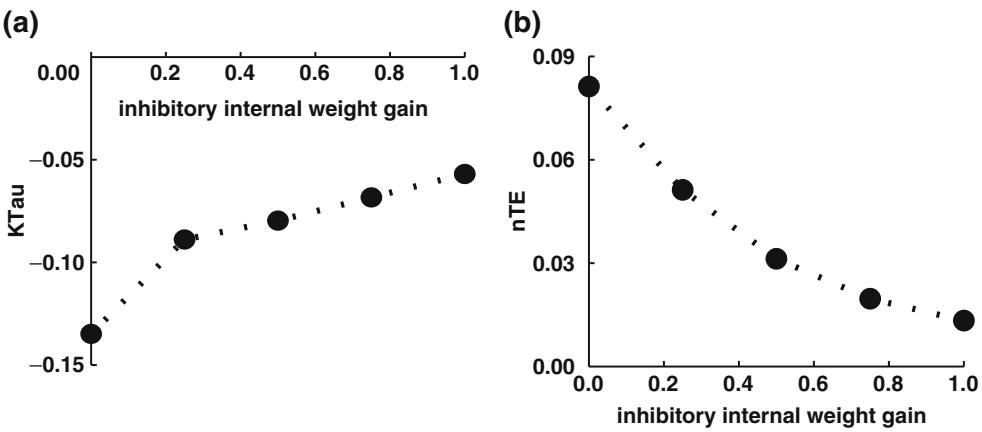

Fig. 3 Input/output via $GABA_A$ synapses. (**a**) Average Kendall's τ correlations between external $GABA_A$ inputs and output spikes for excitatory cells. (**b**) Average nTE from external $GABA_A$ inputs to output spikes for excitatory cells. Standard error bars within symbol size in (**a**, **b**)

simultaneous inputs were required to produce an output spike, the number varying by cell type and network state. As internal synaptic strengths were increased, lower information transfer from external inputs to outputs paralleled the reduction in correlation, as input was no longer able to predict system output as reliably. In these networks, information transfer at NMDA synapses was negligible even at large delays, demonstrating that NMDA synapses acted more as modulators, setting a baseline level of activity rather than transferring information. This was due to longer-lasting effects at NMDA synapses and lower synapse weights relative to AMPA.

We next examined information transfer from external inhibitory inputs to excitatory outputs ($GABA_A$ synapses onto excitatory cells). We again found consistent decrease in correlations with increase in internal weight, again closely paralleled by *n*TE changes (Fig. 3). Note that in this case the correlations are negative since increased inhibition leads to reduced firing. However a reduction in firing probability still represents an information transmission influence, with *n*TE values similar to those found with alterations in AMPA synapses.

We were initially concerned that higher frequency firing might be associated with greater information throughput, since individual inputs would then be more likely to be associated with an output spike by chance. On average, the results contradicted our hypothesis, with the generally increasing rates with increasing gain (seen for example in the RS cells: L2RS, L4RS, L5RS, L6RS in Table 3) being associated with decreasing *n*TE in aggregate and for these individual cell groups. Looking at particular examples, we found that the L4RS cells, with similar rates (∼1.7 Hz) in the disconnected and gain 0.5 networks, had cell-type *n*TE which decreased from 0.10 ± 0.003 (average ± standard error of the mean) to 0.03 ± 0.002 at AMPA synapses and from 0.09 ± 0.002 to 0.06 ± 0.003 at $GABA_A$ synapses. L2RS, the dominant cell type by number, had this same average firing rates in the 0, 0.75, 1.0 gain networks associated with *n*TE of 0.10 ± 0.001, 0.008 ± 0.0007, 0.005 ± 0.0009 at AMPA synapses respectively, and 0.10 ± 0.001, 0.02 ± 0.0007, 0.01 ± 0.0007 at $GABA_A$ synapses respectively. Looking in detail at the relationships for individual cells demonstrated significant weak negative correlations of rate and *n*TE in almost all cases. Figure 4 shows the relations at different gains for L2RS cells. Correlations between *n*TE and rate at AMPA synapses were larger than the corresponding correlations at $GABA_A$ synapses, at all gains.

To further address the possibility that our results might reflect a sensitivity of the *n*TE measure to changes in firing rate, we ran controls where we modulated the external weights in parallel with increasing internal weight gain in order to maintain comparable rates. External excitatory weight was typically increased for a given cell type, in order to overcome the increased internal inhibition generated from within the network. We observed the same pattern of *n*TE and correlation changes with these controls. In another set of controls, we sought to determine whether network

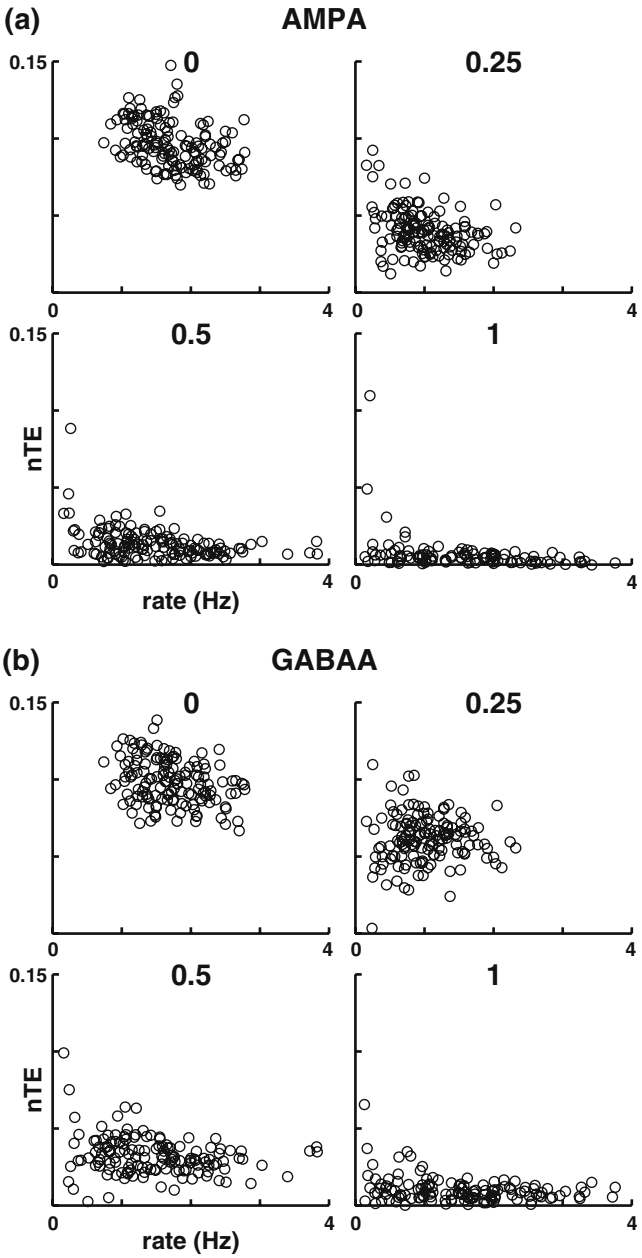

Fig. 4 *n*TE vs rate for individual L2RS cells at four gain factors (0, 0.25, 0.5, 1). (**a**) Transfer from AMPA synapses. (**b**) Transfer from $GABA_A$ synapses. Correlations are significant and negative [−0.36, −0.18] in all cases except $GABA_A$,0.25

effects would dominate over the increase in *n*TE expected with increasing input strength. We independently increased external weights by a given factor across cell types and synapses, in parallel with increasing internal weight gain. Although the results showed slight increases in input/output correlation, the changes in *n*TE were indistinguishable from our primary results.

To determine interneuronal information transfer, we measured *n*TE from inputs generated within the network to outputs of those cells, quantifying the decrease in entropy from knowing the past activity in presynaptic cells within the network itself (Fig. 5). Our external firing rates (80–240 Hz) represent combined inputs from many cells. The internal synaptic input rates are from single cells and therefore much lower (0.5–20 Hz), yielding small *n*TEs which nonetheless showed consistent patterns of increase. In order to provide multiunit internal input which could be directly compared with the external information transfer, we combined inputs to a given synaptic location to produce a single input vector. Using this measure, internal information transfer between neurons increased, and external information transfer decreased. At the higher gains, internal *n*TE surpassed external. Internal spike time correlations also increased at higher gains (not shown).

We also assessed information transfer by varying synaptic gain across different synapses, comparable to that which would be produced with independent synaptic depression and potentiation produced by a learning paradigm (Rao and Sejnowski 2001). We started with randomized internal wiring (gain 0.25) to represent a low level of information content in the network. We then multiplied each synaptic weight by a random value to produce weights with individual gains from 0.1 to 0.9, with an overall gain increase in the network. Table 3 shows firing rates for this randomly potentiated/depressed (RPD) network. Decrease in external information flow and increase in internal information flow with the overall gain increase was similar in magnitude to that seen with uniform gain increase (data not shown).

3.2 Information transmission delays

*n*TE can be used to measure the characteristic response times to incoming information, reaching a maximum when bin size matches input/output delay (Gourevitch and Eggermont 2007). Figure 6 shows that *n*TE peaks between 5–10.5 ms bin sizes for both AMPA and $GABA_A$ inputs. By contrast, Kendall's τ did not show a peak but instead gradually increased with bin size. With internal weights off, AMPA inputs peaked at 6.5 ms and $GABA_A$ inputs peaked at 7 ms. Because time constants and localizations are identical for the two synapse types, the slightly shorter delay for AMPA is due to the difference in polarity, likely a result of AMPA recruitment of rapid depolarization-gated currents.

With internal weights off, the similarity in delay times and *n*TE values suggests a close balance between inhibitory and excitatory effectiveness for information transmission. When internal weights were increased, the AMPA synapse peak delay time increased from 6.5 to 10.5 ms, and the $GABA_A$ synapse peak delay time decreased from 7 to 5 ms. *n*TE value declined more for AMPA than for $GABA_A$. The increasing delay of information throughput from AMPA inputs reflects the activation of network dynamics from excitatory inputs, resulting in an influence that, though weaker, lasts further into the future.

To test *n*TE and information transmission delays at NMDA synapses, we developed several new networks by gradually increasing NMDA strength externally

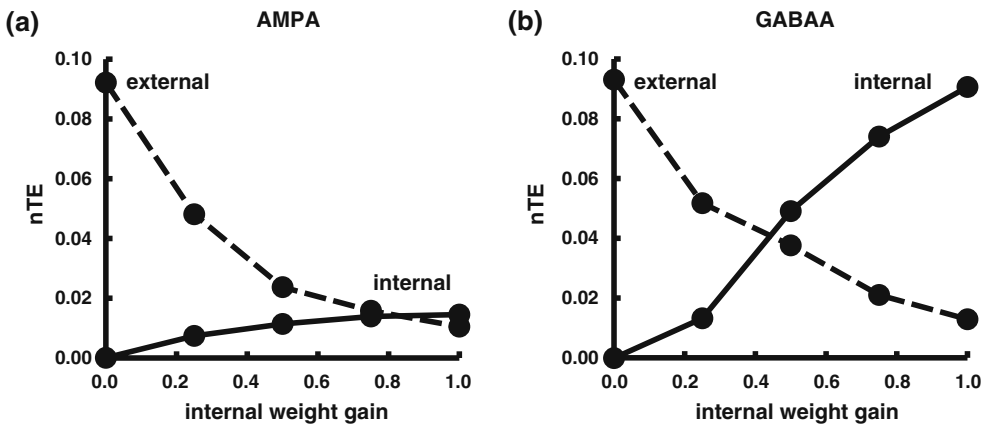

Fig. 5 Information transfer from external (*dashed lines*) and internal (*solid lines*) sources with increasing internal gain. (**a**) Transfer at AMPA synapses onto excitatory cells. (**b**) Transfer at $GABA_A$ synapses onto excitatory cells

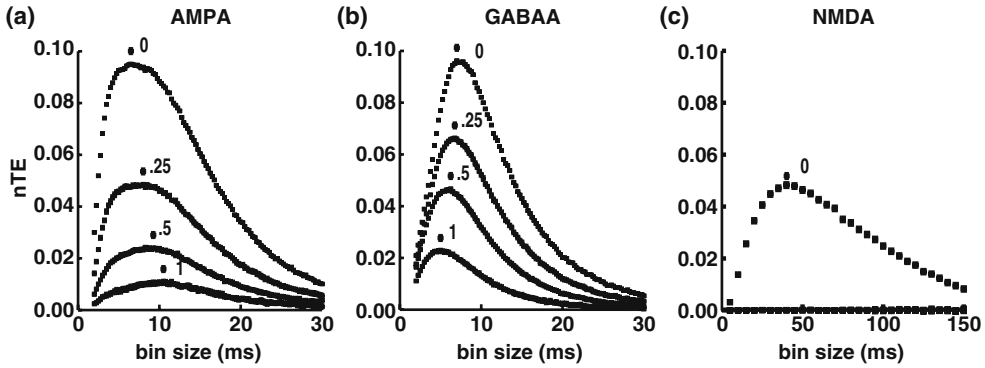

Fig. 6 Average *n*TE as a function of bin size at intrinsic weight gain 0, 0.25, 0.5, 1. Height of shape indicates standard error of the mean from five 100-s simulations with different wiring. *Marker symbols* indicate peaks. (**a**) AMPA synapses on excitatory and inhibitory cells. (**b**) $GABA_A$ synapses on excitatory and inhibitory cells. (**c**) NMDA synapses on excitatory cells only provide *n*TE when internal connectivity strengths are set to 0. Traces where intrinsic gains are non-0 are centered about a *n*TE of 0

and/or internally while avoiding an epileptic network. We found that we could demonstrate information transmission via external NMDA inputs (Fig. 6(c)). As seen in this figure, the information could only be detected at relatively long delays (~40 ms) consistent with the slow peaking and long duration of NMDA current. The data shown in the figure come from a network where the weight of external inputs to NMDA synapses of excitatory cells was increased by a factor of 10 over baseline, matching the strength of external AMPA inputs. In this network, we also reduced the rate of external inputs at NMDA synapses to 10 Hz, because overlap between inputs prevents expression of *n*TE when using higher rates. When internal connectivity was turned on (non-0 traces in Fig. 6(c)), no information flow-through was detected. This was demonstrated over several simulation sets.

The network activation provided significant information flow-through via the fast synapses which overwhelmed any temporally specific effect of NMDA inputs, which in this context become modulatory. Similarly, increasing the synaptic weight of NMDA synapses between excitatory cells did not show significant *n*TE between these cells.

3.3 Changes in local field potentials

We next investigated how macroscopic measure of network activity—local field potential (LFP)—would be expected to relate to network information transfer in detail. In the simulations, internal information transfer via inhibitory $GABA_A$ synapses was of much greater magnitude than via excitatory synapses at high gain, demonstrating that different synapse types produced differential information transfer in the network. Excitatory and inhibitory cells fire at different rates and play different roles in the generation of LFPs of different frequencies (Lytton and Sejnowski 1991; Wang and Buzsaki 1996; Bartos et al. 2007; Tiesinga and Sejnowski 2009). Since information transfer at different

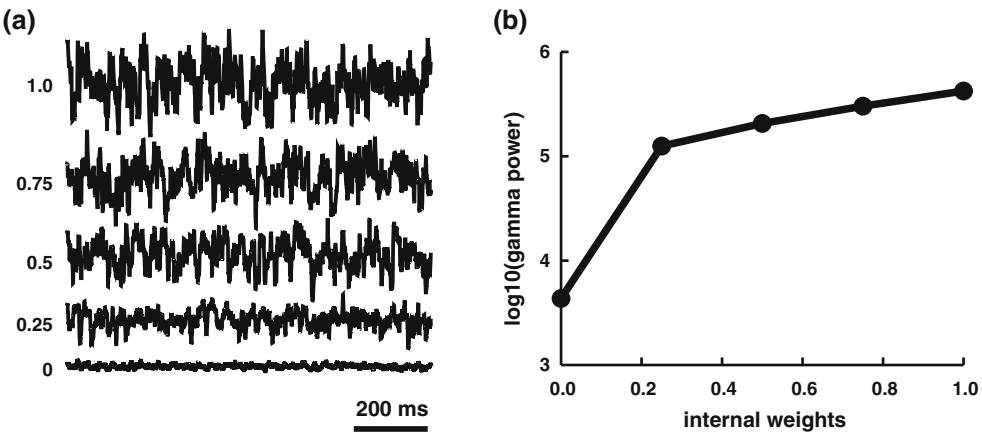

Fig. 7 Simulated local field potential (LFP). (**a**) 1 s of LFP with increasing internal synaptic weights from *bottom* to *top*. (**b**) Total power over 100 s in 40–80 Hz band

types of synapses occurs at different rates, that also implies that LFPs of certain frequencies carry different amounts of information, relating the LFP activity to synaptic activity.

We applied the same increases in internal weights and measured the LFP oscillations of the pyramidal cells (Fig. 7). Dividing the spectrum into bands (delta: 2–4 Hz, theta: 4–8 Hz, alpha: 8–13 Hz, beta: 13–25 Hz, low-gamma: 25–40 Hz, high-gamma: 40–80 Hz) demonstrated monotonic increase in all bands with increasing gain factor. LFP power for high-gamma, which showed the largest increase, is shown in Fig. 7(b). Increases in LFP power are a result of increased cooperative network activity. Since the increases in LFP power were accompanied with increases of information throughput generated within the network, this supports the information processing functions of different bands in LFP or electroencephalogram (Belitski et al. 2008; Mazzoni et al. 2008).

4 Discussion

We used *n*TE to demonstrate that recurrence strength in a simulated neuronal network column gradually reduces information flow from external inputs to outputs of single neurons at both AMPA and $GABA_A$ synapses, and increases information transfer between cells within the network. Information transfer was also reliably demonstrated between spike trains of individual neurons within the network. We found that the information transferred between individual neurons was small, but that multiunit activity transferred more information. This supports the hypothesis that the neural code utilizes populations of neurons. We also note that, although information is generally measured at the single neuron level, separate streams of information can come in at each synapse and have different efficacies.

Measuring *n*TE at different bin sizes demonstrated changes in information transmission delays, with increases for AMPA and decreases for $GABA_A$ inputs with increasing network recurrence. Increasing recurrence was also associated with augmentation in LFP magnitude, most pronounced in the gamma band (40–80 Hz).

In our networks with increased NMDA weights, NMDA information transmission delays were considerably larger than for AMPA and $GABA_A$ synapses, consistent with the longer activation time constants associated with these receptors. The *n*TE at NMDA synapses was negligible except in the case of pure feedforward network with all internal dynamics suppressed. The more immediate AMPA and $GABA_A$ inputs obscured the delayed effects of the NMDA inputs. Similarly, we were not able to demonstrate any information flow via internal NMDA. These findings suggest an NMDA role that primarily modulates activity patterns rather than directly communicating external or stored information.

4.1 Limitations of the model

In designing our model, we specifically set the weight of each external excitatory input to be sub-threshold, so that cells spiked only with temporal integration of inputs. Using a different balance between weights and thresholds would be expected to alter our results. In the extreme, if the weight of each external input had been super-threshold at each AMPA synapse, then the internal weight gains used would have had a significantly lower effect on the *n*TE and correlation between inputs and outputs. However, we expect that our basic result will apply under neocortical conditions of high convergence and low monosynaptic postsynaptic firing efficacy.

Adding randomly-timed noise inputs to neurons would be expected to reduce the *n*TE of external inputs through random disruption of signal throughput. However, the generation of network coordination as measured by increased amplitude LFP (Fig. 7) and increased interneuronal information transfer suggests the development of higher-order network dynamics rather than simply disruption by noise. Additionally, we performed two types of controls to ensure that our results were not simple happenstance of randomness. (1) We used five different input signal sequences and ten different wirings for each gain sequence to control for the possibility that specific variations in external or internal signal, whether in time or within cell groups, produced the alterations in *n*TE. (2) We performed controls where increasing external weight gain paralleled increasing internal weight gain. In these simulations the internal dynamics still dominated *n*TE despite increased drive from the inputs, suggesting that the effect was associated with internal dynamics rather than with a simple increased noise effect.

In this paper, we utilized random external inputs as a stand-in for a signal and random-wiring as a stand-in for learned connectivity. We then varied gain uniformly. In the RPD networks, we did single synapse depression and potentiation to more closely match the alterations that would occur with use of a learning algorithm. In future research, we will explore patterned input signals with specific learning algorithms to confirm that these results hold in the context of purposely structured

dynamics. We expect the *n*TE effects to be still more pronounced with greater internal dynamical structure.

4.2 Network access of stored information

The data processing inequality theorem of information theory states that no more information can be obtained out of a set of data then was there to begin with (McDonnell et al. 2003; Quiroga and Panzeri 2009). However, this theorem applies to Markov chains, where the data processing is done through a historyless point process. Clearly, one can add one information stream to another and in the end obtain more information than either single stream contained. Similarly, neuronal networks can contribute streams of stored information which can then be combined with information from external stimuli. *n*TE is an ideal tool for quantifying the information stored in neuronal networks. It measures the direction of information transfer, distinguishing senders and receivers, and distinguishing between external and internal/intrinsic sources of information. Thus, *n*TE quantified the stored information added at network synapses.

Accessing information is analogous in both computers and in the brain; a simple low-information signal can be used to access complex stored information for output from the system. In computer systems, information can be accessed from a database using a simple key whose entropic distance from other keys is minimal. In such a database system, the user makes a query with key *key*, which does not itself contain information, but simply provides access to information. In brain science, a classic comparable example is Penfield's electrical stimulation of medial temporal lobe in patients undergoing surgery (Halgren et al. 1978; Penfield 1958). A simple stimulation, here characterized primarily by its location, evokes an outpouring of stored information. Similarly, in our simulations, the external stimuli enabled extraction of information stored within the network, measured as increased inter-neuronal information transfer. In all these examples, with a simple stimulus, the information stored within the given medium can be extracted.

4.3 Use of *n*TE for simulation validation

A common problem in simulations of neuronal networks is difficulty in generating and sustaining intrinsic activity in the network (Lytton et al. 2008; Vogels et al. 2005). This is particularly difficult if one wishes to replicate the relatively slow firing rates of cortical neurons—many excitatory neurons fire at 1 Hz or less when not specifically stimulated. Only with extremely large networks is it possible to get realistic firing rates with activity sustained by "circus rhythms," activity circulating through different circuits with different loop delays (Izhikevich and Edelman 2008).

More typically, activity is artificially sustained by providing continuous random input to the system. This artifice raises the question as to whether the network, the object of research interest, is actually producing the activity being measured, or simply parroting the pseudo-random activity that sustains it. By measuring the information transferred between cells within the network, *n*TE can be used to demonstrate whether or not the network is producing "added value" (added intrinsic activity) in this situation.

Simply measuring the information transferred from external inputs to outputs is not sufficient, because adding noise inputs to the neurons will reduce this *n*TE and not reflect anything about activity intrinsic to the network. By contrast, measurement of *n*TE from presynaptic to postsynaptic cell can distinguish intrinsic network activity from added noise. Intrinsic network activity will increase internal information transfer while synaptic noise will decrease it. Measuring significant information transferred between cells within the network serves as a marker for intrinsic activity, and can be used as a validation step in simulations.

4.4 Excitatory vs inhibitory synapses as information carriers

Excitatory and inhibitory synapses showed different efficacies in information transfer, particularly evident with internal transfer (Fig. 5) but also consistently seen as a small difference for the external inputs. This $GABA_A$-mediated increase is associated with an increase in LFP power (Fig. 7), increasing most in the high gamma band (40–80 Hz), which is also largely driven and determined by network-intrinsic $GABA_A$ connectivity (Lytton and Sejnowski 1991; Wang and Buzsaki 1996; Tiesinga and Sejnowski 2009; Börgers and Kopell 2005; Traub et al. 1999). The observation of $GABA_A$ as a particularly good carrier of information makes a connection between inhibitory synaptic inputs as a carrier of information and driver of oscillations, and suggests how oscillation in the associated gamma band might serve as a surrogate for information carriage (Gray and Singer 1989; Von der Malsburg and Schneider 1986). Because neurons are highly sensitive to the correlation structure of inputs (Salinas and Sejnowski 2002), it might further be expected

that drivers of correlated oscillatory inputs would provide larger information throughput through the network than would drivers that function independently. Experimental work has shown that gamma and theta oscillations interact to allow for enhanced information transfer between different areas of the brain (Sirota et al. 2008).

4.5 Implications for artificial networks

The results reported here could likely be replicated using a simple integrate-and-fire model of excitatory and inhibitory cells with a single-layer random-wiring architecture (Brunel 2000; Brunel and Wang 2003). This would have advantages in terms of transparency for further analysis. Instead, we utilized a neocortical column network with moderately complex neuron models and moderately detailed connectivity. Our motivation here is to prepare for additional studies where we will analyze information flow across layers and between cell types in the network to better understand the column as a computational unit of neocortex. It will be valuable to explore the current canonical column pathway, layer 4 $\rightarrow$ layer 2/3 $\rightarrow$ layer 5, in terms of information flow.

Our present results can, however, be viewed in the context of artificial neural networks (ANNs) and machine learning. Our control network (0 gain) is a feedforward network with the architecture of the classic single-layer perceptron network. Information is effectively summed (here nonlinearly unlike in a perceptron or a back-propagation algorithm unit) through the single units and then passed through a nonlinearity to produce an output. In such a feedforward network, unit outputs contain considerable information derived directly from external inputs (Fig. 2(b)).

Our connected networks represent recurrent architectures, more complex than most ANN architectures, since subpopulations are distinguishable not only by connectivity but also by having distinct dynamical personalities. We view our lightly-connected (gain 0.25) network as analogous to an ANN before learning—weights are low and random. This would be a condition where the network contains relatively little information. Therefore information flow-through is only slightly impacted. In this context, our gain 0.5 network might correspond to a network after learning, where stronger interconnects have formed to store information. Although we have not yet assessed learning in our networks, random alterations (random decreases and increases producing net increase) in internal synaptic strength produced the same results as we found with uniform increase in synaptic gain.

The final network we assessed (gain 1.0) is an overweighted network that showed little activity flow from external sources. From the network learning perspective, this would correspond to a saturated network where continuous potentiation has resulted in overlearning (saturation) which blocks further learning and can cause catastrophic forgetting (French 1991; McCloskey and Cohen 1989; Moser and Moser 1999; Ratcliff 1990).

4.6 Relation to echo state networks

Recent work on echo state networks (ESNs, Jaeger and Haas 2004; Buonomano and Maass 2009; Buonomano 2009) showed that these networks do extremely well with time-series prediction. ESNs consist of a pool of recurrently connected neurons (the liquid state) that have randomly assigned weights between them. The many possible dynamical subnetworks within the overall network provide a reservoir from which particular dynamical patterns can be selected (Edelman 1987). In the ESN model, only the connections from network to readout are trained, there is no internal weight learning. Viewing our network in this context, increasing internal weight gain (recurrence strength) produced networks with increasing degrees of internal dynamical richness, which was then measured by *n*TE (Fig. 5). With increased dynamical richness, the output increasingly reflects internal dynamics and recedes from external signal dynamics. Thus, in dynamically rich networks, the external signal primarily serves as a selector signal.

Viewing the neocortical column as an ESN suggests a number of interesting speculations. Insofar as neocortex offers a number of different through-paths, the column could be a set of interacting ESNs that can modify each other's dynamics. Rather than being purely inductive or purely selective, neocortex may well be both, allowing initial selection to greatly reduce the training complexity for a cortical column born with an initial pool of useful recurrent connectivity.

When intrinsic activity is not present, the network is a feedforward network. With recurrence added, the increasing proportion of internal information transmitted via inhibitory inputs might provide a balance between primarily excitatory external information and GABAergic internal information. When intrinsic activity is low, external information passes through. When intrinsic activity is high, external information serves as a selector. The increasing information delay associated

with excitatory inputs (Fig. 6) is consistent with this selector hypothesis.

4.7 Transmission, translation, corruption

We capture the spectrum of recurrence in networks in a descriptive schema: transformation, translation, and corruption of inputs:

$$\text{feedforward} : \textit{input} \xrightarrow{\text{transfer}} \textit{net} \xrightarrow{\text{transform}} \textit{output}$$

$$\text{recurrent} : \textit{input} \xrightarrow{\text{transfer}} \textit{net} \xrightarrow{\text{translate}} \textit{output}$$

$$\text{pathological} : \textit{input} \xrightarrow{\text{transfer}} \textit{net} \xrightarrow{\text{corrupt}} \textit{output}$$

The spectrum shows increasing intrinsic activity of the network with concomitant reduction in external information flow from input to output. Transformation networks have loosely or uncoupled elements. Each individual neuron in the network is not influenced by the other neurons. At this end of the spectrum, each neuron fires asynchronously, and does not influence other neurons. In these networks, n external inputs convergent on the network are thresholded, and produce m transformed outputs (Lytton 1998). This type of network can nominally transmit more information from external inputs to outputs, since there is a larger ensemble of possible network firing states. However, without any interaction between neurons in the network, the output can only be nominally considered a network effect.

Translation networks have higher coupling between the constituent elements. This allows for the external inputs to access information in the weights stored between neurons in the network, adding it to the external input signal. The accessed information is made apparent dynamically via the output activity of the neurons and increased information flow inter-neuronally. Intrinsic weights in translation networks can subserve different purposes, including enhancing a stimulus or feature extraction. For example, topographic input may be sent to a column in visual cortex. The neurons in the network work together, translating the inputs to a decision on what is seen.

In a corrupting network, the network dominates and external inputs are relatively ignored. Such networks have been used to model "inattentional blindness", the physiological phenomenon that causes sensory stimuli to be missed when a subject is concentrating on a mental task (Dehaene and Changeux 2005). Pathologically, a greater degree of network independence or network isolation could be associated with brain disease.

4.8 Information flow and brain disease

We expect that the gradual reduction in information flow-through with increased internal weight strength will be asymptotic: at the most extreme, the highly interconnected network will be epileptic and any input will ignite a seizure that is unrelated to the input (Uhlrich et al. 2005). At a lower level of dynamic pathology, abnormalities in power and coordination in gamma bands might be seen, as has been noted in schizophrenia (Uhlhaas et al. 2006; Spencer et al. 2004, 2003; Uhlhaas and Singer 2006). A network with a pathologically high level of coherence (reduced complexity) will no longer accept external information and will no longer coordinate well with other brain areas. Information from either the outside world or other brain areas would then be drowned out. In schizophrenia, the core cognitive deficit is believed to be a problem of coordinating external and internal states to provide a coherent view of the world, and the person's relation to the world (Phillips and Silverstein 2003). This is most evident with hallucinations, a mistaken attribution of internal activity to the external world. Our simulations provide a framework for capturing aspects of such disorders by quantifying the extent to which external information may be corrupted, leading to an over-reliance on information stored internally in a network.

Acknowledgements We thank the anonymous reviewers for their insightful comments and suggestions on improvement of the manuscript.

We also thank Michael Hines and Ted Carnevale (Yale) for continuing support and assistance with the NEURON simulator, Andrey Olypher (Emory) for helpful discussions, Matthew Lazenka (VCU) for assistance with development of the neocortical model, and Larry Eberle (SUNY Downstate) for administration and network support at the Neurosimulation laboratory.

References

Aldworth, Z., Miller, J., Gedeon, T., Cummins, G., & Dimitrov, A. (2005). Dejittered spike-conditioned stimulus waveforms yield improved estimates of neuronal feature selectivity and spike-timing precision of sensory interneurons. *Journal of Neuroscience, 25*(22), 5323–5332.

Bartos, M., Vida, I., & Jonas, P. (2007). Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews. Neuroscience, 8*(1), 45–56.

Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M., Logothetis, N., et al. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *Journal of Neuroscience, 28*(22), 5696–5709.

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

Brunel, N. (2000). Dynamics of networks of randomly connected excitatory and inhibitory spiking neurons. *Journal of Physiology (Paris), 94*, 445–463.

Brunel, N., & Wang, X. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. *Journal of Neurophysiology, 90*, 415–430.

Buonomano, D. (2009). Harnessing chaos in recurrent neural networks. *Neuron, 63*, 423–425.

Buonomano, D., & Maass, W. (2009). State-dependent computations: Spatiotemporal processing in cortical networks. *Nature Reviews. Neuroscience, 10*(2), 113–125.

Carnevale, N., & Hines, M. (2006). *The NEURON book*. Cambridge: Cambridge University Press.

Dehaene, S., & Changeux, J. (2005). Ongoing spontaneous activity controls access to consciousness: A neuronal model for inattentional blindness. *PLoS Biology, 3*(5), e141.

Destexhe, A., & Contreras, D. (2006). Neuronal computations with stochastic network states. *Science, 314*(5796), 85–90.

Destexhe, A., Mainen, Z., & Sejnowski, T. (1994). An efficient method for computing synaptic conductances based on a kinetic model of receptor binding. *Neural Computation, 6*, 14–18.

Douglas, R., Martin, K., & Whitteridge, D. (1989). A canonical microcircuit for neocortex. *Neural Computation, 1*, 480–488.

Edelman, G. (1987). *Neural Darwinism: The theory of neuronal group selection*. New York: Basic Books.

French, R. (1991). Using semi-distributed representations to overcome catastrophic forgetting in connectionist networks. In *Proceedings of the 13th annual cognitive science society conference* (pp. 173–178). Hillsdale: Erlbaum.

Friesen, W., & Friesen, J. (1994). *NeuroDynamix, a computer-based system for simulating neuronal properties*. New York: Oxford Univ. Press.

Gourevitch, B., & Eggermont, J. (2007). Evaluating information transfer between auditory cortical neurons. *Journal of Neurophysiology, 97*(3), 2533–2543.

Gray, C., & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences of the United States of America, 86*, 1698–1702.

Halgren, E., Walter, R., Cherlow, D., & Crandall, P. (1978). Mental phenomena evoked by electrical stimulation of the human hippocampal formation and amygdala. *Brain, 101*(1), 83.

Hill, S., & Tononi, G. (2005). Modeling sleep and wakefulness in the thalamocortical system. *Journal of Neurophysiology, 93*, 1671–1698.

Hines, M., & Carnevale, N. (2001). NEURON: A tool for neuroscientists. *The Neuroscientist, 7*, 123–135.

Hlaváckova-Schindler, K., Palus, M., Vejmelka, M., & Bhattacharya, J. (2007). Causality detection based on information-theoretic approaches in time series analysis. *Physics Reports, 441*, 1–46.

Izhikevich, E., & Edelman, G. (2008). Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences of the United States of America, 105*(9), 3593–3598.

Jaeger, H., & Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science, 304*, 78–80.

Jumarie, G. (1990). *Relative information: Theories and applications*. New York: Springer.

Kendall, M. (1938). A new measure of rank correlation. *Biometrika, 30*(1–2), 81–93.

Knight, W. (1966). A computer method for calculating Kendall's tau with ungrouped data. *Journal of the American Statistical Association, 61*(314), 436–439.

Lazar, A., & Pnevmatikakis, E. (2008). Faithful representation of stimuli with a population of integrate-and-fire neurons. *Neural Computation, 20*(11), 2715–2744.

Lytton, W. (1996). Optimizing synaptic conductance calculation for network simulations. *Neural Computation, 8*, 501–510.

Lytton, W. (1998). Adapting a feedforward heteroassociative network to Hodgkin–Huxley dynamics. *Journal of Computational Neuroscience, 5*, 353–364.

Lytton, W. (2006). Neural query system: Data-mining from within the NEURON simulator. *Neuroinformatics, 4*, 163–176.

Lytton, W., & Omurtag, A. (2007). Tonic-clonic transitions in computer simulation. *Journal of Clinical Neurophysiology, 24*, 175–181.

Lytton, W., & Sejnowski, T. (1991). Inhibitory interneurons may help synchronize oscillations in cortical pyramidal neurons. *Journal of Neurophysiology, 66*, 1059–1079.

Lytton, W., & Stewart, M. (2007). Data mining through simulation. *Methods in Molecular Biology, 401*, 155–166.

Lytton, W., Neymotin, S., & Hines, M. (2008). The virtual slice setup. *Journal of Neuroscience Methods, 171*, 309–315.

Marschinski, R., & Kantz, H. (2002). Analysing the information flow between financial time series. *European Physical Journal B, Condensed Matter Physics, 30*(2), 275–281.

Mazzoni, A., Panzeri, S., Logothetis, N., & Brunel, N. (2008). Encoding of naturalistic stimuli by local field potential spectra in networks of excitatory and inhibitory neurons. *PLoS Computational Biology, 4*(12), e1000239.

McCloskey, M., & Cohen, N. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. Chapter in *The psychology of learning and motivation: Advances in research and theory* (Vol. 24). Maryland Heights: Academic.

McDonnell, M., Stocks, N., Pearce, C., & Abbott, D. (2003). Stochastic resonance and data processing inequality. *Electronics Letters (IEEE), 39*(17), 1287–1288.

Moser, E., & Moser, M. (1999). Is learning blocked by saturation of synaptic weights in the hippocampus. *Neuroscience and Biobehavioral Reviews, 23*, 661–672.

Nelson, S. (2002). Cortical microcircuits: Diverse or canonical. *Neuron, 36*, 19–27.

Paluš, M. (1996). Detecting nonlinearity in multivariate time series. *Physics Letters A, 213*(3–4), 138–147.

Penfield, W. (1958). Some mechanisms of consciousness discovered during electrical stimulation of the brain. *Proceedings of the National Academy of Sciences of the United States of America, 44*(2), 51–66.

Phillips, W., & Silverstein, S. (2003). Convergence of biological and psychological perspectives on cognitive coordination in schizophrenia. *Behavioral and Brain Sciences, 26*(01), 65–82.

Press, W., Teukolsky, S., Vetterling, W., & Flannery, B. (2007). *Numerical recipes: The art of scientific computing*. Cambridge: Cambridge University Press.

Quiroga, R., & Panzeri, S. (2009). Extracting information from neuronal populations: Information theory and decoding approaches. *Nature Reviews Neuroscience, 10*(3), 173–185.

Rao, R., & Sejnowski, T. (2001). Spike-timing-dependent Hebbian plasticity as temporal difference learning. *Neural Computation, 13*(10), 2221–2237.

Ratcliff, R. (1990). Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review, 97*(2), 285–308.

Rieke, F., Warland, D., & Bialek, W. (1999). *Spikes: Exploring the neural code*. Cambridge: MIT.

Salinas, E., & Sejnowski, T. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience, 2*(8), 539–550.

Salinas, E., & Sejnowski, T. (2002). Integrate-and-fire neurons driven by correlated stochastic input. *Neural Computation, 14*(9), 2111–2155.

Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters, 85*(2), 461–464.

Sirota, A., Montgomery, S., Fujisawa, S., Isomura, Y., Zugaro, M., & Buzsáki, G. (2008). Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron, 60*, 683–697.

Spencer, K., Nestor, P., Niznikiewicz, M., Salisbury, D., Shenton, M., & Carley, R. (2003). Abnormal neural synchrony in schizophrenia. *Journal of Neuroscience, 23*, 7407–7411.

Spencer, K., Nestor, P., Perlmutter, R., Niznikiewicz, M., Klump, M., Frumin, M., et al. (2004) Neural synchrony indexes disordered perception and cognition in schizophrenia. *Proceedings of the National Academy of Sciences of the United States of America, 101*, 17288–17293.

Sporns, O., Tononi, G., & Kotter, R. (2005). The human connectome: A structural description of the human brain. *PLoS Computational Biology, 1*(4), e42.

Tiesinga, P., & Sejnowski, T. (2009). Cortical enlightenment: Are attentional gamma oscillations driven by ing or ping? *Neuron, 63*(6), 727–732.

Traub, R., Jefferys, J., & Whittington, M. (1999). *Fast oscillations in cortical circuits*. Cambridge: MIT.

Uhlhaas, P., & Singer, W. (2006). Neural synchrony in brain disorders: Relevance for cognitive dysfunctions and pathophysiology. *Neuron, 52*, 155–168.

Uhlhaas, P., Linden, D., Singer, W., Haenschel, C., Lindner, M., Maurer, K., et al. (2006). Dysfunctional long-range coordination of neural activity during gestalt perception in schizophrenia. *Journal of Neuroscience, 26*, 8168–8175.

Uhlrich, D., Manning, K., Laughlin, M., & Lytton, W. (2005). Photic-induced sensitization: Acquisition of an augmenting spike-wave response in the adult rat through repeated strobe exposure. *Journal of Neurophysiology, 94*, 3925–3937.

Victor, J. (2006). Approaches to information-theoretic analysis of neural activity. *Biological Theory, 1*(3), 302–316.

Vogels, T., Rajan, K., & Abbott, L. (2005). Neural network dynamics. *Annual Review of Neuroscience, 28*, 357–376.

Von der Malsburg, C., & Schneider, W. (1986). A neural cocktail-party processor. *Biological Cybernetics, 54*, 29–40.

Wang, X., & Buzsaki, G. (1996). Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *Journal of Neuroscience, 16*, 6402–6413.

Zhu, J., Lytton, W., Xue, J., & Uhlrich, D. (1999a). An intrinsic oscillation in interneurons of the rat lateral geniculate nucleus. *Journal of Neurophysiology, 81*, 702–711.

Zhu, J., Uhlrich, D., & Lytton, W. (1999b). Burst firing in identified interneurons of the rat lateral geniculate nucleus. *Neuroscience, 91*, 1445–1460.