The A-current and Type I / Type II transition determine collective spiking from common input

Andrea K. Barreiro, Evan L. Thilo, Eric Shea-Brown[*]

Department of Applied Mathematics and

Program in Neurobiology and Behavior

University of Washington,

Box 352420, Seattle, WA 98195

[*] *Correspondence: etsb@uw.edu*

1

ABSTRACT

The mechanisms and impact of correlated, or synchronous, firing among pairs and groups of neurons is under intense investigation throughout the nervous system. A ubiquitous circuit feature that can give rise to such correlations consists of overlapping, or common, inputs to pairs and populations of cells, leading to common spike train responses. Here, we use computational tools to study how the transfer of common input currents into common spike outputs is modulated by the physiology of the recipient cells. We focus on a key conductance — g_A, for the A-type potassium current — which drives neurons between "Type II" excitability (low g_A), and "Type I" excitability (high g_A).

Regardless of g_A, cells transform common input fluctuations into a tendency to spike nearly simultaneously. However, this process is more pronounced at low g_A values, as previously predicted by reduced "phase" models. Thus, for a given level of common input, Type II neurons produce spikes that are relatively *more* correlated over short time scales. Over long time scales, the trend reverses, with Type II neurons producing relatively *less* correlated spike trains. This is because these cells' increased tendency for simultaneous spiking is balanced by opposing tendencies at larger time lags. We demonstrate a novel implication for neural signal processing: downstream cells with long time constants are selectively driven by Type I cell populations upstream, and those with short time constants by Type II cell populations. Our results are established via high-throughput numerical simulations, and explained via the cells' filtering properties and nonlinear dynamics.

INTRODUCTION

Neurons throughout the nervous system — from the retina (Shlens et al., 2008), thalamus (e.g., (Alonso et al., 1996)), and cortex (e.g., (Zohary et al., 1994)) to motoneurons (Binder and Powers, 2001) — show temporal correlation between the discharge times of their spikes. This correlated spiking can impact sensory discrimination (Averbeck et al., 2006) and signal propagation (Salinas and Sejnowski, 2000).

How do these correlations arise? We study a simple mechanism in which the inputs to a pair or population of neurons has a common component that is shared across multiple cells (Figure 1). On an anatomical level, the large number of divergent connections that span layers and areas makes shared
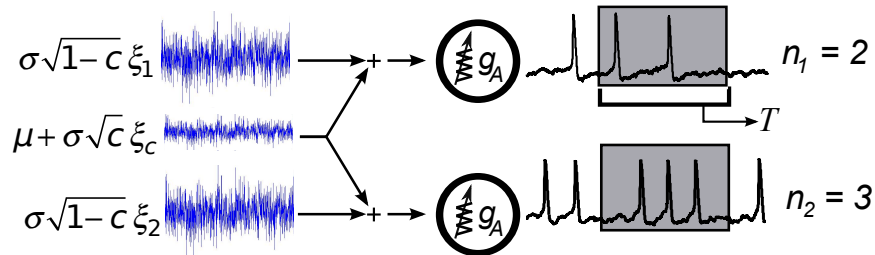
Fig. 1: Shared input microcircuit, in which two neurons receive input currents with a common component that represents correlated activity or shared afferents upstream. Each neuron is a single-compartment Connor-Stevens model (see Methods), with a maximal A-current conductance g_A that we vary, eliciting a full range of Type I to Type II spiking dynamics. Shared input currents lead to correlated spikes, which are quantified as shown via spike counts n_1, n_2 over sliding time windows of length T. The input currents received by each cell have mean μ and fluctuate with total variance σ^2; the common noise is chosen with variance $c\sigma^2$, and independent noise terms with variance $(1-c)\sigma^2$.

afferents to pairs of nearby cells unavoidable (Shadlen and Newsome, 1998). Correlated spiking in areas upstream from the target cells can add to this anatomical factor. In fact, for some neural circuits, shared inputs are themselves the dominant source of correlated spiking (Trong and Rieke, 2008). In general, correlating effects of shared input interact with effects of recurrent coupling (cf. (Ostojic et al., 2009)).

What makes shared input circuitry especially interesting is the pivotal role of spike generating dynamics. For a given fraction of shared input, these dynamics control the fraction of shared output — that is, the fraction of spikes that will be shared across the two cells. This *correlation transfer* depends on two factors. The first is the mechanism of spike generation. The second is the operating point of the neurons (i.e., their rate and variability of firing (Binder and Powers, 2001; de la Rocha et al., 2007)). Excepting Hong and De Schutter (2008), studies of correlation transfer have mostly focused on simplified neuron models, such as integrate-and-fire, phase, or threshold crossing systems, leaving open allied questions for models with more complex

subthreshold and after-spike dynamics.

Here, we study correlation transfer for a family of conductance-based neuron models that varies from Type I excitability, in which firing can occur at arbitrarily low rates in response to a DC current (as for cortical pyramidal cells), to Type II excitability, in which firing occurs at a nonzero "onset" rate (as for fast-spiking interneurons or the Hodgkin-Huxley model) (Rinzel and Ermentrout, 1998; Hodgkin, 1948). We use the Connor-Stevens model (Connor and Stevens, 1971), which transitions between Type I and Type II as g_A – the maximal conductance of the A-type potassium current – is varied (see Fig. 2). Beyond firing rates, Type I vs. Type II neurons differ in single-cell computation (Ermentrout et al., 2007) and synchronization under reciprocal coupling (Rinzel and Ermentrout, 1998).

We test the hypothesis — based on predictions from simplified "normal form" phase models (Barreiro et al., 2010; Marella and Ermentrout, 2008; Galán et al., 2007) — that the Type I to Type II transition will produce strong differences in levels and time scales of correlation transfer. Upon finding a positive result, and explaining it via the filtering properties of individual neurons, we ask how the distinct features of correlated spiking in Type I vs. Type II neurons manifest in signal transmission through feedforward neural circuits. Preliminary versions of some findings have appeared in abstract form (Barreiro, 2009; Shea-Brown et al., 2009).

METHODS

Circuit setup

We primarily consider the feedforward circuit of Figure 1. Here, each of two neurons receives two sources of fluctuating current: a common, or "shared" source $I_c = \sigma\sqrt{c}\,\xi_c(t)$, and a "private" source $I_1 = \sigma\sqrt{1-c}\,\xi_1(t)$ or $I_2 = \sigma\sqrt{1-c}\,\xi_2(t)$. Each of these inputs is chosen to be a scaled statistically independent, Gaussian white noise process (uncorrelated in time) — that is, $\langle \xi_i(t)\xi_i(t+\tau)\rangle = \delta(\tau)$ for $i = 1, 2, c$. This is for simplicity and agreement with prior studies of correlated spiking (Lindner et al., 2005; de la Rocha et al., 2007; Shea-Brown et al., 2008; Marella and Ermentrout, 2008; Vilela and Lindner, 2009; Barreiro et al., 2010). The common current I_c has variance $\sigma^2 c$; each private current has zero mean and variance $\sigma^2(1-c)$. Note that these scalings are chosen so that the total variance of current injected into

each cell is always σ^2, while the parameter c gives the fraction of this variance that arises from common input sources. For example, when $c = 0.5$, 50% of each neuron's presynaptic inputs come from the shared and 50% from the independent input. Finally, the mean of the total current received by each cell is given by μ. This term represents the total bias toward negative or positive currents from all sources; in Fig. 1, it is illustrated as part of the common input for simplicity.

The combined currents,

$$
\begin{aligned}
I_{app,i}(t) &= \mu + I_c(t) + I_i(t) \\
&= \mu + \sigma\sqrt{c}\,\xi_c(t) + \sigma\sqrt{1-c}\,\xi_i(t)
\end{aligned}
$$

($i = 1, 2$) are injected into identical single-compartment, conductance-based membrane models (see Methods, "Neuron model"); spike times are identified from the resulting voltage trace.

Neuron model

We investigate correlation transfer in the Connor-Stevens model, which was designed to capture the low firing rates of a crab motor axon (Connor and Stevens, 1971; Connor et al., 1977). This model adds a transient potassium current, or A-current, to sodium and delayed-rectifier potassium currents of Hodgkin-Huxley type. The A-Type channel provides extended after-spike hyperpolarization currents, which lead to arbitrarily low firing rates and hence Type I excitability (see Introduction and Methods, "Characterizing the dynamics of spike generation," below).

The voltage equation is

$$
\begin{aligned}
C_M \frac{dV}{dt} &= -g_L(V - E_L) - g_{Na}m^3h(V - E_{Na}) - g_K n^4(V - E_K) \quad (1) \\
&- g_A A^3 B(V - E_K) + I_{app},
\end{aligned}
$$

The gating variables m, n, h, A, and B each evolve according to the standard voltage-gated kinetics; e.g., for m:

$$
\frac{dm}{dt} = \frac{m_\infty(V) - m}{\tau_m(V)} \quad (2)
$$

where $m_\infty(V)$ is the steady-state value and $\tau_m(V)$ is the (voltage-dependent) time constant. All equations and parameters are exactly as specified as in
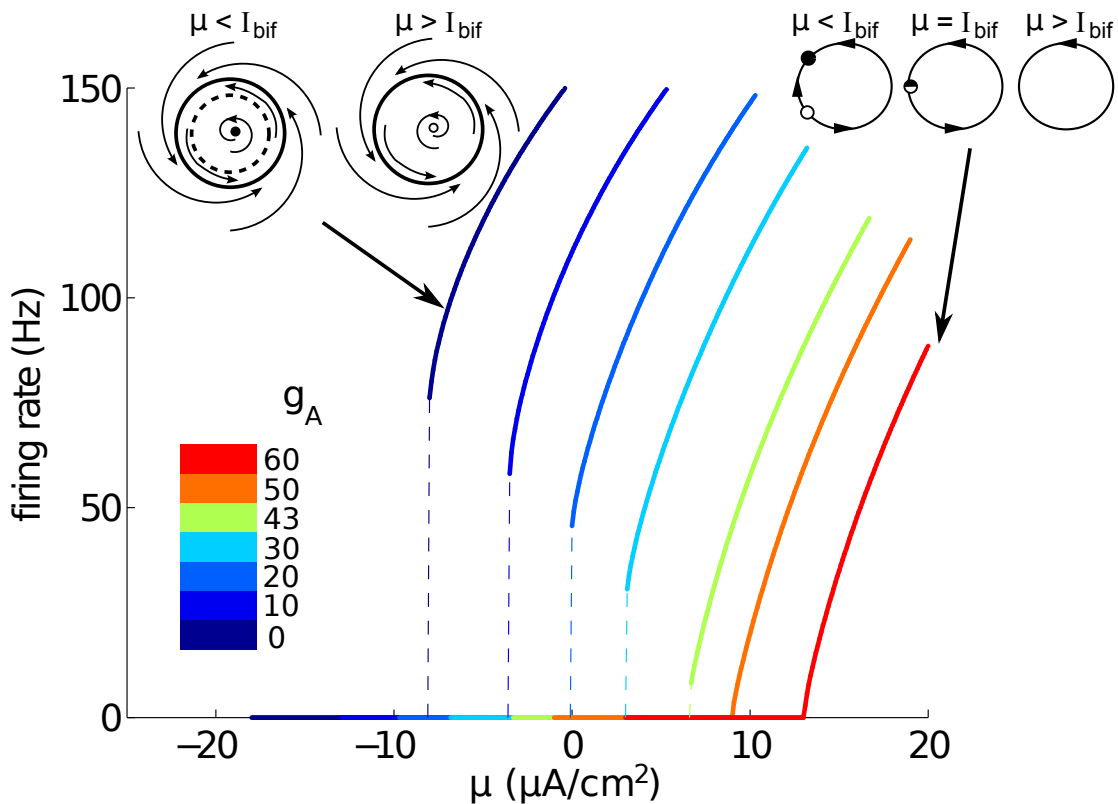
Fig. 2: Firing rate vs. injected current $(f - I)$ curves, for the deterministic $(\sigma = 0)$ Connor-Stevens model. Several values of g_A, yielding a range from Type II to Type I excitability, are shown — note the nonzero "onset" firing rates and Type II excitability for $g_A \approx 0$ mS/cm^2, zero onset rate and Type I excitability for $g_A \approx 60$ mS/cm^2, and a gradual transition between. Insets show cartoons of dynamical transitions that lead to non-zero vs. zero onset rates: a subcritical Hopf bifurcation (left) and a saddle-node on invariant circle bifurcation (right).

Connor et al. (1977), with the exception that we vary the maximal A-current conductance, g_A, over the range of values reported below. As g_A is decreased from the value set by Connor et al. (1977), the neuron transitions from Type I to Type II excitability; we describe this phenomenon next.

Measuring spike train correlation

We represent the output spike trains as sequences of impulses $y_i(t) = \sum_i \delta(t - t_i^k)$, where t_i^k is the time of the kth spike of the ith neuron. The firing rate of the ith cell, $\langle y_i(t) \rangle$, is denoted ν_i. To quantify correlation over a given time scale T, we compute the Pearson's correlation coefficient of spike counts over a time window of length T (as in, e.g., (Zohary et al., 1994; Bair et al., 2001)):

$$\rho_T = \frac{\text{Cov}(n_1, n_2)}{\sqrt{\text{Var}(n_1)}\sqrt{\text{Var}(n_2)}}$$

where n_1, n_2 are the numbers of spikes simultaneously output by neurons 1 and 2 respectively in a time window of length T; i.e. $n_i(t) = \int_t^{t+T} y_i(s)\,ds$. If ρ_T is measured at values of T that are less than a typical inter-spike interval (ISI), we are essentially measuring the degree of synchrony between individual spikes. For larger T values, ρ_T assesses total correlation between numbers of spikes emitted by each cell.

A short calculation (cf. (Bair et al., 2001; Cox and Lewis, 1966)) shows that $\text{Cov}(n_1, n_2)$ is

$$\text{Cov}(n_1, n_2) = T \int_{-T}^{T} C_{12}(t) \frac{T - |t|}{T}\,dt \qquad (3)$$

where the spike train cross-covariance $C_{12}(\tau) = \langle y_1(t) y_2(t + \tau) \rangle - \nu_1 \nu_2$. Similarly, the variance $\text{Var}(n_1)$ can be given in terms of the spike train auto-covariance function. The autocovariance function of neuron 1, defined as $A_1(\tau) = \langle y_1(t) y_1(t + \tau) \rangle - \nu_1^2$, satisfies

$$\text{Var}(n_1) = T \int_{-T}^{T} A_1(t) \frac{T - |t|}{T}\,dt \,,$$

and similarly for neuron 2.

Characterizing the dynamics of spike generation

We give a brief discussion of spike-generation dynamics in the Connor-Stevens system. We organize this by focusing on how the neurons transition from quiescent (i.e., rest) behavior to periodic spiking in response to constant, noiseless (DC) current (i.e., when fluctuation terms $\sigma = 0$). While these

results are well-established (Rush and Rinzel, 1995) we review critical aspects that will help us understand how correlated spiking arises in the circuit of Fig. 1.

For any value of g_A, there is a critical value of the DC current, $\mu = I_{bif}(g_A)$, such that the neuron has a stable stationary state — that is, it remains at a resting voltage — if $\mu < I_{bif}(g_A)$. If $\mu > I_{bif}(g_A)$, then the neuron instead spikes periodically. We refer to these two regimes as *subthreshold* and *superthreshold* respectively. If the current μ is ramped slowly from a subthreshold value, the membrane potential will shift slowly until $\mu = I_{bif}(g_A)$, at which point periodic spikes begin; the rate continues to increase with μ over some range, as quantified by familiar firing rate-current ($f - I$) curves (Fig. 2).

As discussed above, neurons are often classified as Type I vs. Type II based on whether their $f - I$ curves are continuous vs. discontinuous (with a jump) at $\mu = I_{bif}$. Figure 2 shows that the Connor-Stevens model is Type II for $g_A \approx 0$ mS/cm^2, Type I for $g_A \approx 60$ mS/cm^2, and displays a gradual transition in between. This qualitative shift in behavior is related to the underlying model (Equations 1,2) through *bifurcation theory* (Izhikevich, 2007; Rinzel and Ermentrout, 1998). The key fact is that all possible changes from resting to periodic spiking behavior fall into a few qualitatively equivalent categories of bifurcation. We next describe two of these categories that we will put to use below.

In a *saddle-node on invariant circle* (SNIC) bifurcation, the stable resting state merges with an *unstable* resting state to create a single steady state precisely when $\mu = I_{bif}(g_A)$. A solution to Equations (1,2) that begins near this steady state can return to it, via an excursion through state space that traces out a spike. This spiking trajectory persists for $\mu > I_{bif}$. However, due to the steady states at nearby values of μ, spiking trajectories take a long time to return to where they started — technically an infinite amount of time when $\mu = I_{bif}(g_A)$; see the righthand schematic in Figure 2. As a result, the firing rate for μ near $I_{bif}(g_A)$ is arbitrarily low, resulting in a continuous $f - I$ curve and Type I excitability.

In a *subcritical Hopf* bifurcation, by contrast, the stable resting state instead loses stability as an unstable periodic orbit shrinks into this point. For neural systems, there is typically also a stable periodic orbit in the state space. Once the resting state becomes unstable, trajectories quickly diverge from the resting state to the stable periodic orbit, which has a non-zero frequency f_{bif}; see the lefthand schematic in Figure 2. This results in a

discontinuous $f - I$ curve that jumps from zero (at rest) to f_{bif} at $I_{bif}(g_A)$ — the characteristic of Type II excitability.

We use software tools to automate the bifurcation analysis of the Connor-Stevens model: specifically XPP (Ermentrout, 2002), and MatCont (Dhooge et al., 2003). This allows us to track fixed points and limit cycles as system parameters g_A, μ vary, and to check the mathematical conditions that define bifurcation types (Guckenheimer and Holmes, 1983).

Linking linear response theory, spike-triggered averages, and spike count correlations

When the variance c of the shared input is small (see Fig. 1), then we can treat the circuit with a shared input as a perturbation from two independently firing neurons. We describe this perturbation via linear response theory (Lindner et al., 2005; de la Rocha et al., 2007; Ostojic et al., 2009; Ostojic and Brunel, 2011), which is related to classical Linear-Poisson (LP) models of neural spiking (Perkel et al., 1967). That is, we make the assumption that the *change* in a neuron's instantaneous firing rate $\nu_i(t)$ due to the shared input signal can be represented by linearly filtering the common (perturbing) input I_c:

$$
\begin{aligned}
\nu_i(t) &\equiv \langle y_i(t)|I_c\rangle \\
&= \nu_{0,i} + \int_0^\infty K(s)I_c(t-s)\,ds \\
&= \nu_{0,i} + (K * I_c)(t) \quad\quad (4)
\end{aligned}
$$

where the filter $K(t) = 0$ for $t < 0$ (causality) and $\nu_{i,0}$ is the "background" average firing rate of the independently firing neuron.

Eqn. (4) is extremely useful, because it isolates the common component of the response of neurons $i = 1$ and $i = 2$, which is an enhanced (or depressed) tendency to emit spikes, at a rate determined by the filtered, common input. As a result, the cross-covariance function,

$$
C_{12}(\tau) = \langle(\nu_1(t) - \nu_{0,1})(\nu_2(t+\tau) - \nu_{0,2})\rangle
$$

becomes

$$
C_{12}(\tau) = \int_0^\infty \int_{-s}^\infty K(s)K(s+r)\langle I_c(t-s)I_c(t-s+\tau-r)\rangle\,dr\,ds
$$

This simplifies: because we assume I_c is white (uncorrelated in time), $\langle I_c(t)I_c(t+\tau)\rangle = \text{Var}(I_c)\delta(\tau)$. $\text{Var}(I_c)$ is the instantaneous variance of the noisy process (here assumed to be constant), so that

$$
\begin{aligned}
C_{12}(\tau) &= \text{Var}(I_c) \int_0^\infty K(s)K(s+\tau)\, ds \\
&= \text{Var}(I_c)(K * \tilde{K})(\tau) \quad\quad (5)
\end{aligned}
$$

where $\tilde{K}(t) \equiv K(-t)$; cf. (Ostojic et al., 2009).

Thus, the cross-covariance between our two spike trains has a simple expression in terms of the filter with which neurons process the common input into a firing rate. It remains only to identify this filter $K(t)$. As in classical cases, this is precisely given by a *spike triggered average* (Gabbiani and Koch, 1998). Different from the classical setting, but as in Ostojic et al. (2009), it is only the common component of the input current that is averaged in this procedure.

The connection to the spike-triggered average may be seen by looking at the response to a white noise stimulus. On the one hand,

$$
\begin{aligned}
\langle I_c(t)\nu_i(t+\tau)\rangle &= \langle I_c(t)(K * I_c)(t+\tau)\rangle \\
&= \left(K * \langle I_c(t)I_c(t+\tau)\rangle\right)(\tau) \\
&= \text{Var}(I_c)K(\tau)
\end{aligned}
$$

if the noise is white. On the other, by ergodicity, averaging over repeated presentations *and* different stimuli I_c will yield a response equivalent to averaging over repeated presentations of a single long (duration T_{max}) stimulus. As $T_{max} \to \infty$,

$$
\begin{aligned}
\langle I_c(t)\nu_i(t+\tau)\rangle &= \frac{1}{T_{max}} \int_0^{T_{max}} I_c(t)\langle y(t+\tau)|I_c\rangle\, dt \\
&= \frac{1}{T_{max}} \int_0^{T_{max}} I_c(t) \left\langle \sum_{k=1}^N \delta(t+\tau-t_k) \right\rangle dt \\
&= \frac{N}{T_{max}} \left(\frac{1}{N} \sum_{k=1}^N I_c(t_k - \tau) \right) \\
&= \nu_i \left(\frac{1}{N} \sum_{k=1}^N I_c(t_k - \tau) \right) \quad\quad (6)
\end{aligned}
$$

which is exactly the average stimulus preceding a spike (multiplied by the firing rate). Therefore the linear response function $K(\tau)$ is related to the spike-triggered average $\text{STA}(\tau)$ as follows:

$$
\begin{aligned}
K(\tau) &= \frac{\nu_i}{\text{Var}(I_c)} \left(\frac{1}{N} \sum_{k=1}^{N} I_c(t_k - \tau) \right) \\
&= \frac{\nu_i}{\text{Var}(I_c)} \text{STA}(\tau).
\end{aligned} \tag{7}
$$

Below, we use this expression to derive $K(\tau)$ from numerically computed spike triggered averages.

Relating spike-triggered averages and spike-generating dynamics

In order to relate the common input STA (defined in the previous subsection) to spike generating dynamics, it will be helpful for us to derive an explicit formula for the common input STA of a *phase model*, which captures the response of a tonically spiking neuron to a small-amplitude current $I(t)$. Our formula, and the calculation that yields it, is very similar to a relationship previously derived (Ermentrout et al., 2007) for the STA of a phase oscillator *without* background noise.

We consider a model which tracks only the phase of a neuron as it progresses along its periodic spiking orbit:

$$
\frac{d\theta}{dt} = \omega + Z(\theta)I(t), \qquad \theta \in [0, 2\pi) \tag{8}
$$

The function $Z(\theta)$, called a *phase response curve* or PRC (Ermentrout and Kopell, 1984; Winfree, 2001; Ermentrout and Terman, 2010; Reyes and Fetz, 1993), determines how a brief current injection applied at a specific phase of the cycle affects the timing of the next spike. By convention, the neuron is said to "spike" when θ crosses 2π.

To begin, we assume that the phase model is forced by scaled zero-mean, stationary stochastic processes, which we also label $\xi_c(t)$ and $\xi_i(t)$. For now, $\xi_c(t)$ and $\xi_i(t)$ have unit variance and are differentiable with some finite correlation time τ, although we will consider the limit $\tau \to 0$ (i.e. the white noise limit). We are interested in the average value of $\xi_c(t)$ that precedes

a spike; the term $\xi_i(t)$ will play the role of a background noise. Assuming that the background noise process is scaled by a small constant σ, and that ξ_c is scaled by an order of magnitude ϵ smaller still, we write the evolution equation of the phase model as

$$\frac{d\theta}{dt} = 1 + Z(\theta)(\sigma\xi_i(t) + \sigma\epsilon\xi_c(t)), \qquad \theta \in [0, T)$$

where $\sigma, \epsilon \ll 1$. Note that we have chosen our phase variable to have unit speed; i.e $\theta \in [0, T)$, where T is the period of the unperturbed ($\sigma = 0$) oscillator. We proceed as in Ermentrout et al. (2007): writing θ as a series in the small parameters σ and ϵ

$$\theta(t) = \theta_0(t) + \sigma\theta_{10}(t) + \epsilon\theta_{01}(t) + \sigma^2\theta_{20}(t) + \sigma\epsilon\theta_{11}(t) + \epsilon^2\theta_{02}(t) + ...$$

and matching terms of same order in the evolution equation, we find $\theta_0(t) = t$. We additionally find that $\theta_{01} = 0$, $\theta_{02} = 0$, and

$$\begin{aligned}
\theta'_{10} &= Z(t)\xi_i(t) \\
\theta'_{20} &= Z'(t)\xi_i(t)\theta_{10}(t) \\
\theta'_{11} &= Z(t)\xi_c(t)
\end{aligned}$$

so that

$$\begin{aligned}
\theta_{10}(t) &= \int_0^t Z(s)\xi_i(s)\,ds \\
\theta_{20}(t) &= \int_0^t Z'(s)\xi_i(s)\int_0^s Z(r)\xi_i(r)\,dr\,ds \\
\theta_{11}(t) &= \int_0^t Z(s)\xi_c(s)\,ds \ .
\end{aligned}$$

In order to compute the spike-triggered average, we need to find the time of the next spike, assuming the neuron has just spiked ($\theta(0) = 0$); in other words, the time τ when $\theta(\tau) = T$. As above, we expand

$$\tau = T + \sigma\tau_{10} + \sigma^2\tau_{20} + \sigma\epsilon\tau_{11} + ... \tag{9}$$

Using our previous expressions for $\theta(\tau)$, and using the fact that $\tau = T +$

$\sigma\tau_{10} + O(\sigma^2, \sigma\epsilon)$ to decompose the stochastic integrals, we find

$$\tau_{10} = -\int_0^T Z(s)\xi_i(s)\,ds$$

$$\tau_{20} = -\int_T^{T+\sigma\tau_{10}} Z(s)\xi_i(s)\,ds - \int_0^T Z'(s)\xi_i(s) \int_0^s Z(r)\xi_i(r)\,dr\,ds$$

$$\tau_{11} = -\int_0^T Z(s)\xi_c(s)\,ds\ .$$

Next, we use Taylor's theorem for smooth functions to expand ξ_c about $T-t$ to compute

$$
\begin{aligned}
\mathrm{STA}(t) &= \langle I_c(\tau - t)\rangle \\
&= \sigma\epsilon\langle \xi_c(T + \sigma\tau_{10} + \sigma^2\tau_{20} + \sigma\epsilon\tau_{11} - t)\rangle \\
&= \sigma\epsilon\langle \xi_c(T-t) + (\sigma\tau_{10} + \sigma^2\tau_{20} + \sigma\epsilon\tau_{11})\xi_c'(T-t) \\
&\quad + \frac{1}{2}(\sigma\tau_{10} + \sigma^2\tau_{20} + \sigma\epsilon\tau_{11})^2\xi_c''(T-t)\rangle \\
&= \sigma\epsilon\langle \xi_c(T-t) + (\sigma\tau_{10} + \sigma^2\tau_{20} + \sigma\epsilon\tau_{11})\xi_c'(T-t) \\
&\quad + \frac{1}{2}\sigma^2\tau_{10}^2\xi_c''(T-t)\rangle + O(\sigma^3, \sigma^2\epsilon).
\end{aligned}
$$

where we have kept terms up to second order both in our expression for τ, and in our Taylor expansion of ξ_c. We can use the independence of ξ_c and ξ_i to eliminate a large number of terms, as

$$\langle \xi_c(t)\xi_i(t+s)\rangle = \langle \xi_c(t)\rangle\langle \xi_i(t+s)\rangle = 0.$$

Similarly,

$$\langle \xi_c(t)\xi_i'(t+s)\rangle = \langle \xi_c(t)\rangle\langle \xi_i'(t+s)\rangle = 0, \tag{10}$$

and so forth for expressions with higher derivatives. The only term that survives is

$$
\begin{aligned}
\mathrm{STA}(t) &= \sigma\epsilon\langle \xi_c'(T-t) \times -\sigma\epsilon\int_0^T Z(s)\xi_c(s)\,ds\rangle \\
&= -(\sigma\epsilon)^2 \int_0^T Z(s)\langle \xi_c'(T-t)\xi_c(s)\rangle\,ds \\
&= -(\sigma\epsilon)^2 \int_0^T Z(s)A_c'(T-t-s)\,ds \\
&= -(\sigma\epsilon)^2 \int_0^T Z'(s)A_c(T-t-s)\,ds,
\end{aligned}
$$

where A_c is the autocovariance function of ξ_c and we used integration by parts in the final step. Taking the white noise limit $(A_c(T - t - s) \rightarrow \delta(T - t - s))$ and using the periodicity of the PRC $(Z(T - t) = Z(-t))$, we recover a very similar expression as Ermentrout et al. (2007):

$$\text{STA}(t) \;\; = \;\; -(\sigma\epsilon)^2 Z'(-t). \tag{11}$$

Numerical simulations and estimates of spike count statistics

To compute spike count correlations and other statistical quantities, we performed Monte Carlo simulations of the circuit in Fig. 1. The governing Connor-Stevens equations (1, 2) were integrated using the stochastic Euler method with time step $\Delta t = 0.01$ ms, for a total time T_{max} of 8×10^6 ms. Random input currents were chosen at each time step using a standard random number generator (Marsaglia and Zaman, 1994). To facilitate exploration of parameter space μ, σ, c, and g_A, we distributed computations on parallel machines through the NSF Teragrid program (http://www.teragrid.org). The simulation code was implemented in FORTRAN90, and distribution scripts in Python for running on clusters with and without PBS submission protocols. All code and scripts will be available at the modelDB site upon publication (http://senselab.med.yale.edu/modeldb/).

We register spikes in our simulations at times when the membrane voltage exceeds -30 mV and maintains a positive slope in voltage for the next three time steps (0.03 ms). To avoid counting each spike more than once, we omit a 2 ms refractory period after each spike.

Spike count statistics were computed directly from the recorded spike times, based on a single long simulation, after discarding an initial transient (200 ms). When sampling spike counts over a time window T, we advance the window by $\frac{1}{4}T$, resulting in approximately $4T_{max}/T$ (correlated) samples; consequently, our estimates of spike counts become noisier as T increases. To estimate standard errors on spike count statistics, we further divided the simulation into 10 equal time intervals (8×10^5 ms each) and computed statistics on each sub-simulation; the standard deviation, divided by $\sqrt{10}$, gives us an estimated standard error of the mean. When appropriate, these are presented along with the mean estimates, as error bars.

Below, we also report spike triggered averages (STAs) described above; these were computed using long simulations of length 8×10^7 ms for several

sets of parameter values μ, σ, and g_A. To compute these, the common input current I_c was treated as the "signal" which was averaged and the private input as a "background" which was not. We used $c = 0.10$ in this computation. In our code, the history I_c was continuously recorded for a duration into the past; when a spike was recorded, the STA was augmented by this current.

Finally, we generate auto- and cross-correlograms (shown in Figure 4) by collecting inter-spike intervals (ISIs) from our simulations in 1 ms-long bins. These are used, after the standard normalization, as auto- and cross-covariance functions.

RESULTS

Rich structure of spike count correlations over short and long time scales

Our central findings contrast how different conductance-based neuron models produce correlated spiking when they receive overlapping fluctuating inputs, via the shared-input circuitry in Fig. 1. Specifically, we show how this correlation depends on the Type I vs. Type II excitability class of a neuron described by the well-studied Connor-Stevens model. As discussed above, neurons are often classified as Type I vs. Type II based on whether their firing rate-current curves are continuous vs. discontinuous at $\mu = I_{bif}$. Figure 2 demonstrates — as shown in (Rush and Rinzel, 1995) — that the Connor-Stevens model is Type II when the maximal A-current conductance $g_A \approx 0$ mS/cm^2, Type I for $g_A \approx 60$ mS/cm^2, and displays a gradual transition in between. Thus, we fix the neurons in the shared-input circuit to a point along the spectrum from Type I to Type II excitability by choosing different values of g_A.

To compute levels of correlated spiking, we then fix the correlation in the input currents — that is, the fraction of the current variance that is shared vs. private to the two cells — to a preset value c. For each value of c and g_A, we compute spike count correlations for wide range of operating points for the neurons, as determined by a $\sim 200 \times 50$ grid of values for the mean current μ and variance σ (both μ and σ are sampled at a resolution of 0.1 μA/cm^2). Specifically, we vary μ over values centered at the threshold current $I_{bif}(gA)$, from a minimum $\mu = I_{bif}(gA) - 10$ (μA/cm^2) to a maximum
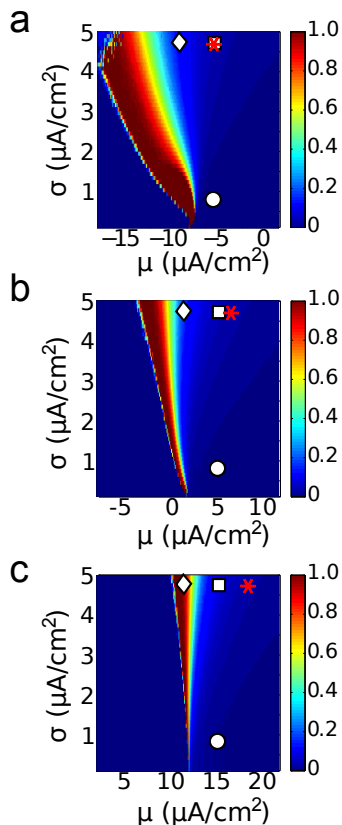
Fig. 3: Fano factor of spike counts over a long time window ($T = 256$ ms) for a $\sim 200 \times 50$ grid of values for the mean current μ and variance σ. From top to bottom, Type II to Type I: (a) $gA = 0$, (b) $gA = 30$, and (c) $gA = 60$ mS/cm^2. Markers indicate relative location of (μ, σ)-pair; subthreshold by 1 μA/cm^2(diamond), superthreshold by 2 μA/cm^2 with low noise (circle) and high noise (square), and superthreshold with matched Fano factors (asterisk, see text).

$\mu = I_{bif}(gA) + 10$ (μA/cm^2) for each value of g_A. This enables us to cover, respectively, both subthreshold (i.e., fluctuation-driven, $\mu < I_{bif}(gA)$) and superthreshold (i.e., mean-driven, $\mu > I_{bif}(gA)$) firing regimes for each value of g_A. We additionally vary σ over $0 < \sigma \leq 5$ (μA/cm^2), so that we cover the range from nearly Poisson, irregular spiking to nearly periodic, oscillatory spiking. This is demonstrated by Figure 3, which shows the Fano factor of spike counts over a long time window ($T = 256$ ms) — a proxy for the squared
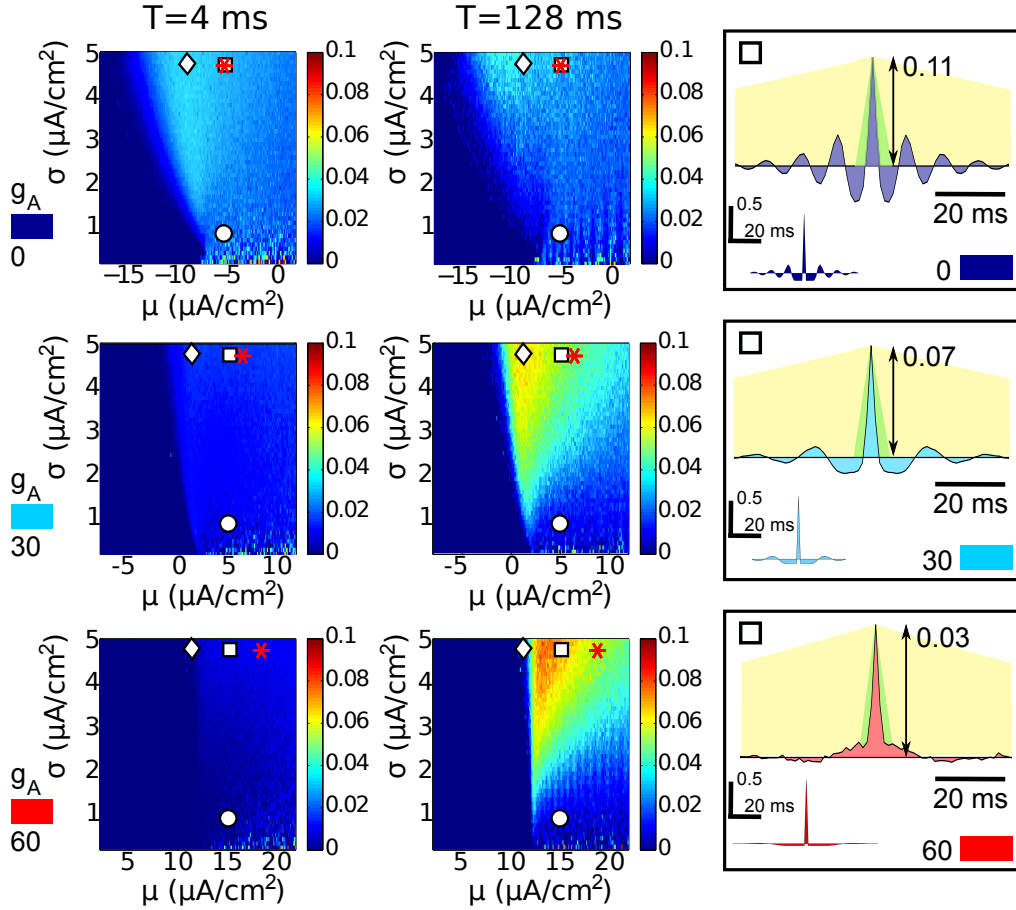
Fig. 4: Spike count correlations for three models at both short and long time scales. Each row displays data from a value of g_A: from top to bottom, Type II ($g_A = 0$ mS/cm^2), intermediate ($g_A = 30$ mS/cm^2), and Type I ($g_A = 60$ mS/cm^2). Left column: Spike count correlations ρ_T, for short windows $T = 4$ ms. Center: Spike count correlations ρ_T, for long windows $T = 128$ ms. Markers indicate points used for cross-model comparison: subthreshold by 1 μA/cm^2 (diamonds), superthreshold by 2 μA/cm^2 and low noise (circles), superthreshold by 2 μA/cm^2 and high noise (squares), and superthreshold with matched Fano factor (stars). Right: Cross-covariance and autocovariance (inset) functions for the superthreshold high noise points (squares). Behind cross-covariance functions, the shape of the triangular kernel that relates this function to spike count covariance (as in Eqn. 3) is illustrated for $T = 4$ ms (green) and $T = 128$ ms (yellow). For each value of g_A, autocovariance functions are given in normalized units (so that $A(0) = 1$).

inter-spike interval coefficient of variation (Gabbiani and Koch, 1998) — over the entire μ, σ parameter space for three representative values of g_A (0, 30, 60 mS/cm^2). For each value of g_A, the Fano factor spans a range from near zero (periodic) to one (i.e. Poisson-like) or higher.

For each set of parameters g_A, c, μ and σ, we compute the Pearson's correlation coefficient ρ_T between the spike counts that the neuron pair in Fig. 1 produces in time windows of length T. Figure 4 summarizes the results, for inputs with 10% shared variance ($c = 0.1$). Here, we view ρ_T over the entire μ, σ parameter space for three representative values of g_A (0, 30, 60 mS/cm^2) and two different time windows ($T = 4$ ms and $T = 128$ ms). Values of ρ_T depend in a strong but systematic way on all of the parameters we have introduced. As we move down a column, we see major qualitative differences in levels of correlation that emerge at different points through the Type II ($g_A = 0$) to Type I ($g_A = 60$) spectrum. Within each panel, the operating point set by input mean and variance (μ, σ) have a strong impact on ρ_T. Finally, the levels and trends in ρ_T depend strongly on the time scale T. We now describe these trends in more detail; the sections that follow will give an explanation for how they arise.

We begin with the upper panels in Fig. 4, which show correlation ρ_T for $g_A = 0$ and hence Type-II excitability. First, note that correlations are overall quite weak. The largest values of ρ_T obtained are ≈ 0.04, indicating that $\approx 40\%$ or less of correlations in input currents are ever transferred into correlations in output spikes. Moreover, the level of correlations ρ_T and their dependence on input parameters μ and σ appear roughly similar for both short and long time scales T. In both cases, for a fixed value of DC input μ, a general trend is that ρ_T gradually increases with fluctuation strength σ. For a fixed value of σ, in general ρ_T first increases and then decreases with μ; the dependence is slightly more complex at longer T. Significantly non-zero values of ρ_T are present for $\mu < I_{bif}$, as σ becomes appreciably high; this reflects the bistable firing dynamics of the underlying deterministic system, which supports both a stable resting state and a stable spiking trajectory for $\mu < I_{bif}$ (see Methods, "Characterizing the dynamics of spike generation").

For Type I excitability at $g_A = 60$ (lower panels in Fig. 4), the picture is dramatically different. First, there is a marked difference between correlation elicited on short vs. long time scales T, with much stronger correlations observed for larger T. Moreover, correlations produced by Type I neurons over longer time scales T are much higher than those observed for Type II neurons at any time scale: the largest values of ρ_T obtained for Type

I are ≈ 0.8, indicating that $\approx 80\%$ of correlations in input currents can be transferred into spike correlations. Conversely, correlations for Type I neurons are strongly suppressed at short time scales, where $\lesssim 10\%$ of input correlations are transferred. Overall, trends in ρ_T as μ and σ vary are similar to those found previously: correlations increase with σ, and first increase — then decrease — with μ.

Correlation transfer in the intermediate model, $g_A = 30$, displays trends between those of the Type I ($g_A = 60$) and Type II ($g_A = 0$) cases. As when $g_A = 60$, spike count correlations ρ_T are very low for short time windows T and attain intermediate to high values $\approx 60\%$ of input correlations transferred for longer T. As for both $g_A = 60$ and $g_A = 0$, ρ_T increases with noise magnitude σ and displays a nonmonotonic trend with mean current μ.

We obtain additional insight into how spike count correlations depend on Type I vs. Type II spike generation and the time scale T by choosing matched values of input parameters μ and σ, and comparing spike count correlations produced for different values of the A-current conductance g_A. We first concentrate on the μ and σ values indicated by squares and circles in Fig. 4. Both of these points indicate superthreshold inputs $\mu = I_{bif}(g_A) + 2$ (μA/cm^2) for all g_A values. The square corresponds to higher noise $\sigma = 5$ μA/cm^2, and the circle to lower noise $\sigma = 1$ μA/cm^2. In Figure 5, we plot ρ_T for a full range of T values from 1 to 200 ms, for nine values of g_A between $g_A = 0$ and $g_A = 60$ (thus filling in intermediate values of g_A and T between those in Fig. 4). For both superthreshold cases, we see that Type II neurons transfer more input correlation into output (spike) count correlation at small T, while Type I neurons transfer more at large T; this transition occurs, roughly, at a value T_{switch} indicated by the dotted line. We note that for the low noise case $\sigma = 1$ (Figure 5(b)), the trends appear less ordered as g_A varies; as we will see in the next section, this is because the cross-covariance function is more oscillatory here, so that ρ_T has not yet converged to its asymptotic large T value.

Subthreshold points, denoted by diamonds in Fig. 4, were also compared: the mean input current is chosen to be $\mu = I_{bif}(g_A) - 1$ (μA/cm^2), and the noise magnitude to be $\sigma = 5$ μA/cm^2. Here, the differing dynamical structure between Type II and Type I neurons is evident in the firing statistics (see Table 1): while the bistable Type II neuron ($g_A = 0$) sustains a substantial firing rate, the monostable Type I neuron ($g_A = 60$) barely fires at this level of input current. The correlation coefficient ρ_T, is also very low for $g_A = 60$ at all time windows (Figure 5(c)); this is consistent with the relationship between
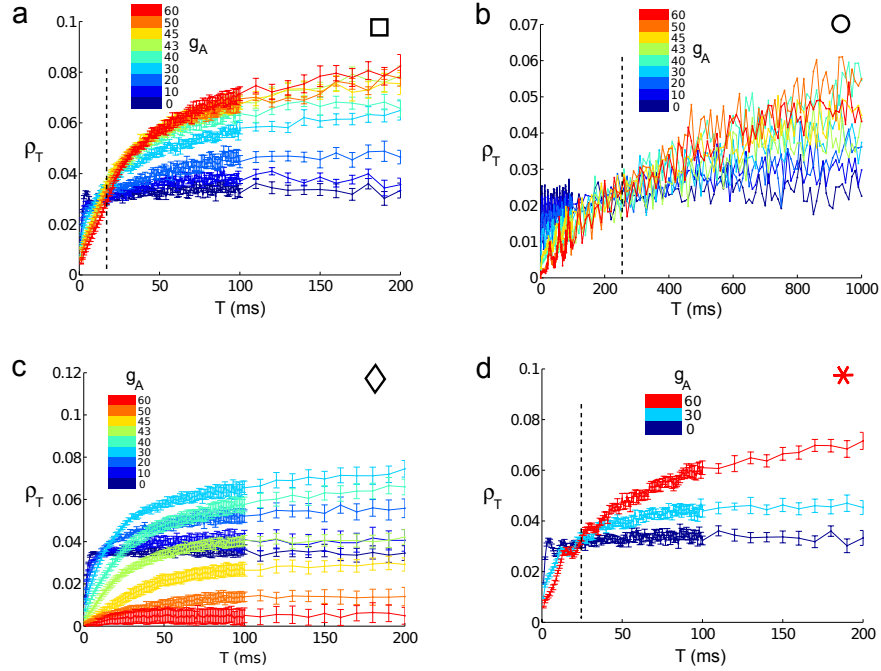
Fig. 5: Correlation coefficient ρ_T vs. time window T. Colors indicate $g_A = 0$ (dark blue) through $g_A = 60$ (red) mS/cm^2. The data from the superthreshold cases (a, b, and d) show the switch from Type II cells transferring more correlations to Type I cells transferring more as T increases. The dotted line indicates the approximate time window T_{switch} where the switch occurs. (a) High noise, superthreshold (b) Low noise, superthreshold, (c) Subthreshold, and (d) high noise, where input current μ has been chosen to match spike train variability across different values of g_A (see text).

correlation and firing rate identified in earlier studies (de la Rocha et al., 2007; Shea-Brown et al., 2008). Overall, note that the correlation coefficient ρ_T increases steadily with T for the Type I neurons (high g_A) but stays roughly constant over a broad range of T for Type II neurons (low g_A). Thus, while we do not observe a clear value of T_{switch} for all values of g_A for the subthreshold point in Figure 5(c), we see the same *relative* trends as for superthreshold points. Below, we will see how this effect follows from filtering properties of Type I vs. Type II cells.

Regime	Statistics	$g_A = 0$	$g_A = 30$	$g_A = 60$
superthreshold, high σ ($\square$)	ν (Hz)	113.2	69.9	31.6
	$FF_{T=256}$	0.059	0.0795	0.195
superthreshold, low σ ($\bigcirc$)	ν (Hz)	108.4	65.0	34.8
	$FF_{T=256}$	0.0107	0.013	0.023
subthreshold ($\lozenge$)	ν (Hz)	81.7	30.0	0.171
	$FF_{T=256}$	0.145	0.282	0.99
superthreshold, fixed variability ($\ast$)	ν (Hz)	113.2	82.5	68.5
	$FF_{T=256}$	0.059	0.059	0.059

Tab. 1: Output firing statistics at each of the comparison points identified in Figure 4 (see text). For each set of matched points, we note the firing rate ν and the Fano factor of spike counts over long time windows (specifically, $T = 256$ ms).

Because spike generation mechanisms vary widely as g_A changes, the neuron models with matched input statistics at different values of g_A in Figs. 5(a, b, c) do not all have the same firing variability. In the final panel of Fig. 5, we address this by showing that the same trends in ρ_T persist if we select values of μ to maintain constant firing variability for each value of g_A (see Table 1; variability measured via large-time ($T = 256$ ms) Fano factor). Here, we fix $\sigma = 5$ μA/cm^2; the required current value μ for each g_A is indicated with a red star in Fig. 4.

In sum, for matched values of the mean and variance of input currents, a pair of superthreshold Type II (vs. Type I) neurons will produce greater spike count correlations ρ_T at short time scales T. For a wide range of choices for the mean and variance, there will be a value of T_{switch} where this relationship reverses, so that Type I (vs. Type II) neurons produce greater ρ_T for $T > T_{switch}$. For matched subthreshold currents, similar trends are present; overall, the presence of a time T_{switch} depends on how the input statistics are chosen.

Finally, we note that the general trends observed here carry over, largely unchanged, to different values of c. Figure 6 shows that, for the range $0.1 < c < 0.5$, trends in how ρ_T changes with excitability type via (g_A) remain consistent. In particular the relationship between input correlation c and spike count correlation ρ_T is roughly linear over this broad range of shared
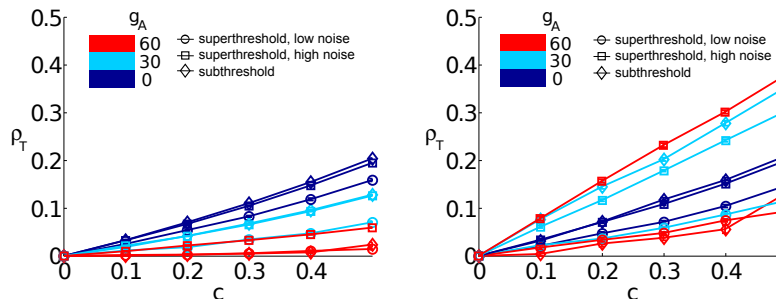
Fig. 6: Output correlation coefficient ρ_T vs. input correlation coefficient c, showing an approximately linear relationship. Left: short time window ($T = 4$ ms). Right: long time window ($T = 150$ ms). Colors indicate $gA = 0$ (dark blue), $gA = 30$ (light blue), and $gA = 60$ (red). Markers indicate relative location of (μ, σ)-pair; subthreshold (diamond), superthreshold with low noise (circle), and superthreshold with high noise (square).

inputs.

Trends in cross-correlation functions for Type I vs. Type II neurons

The trends in spike count correlations that we have just described can be explained from the cross-covariance functions for neuron pairs, and how they differ as the characteristics of input currents and the level of the A-current conductance g_A vary. We now demonstrate this via the cross-covariance functions shown in the right-most column of Fig. 4; these are for the superthreshold high noise cases ($\sigma = 5$) discussed in the previous section.

To make the connection, recall that the spike count covariance, $\text{Cov}(n_1, n_2)$, measured over a window of duration T is given by the integral of the cross-covariance function $C_{12}(\tau)$ against a triangular kernel of width T (see Methods, and Equation 3). Thus, for short windows T, only the central peak of $C_{12}(\tau)$ contributes to spike count covariance. In the limit of long windows $T \to \infty$, the spike count covariance is simply the integral of the cross-covariance function, multiplied by T, over the whole τ axis. Spike-count correlation ρ_T is then given by the *ratio* of spike count covariance to the spike

count variance. As we will show below, ρ_T and $\frac{\text{Cov}(n_1, n_2)}{T}$ often show similar trends with T. Both quantities are of interest: while ρ_T gives a normalized metric of correlation, $\frac{\text{Cov}(n_1, n_2)}{T}$ is the relevant quantity to analyze impact on downstream excitable cells (see below, "Readout of correlated spiking by downstream cells").

Armed with these relationships between $C_{12}(\tau)$, $\text{Cov}(n_1, n_2)$, and ρ_T, we revisit the trends observed in the previous section and explore their origin. Starting in the upper-right of Fig. 4, note that $C_{12}(\tau)$ has a much larger central peak — and hence short-T spike count covariance — for Type II excitability ($g_A = 0$) than for Type I (bottom-right, $g_A = 60$). This is also clear in Fig. 7(a), where we plot spike count covariance vs. T. Over long windows T, the trend reverses. For $g_A = 0$, the cross-covariance function shows oscillations with significant negative and positive lobes. These lobes tend to cancel as $C_{12}(\tau)$ is integrated over long windows T. This cancellation results in little overall change in values of spike-count covariance computed at increasingly long values of T. For Type I excitability, however, $C_{12}(\tau)$ is mostly positive, so that spike count covariance increases with T.

As discussed above, spike-count correlation ρ_T is given by the ratio of spike-count covariance and variance. Comparing Fig. 7(a) and Fig. 5(a) it is clear that spike-count correlation and spike-count covariance display the same trends as g_A is varied. For example, for very short times T spike-count correlation ρ_T is given by the ratio of the peak in $C_{12}(\tau)$ to that in $A_1(\tau)$; this ratio is also larger for Type II vs. Type I excitability.

For other operating points (μ, σ), while trends in spike-count correlation and spike-count covariance do not exactly agree, the some relative trends persist. Specifically, the general pattern that Type II cells produce greater covariance over short time windows, and that this trend disappears or reverses for larger time windows, holds for each of the superthreshold operating points explored here. Overall, the major trends in spike count correlations stem from the presence vs. absence of large negative lobes in cross-covariance functions for Type II vs. Type I neurons. We next describe how this difference arises via the distinct filtering properties of the two neuron types.
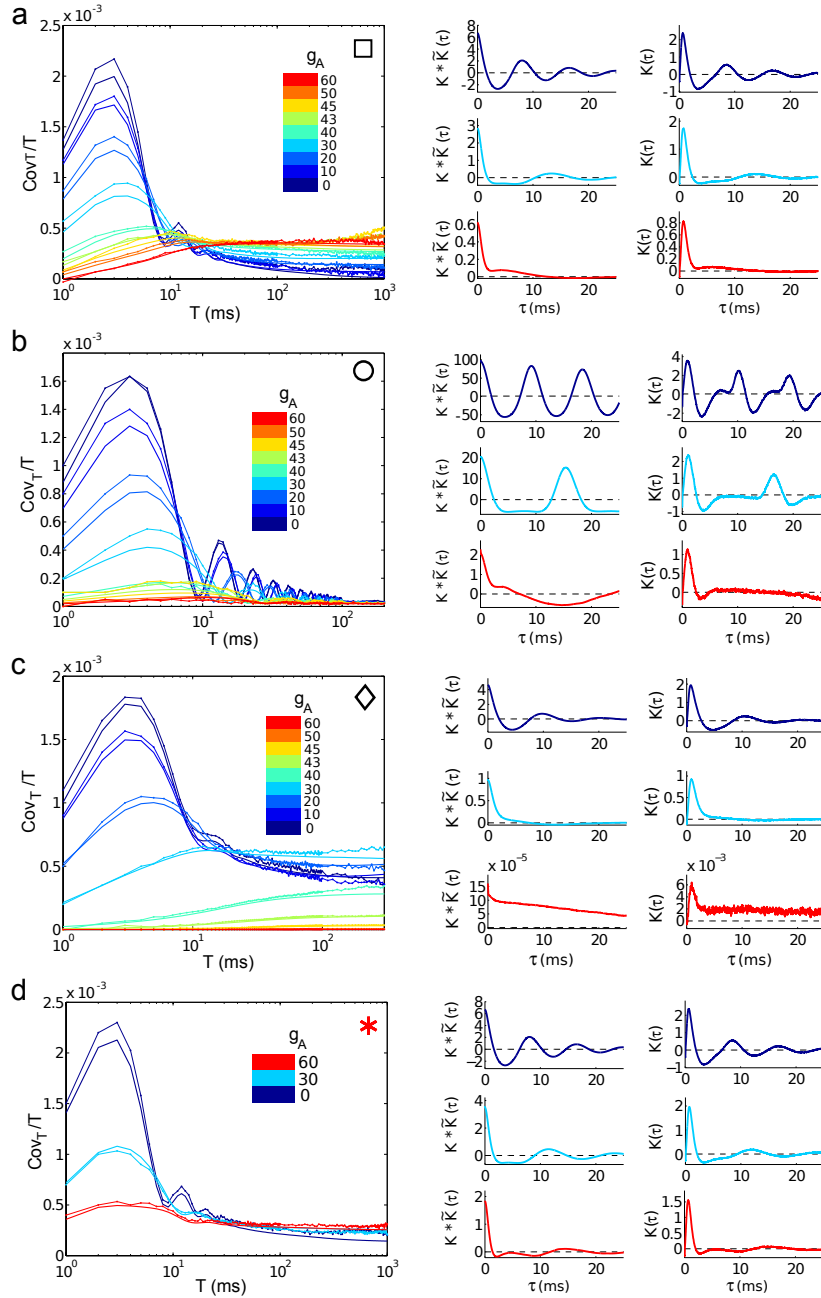
Fig. 7: Spike count covariances, and their relationship to spike triggered averages (STAs). Each row compares data collected at a comparison point for the input current statistics (μ, σ); see text. From top, superthreshold current with high noise (square), superthreshold with low noise (circle), subthreshold (diamond), and superthreshold with high noise and matched variability (red star). (Left column) Actual (dotted lines) and predicted (heavy solid lines) spike count covariances $(\mathrm{Cov}(n_1, n_2)/T)$, for representative points and all g_A values. Colors identify g_A values, which range from $gA = 0$ (dark blue), through $gA = 30$ (light blue), to $gA = 60$ (red) mS/cm^2. (Right column) Select spike-triggered averages (right column) and one-sided cross-covariance functions (left column, derived from the STA using Eqn. 5) used to compute predicted spike count covariances.

Spike triggered average currents reliably predict spike count covariance

The previous section showed how trends in spike count covariances for Type I vs. Type II neurons follow from the presence of both strongly negative and positive lobes in cross-covariance functions for Type II neurons. Here, we explain the origin of this phenomenon. Equations (5,7) (see Methods) provide the key link, in which the cross-covariance function is given in terms of a cell's spike-triggered average (STA), which is an estimate of the filter through which cells turn incoming currents into time-dependent spiking rates. Here, we define the STA in response to common input as an average of the currents that precede spikes over a single long realization:

$$\text{STA} \equiv \left(\frac{1}{N} \sum_{k=1}^{N} I_c(t_k - \tau) \right), \tag{12}$$

where the t_k are the N spike times from the realization.

We first show that the prediction of spike count covariances from STAs is accurate. The left column Fig. 7 shows close agreement between spike count covariances computed from "full" numerical simulation (heavy solid lines) vs. predictions from STAs via Eqn. (5) (dotted lines). Next, we examine the shape of the STAs themselves (Fig. 7, right column).

For each operating point we consider, the Type II STA ($g_A = 0$) has a pronounced negative lobe. Functionally, this corresponds to a "differentiating" mode through which inputs are processed: negative currents sufficiently far in the past tend to drive more vigorous spiking. Biophysically, this corresponds to the kinetics of ionic currents, such as inward currents that can be de-inactivated through hyperpolarization. By contrast, Type I STAs show a less prominent negative lobe, or none at all. The resulting filtering of inputs is characterized as "integrating:" a purely positive filter is applied to past inputs to determine firing rate (cf. (Dayan and Abbott, 2001; Agüera y Arcas et al., 2003; Mato and Samengo, 2008)).

The consequences for spike cross-covariance functions are straightforward. Trends are most pronounced for the superthreshold, high σ case (Figure 7(a)). Here, the pronounced negative lobe in the Type II STA ($g_A = 0$) leads to a similar negative lobe in cross-covariance, and hence a sharp decrease — following an initial increase — of spike count covariance as a function of time window T. The Type I STA ($g_A = 60$) is positive, leading to a spike count

covariance which steadily increases until it overtakes the Type II value at $T \approx 20$ ms. These trends are also reflected in the spike count correlation ρ_T, as described previously. Moreover, analogous plots for superthreshold, high σ points, where spike-count variability (Fano factor) is maintained across g_A values (Fig. 7(d)), show the same trends.

In the low σ case (Fig. 7(b)), the STAs have similar characteristics; however there is signifiant ringing in the STA at the characteristic frequency of the oscillator. By the time the predicted (and actual) covariances $\frac{1}{T}\text{Cov}(n_1, n_2)$ reach a limiting value, they are close to zero, and possibly too variable to order definitively. It appears that covariance is still larger for Type II than for Type I at $T = 200$ ms. In the subthreshold case (Fig. 7(c)), the STA for the Type I neurons is very small, consistent with the very low firing rate here. The Type II neuron shows a more robust response, similar in magnitude but less oscillatory than for the superthreshold regime.

Trends in spike-generating dynamics mirror trends in spike-triggered averages and transferred correlations

The transition from Type II to Type I spike generation in the Connor-Stevens model — as manifest in the progression from discontinuous to continuous spike-frequency vs. current curves in Fig. 2 — can also be characterized via the type of *bifurcation* that governs the transition from quiescence to periodic spiking as increasingly strong currents are injected (see Methods).

For $0 < g_A < 46$ mS/cm^2, the transition occurs via a subcritical Hopf bifurcation, as voltage trajectories jump from a stable rest state to a pre-existing stable periodic orbit (limit cycle). This transition is schematized in the upper-left cartoon in Fig. 2. As this figure shows, for smaller values of g_A in this range, the frequency of this cycle is high ($\gtrsim 60$ Hz). The voltage-conductance dynamics near both stable structures – the stable rest state and the limit cycle – is oscillatory. This creates a *resonator* property (see (Izhikevich, 2007) and references therein): if they are properly timed, both negative and positive inputs cooperatively produce spikes or cause them to occur earlier than they would in the absence of inputs. This is reflected in the negative and positive lobes in the STA $\propto K(\tau)$ for the $g_A = 0$ cases in Figure 7: recall that the STA is the filter applied to incoming currents to determine firing rates.

By contrast, for large $g_A > 58$ mS/cm^2, a saddle-node on invariant circle

(SNIC) bifurcation occurs (see Methods). As sketched in the upper-right cartoon in Fig. 2, in this case there is a pair of fixed points that form a "barrier" to spike generation for subthreshold values of μ, and the shadow, or "ghost" of these fixed points still affects dynamics for superthreshold μ – producing slow dynamics in their vicinity. The inputs that will elicit or accelerate spikes are those that will push trajectories past the fixed points, or their ghost, in a distinguished direction. These inputs therefore tend to have a single (positive) sign. This is referred to as an *integrator* property, and gives rise to the mostly positive STAs seen in Fig. 7 for the $g_A = 60$ cases. (This argument breaks down for very low-variance (low σ) inputs, as we will see in the next section).

Between these two extremes in g_A, the minimum frequency in response to a ramp current decreases steadily, creating a gradual shift between Type I and Type II behavior. This gradual shift is mirrored in the neural dynamics, in which the slow regions in the state space become increasingly dominant. This transition is clear in the spike triggered averages — and therefore spike count covariances —shown in Fig. 7. For example, for $g_A = 30$, we find both distinctly "Type II"-like and "Type I"-like aspects in the high noise (Fig. 7(a)) and subthreshold (Fig. 7(c)) covariance trends respectively. In the former, spike count covariance increases — then decreases — with T; reflecting an oscillating STA; the end result is that the (normalized) covariance at $T = 200$ ms is lower than the covariance at $T = 1$ ms. In the latter, the normalized covariance steadily increases with T, reflecting a non-negative STA.

Phase response curves (PRCs) predict common-input STAs

We next show how, for superthreshold operating points, the key properties of Type I vs. Type II spike generation that determine the filtering of common inputs can be understood via a commonly used and analytically tractable reduced model for tonically spiking neurons. The response of such neurons to an additional small-amplitude current $I(t)$ can be described by a *phase model*, a one-dimensional description which keeps track only of the progress of neuron along its periodic spiking orbit (or limit cycle). Identifying progress along the cycle with a phase $\theta \in [0, 2\pi)$, this model is completely determined by a single function of phase $Z(\theta)$, called a *phase response curve* or PRC (Ermentrout and Kopell, 1984; Winfree, 2001; Ermentrout and Terman, 2010;
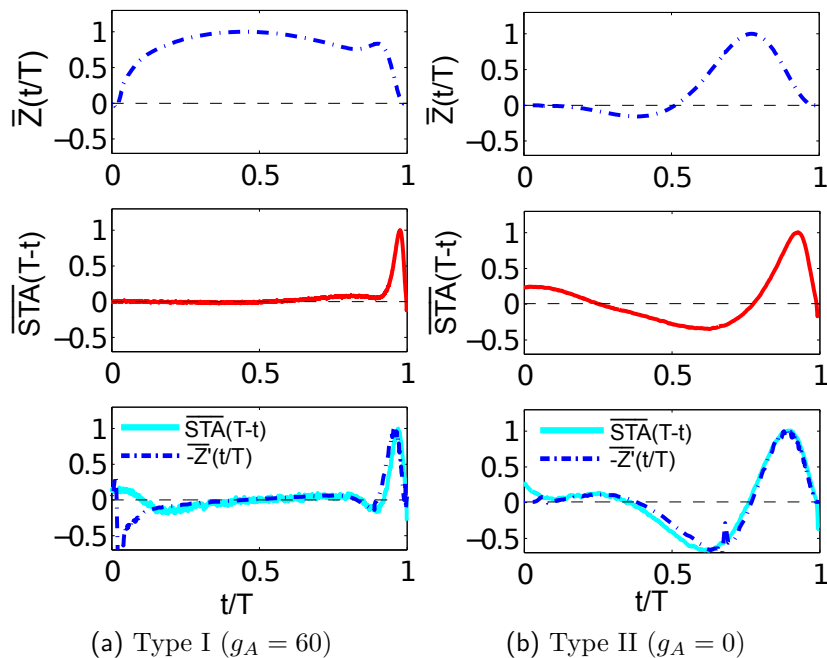
Fig. 8: Comparison of PRCs to spike triggered averages (STAs) computed for both low and high noise, for both Type I (left column) and Type II (right column) neurons. For simplicity of visualization, each curve has been normalized by its maximum; that is $\overline{Z}(t) \equiv Z(t)/\max(Z(t))$, $\overline{\text{STA}}(t) \equiv \text{STA}(t)/\max(\text{STA}(t))$, and $\overline{-Z'}(t) \equiv -Z'(t)/\max(-Z'(t))$. In addition, the time axis has been scaled by the mean period in each case. Top row: PRC, showing monophasic and biphasic shape for Type I and Type II neurons respectively. Middle row: High noise STA; the Type I neuron has lost the negative lobe in its STA, while the Type II neuron retains a negative component. Bottom row: Comparison of (blue dashed) dPRCs with STA for the low noise (cyan) case. Both STAs have negative components.
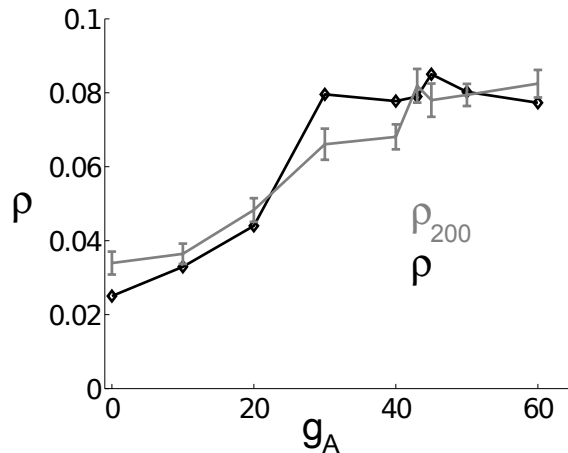
Fig. 9: Correlation coefficient ρ at time window $T = 200$ ms, as g_A is varied. The data (gray solid) are from high-noise, superthreshold points and are the same as reported in Figure 7. The prediction (black solid with diamonds) uses Equation 14. These data show the increase in long time scale correlation as the model transitions from Type II to Type I.

Reyes and Fetz, 1993):

$$\frac{d\theta}{dt} = \omega + Z(\theta)I(t). \tag{13}$$

We can interpret the meaning of this function by considering its effects on the timing of the next spike delivered at a particular phase of the limit cycle ϕ. If $Z(\phi) > 0$, then a positive input delivered at that particular phase will push the neuron further along, advancing the time of the next spike; if $Z(\phi) < 0$, the same input would delay the time of the next spike.

Neurons that display Type I spiking have a purely positive (or Type I) PRC, while Type II neurons show a PRC that has both positive and negative lobes (Ermentrout and Kopell, 1984; Ermentrout, 1996; Hansel et al., 1995; Brown et al., 2004). A purely positive PRC is characteristic of dynamics near a saddle node bifurcation, in which the system lingers near the ghost of its fixed points (as described in the previous section); input in a specific direction is needed to force the system away and elicit a spike. A biphasic PRC reflects oscillatory structure in the phase space, in which correctly timed

negative and positive inputs can cooperate to elicit a spike (as with a Hopf bifurcation).

Strong relationships between the PRC and the STA have been found for neurons close to the threshold for periodic spiking (i.e., $\mu \gtrsim I_{bif}$, see Methods). Spike-triggered covariance analysis of both a Type I phase model and the Wang-Busaki model show that the dominant linear "feature" (corresponding to the STA) qualitatively resembles the PRC (Mato and Samengo, 2008) in the presence of sufficient current noise. In the (Type II) Hodgkin-Huxley model, the two dominant "spike-associated" features identified through covariance analysis closely resemble the STA and its derivative; the STA, in turn, closely resembles the PRC (Agüera y Arcas et al., 2003).

In contrast, phase models in the oscillatory regime (far from the excitability threshold) are known to have an STA proportional to the *derivative* of the PRC (Ermentrout et al., 2007). We generalize this to the case of Fig. 1, where the relevant signal ξ_c is delivered *on top of a noisy background* (see Eqn. (11), Methods): $\text{STA}(t) \propto -Z'(-t)$.

In Figure 8, we test the accuracy of these relationships for the superthreshold points considered above. We show results for Type I ($g_A = 60$, left panels) and Type II neurons ($g_A = 0$, left panels), and compare the STA computed at two different noise levels to the shape of the PRC ($Z(\theta)$) and its derivative, labeled dPRC ($Z'(\theta)$). The time argument of the STA has been scaled so that one period (T) maps onto the unit interval; likewise, the PRC is mapped onto the unit interval. At the lower level of noise, we have good correspondence between the STA and the dPRC in both cases. Notably, both Type I and Type II neurons have biphasic STAs. At high noise levels, while there is not a strong *quantitative* relationship between the STA and the PRC itself (unlike in the excitable regime explored by (Agüera y Arcas et al., 2003)), the PRC carries important clues about the *qualitative* behavior of the STA. The Type II neuron retains the biphasic shape reflective of its PRC, while the Type I neuron has shifted to a purely positive STA. In sum, by predicting the STA shape, the PRC gives important clues to the linear response (and hence common input transfer) that we observe in Figure 7.

Finally, we test an alternate result (cf. (Barreiro et al., 2010; Abouzeid and Ermentrout, 2011)) that, in limited cases, relates PRCs to spike count correlations directly. For for long T and reasonably small σ and c,

$$\rho \approx c \frac{\left[\frac{1}{2\pi} \int_0^{2\pi} Z(x)\, dx\right]^2}{\frac{1}{2\pi} \int_0^{2\pi} (Z(x))^2\, dx} + O(\sigma^2). \tag{14}$$

In Figure 9, we show that this gives a close approximation to simulation results for $T = 200$ ms in the superthreshold, high-noise case (see Table 1). Moreover, we can gain insight into the limitations this asymptotic approximation by comparing with the superthreshold, *low*-noise case. The results of Barreiro et al. (2010); Abouzeid and Ermentrout (2011) are derived by taking the asymptotic limit $T \to \infty$ *before* considering σ finite but small; in practice, the smaller the noise variance σ, the longer T must be in order to see this effect. For our low-noise points, the asymptotic behavior has not been recovered even at $T = 1000$ ms (as may be seen in Figure 5(b)). By using a very large (but probably biologically irrelevant) time window (data not shown), we eventually recover results consistent with the asymptotic prediction (Equation 14).

Readout of correlated spiking by downstream cells

How could the difference in spike count correlation between Type I and Type II cells impact neural circuits? We explore this impact in a simple network, in which correlated Type I or Type II cells collectively converge to drive a neuron downstream (see Fig. 10(a)).

In more detail, the drive comes from a population of $N = 200$ identical Type I ($g_A = 60$ mS/cm^2) or Type II ($g_A = 0$ mS/cm^2) upstream neurons; we refer to these as population I and population II respectively. The upstream populations receive correlated inputs with $c = 0.5$ and values of μ and σ that yield matched levels of variability, as for the parameter set identified with stars in Fig. 4 (for population I, $\mu = 18$ μA/cm^2 and $\sigma = 5$ μA/cm^2; population II, $\mu = -6$ μA/cm^2 and $\sigma = 5$ μA/cm^2). This yields firing rates of $\nu_I = 63.5$ Hz for neurons in population I, and $\nu_{II} = 113$ Hz in population II. Each upstream neuron has a single, instantaneous (delta function) synapse onto the downstream neuron of strength g_I or g_{II}; the relative size of the EPSPs are chosen so that the mean driving current is equal for each population ($\nu_I g_I = \nu_{II} g_{II}$, so that $g_I = 0.825$ mV, $g_{II} = 0.5$ mV).

The total input received by the downstream neuron, I_{ds}, is thus the weighted sum of N upstream spike trains $y_j(t)$:

$$I_{ds}(t) = g_I \sum_j y_j(t) \quad \text{or} \quad I_{ds}(t) = g_{II} \sum_j y_j(t) \,. \tag{15}$$

When the population size N is large, the summed signal has the same temporal characteristics as the the cross-covariance between neuron pairs. Specifi-
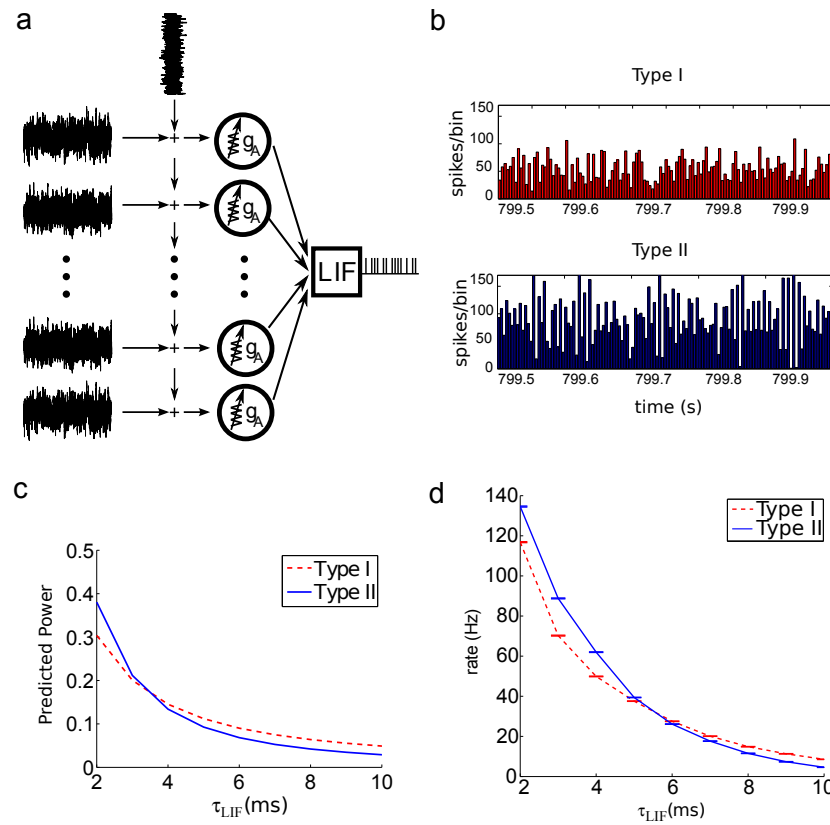
Fig. 10: (a) Schematic of "upstream" Type I or Type II neuron population receiving common and independent inputs, and converging to a Leaky Integrate and Fire (LIF) cell downstream. (b) PSTHs from Type I and Type II upstream populations. (c) Predicted power of the voltage fluctuations in the LIF cell, using STA (see text). (d) Actual firing rates of the LIF cell, showing similar trends with LIF time scale τ_{LIF}.

cally, the autocovariance of the summed input is

$$
\begin{aligned}
A_{ds}(\tau) &= \mathbf{E}\left[(I_{ds}(t) - \langle I_{ds}\rangle)\,(I_{ds}(t+\tau) - \langle I_{ds}\rangle)\right] \qquad &(16)\\
&= N(N-1)\,\mathbf{E}\left[(y_i(t)-\nu_i)(y_j(t+\tau)-\nu_j)\right] + N\,\mathbf{E}\left[(y_i(t)-\nu_i)(y_i(t+\tau)-\nu_i)\right]\\
&\approx N^2\,\mathbf{E}\left[(y_i(t)-\nu_i)(y_j(t+\tau)-\nu_j)\right]\\
&= N^2 C_{12}(\tau)\,. \qquad &(17)
\end{aligned}
$$

This relationship is evident in the peri-stimulus time histograms (PSTHs) in Fig. 10(b). For population II, fast fluctuations above and below the mean population output reflect the negative lobe in $C_{12}(\tau)$ adjacent to its large peak. Meanwhile, fluctuations in the output of population I are less extreme and more gradual in time.

The downstream cell integrates $I_{ds}(t)$ via Leaky Integrate-and-Fire (LIF) voltage dynamics (Dayan and Abbott, 2001):

$$
\tau_{LIF}\frac{dV}{dt} = -(V - V_r) + I_{ds}(t)
$$

where τ_{LIF} is the membrane time constant and $V_r = -60$ mV is the rest voltage. Spikes are produced when the voltage V crosses $V_{thresh} = -45$ mV, at which point V is reset to V_r.

For the parameters we have chosen, the downstream neuron is driven sub-threshold, so that $\langle I_{ds}(t)\rangle$ is not sufficient to excite a spike — any spikes must be driven by *fluctuations* in $I_{ds}(t)$. Thus, the variance of fluctuations in $V(t)$ should give a rough estimate of how often membrane voltage will exceed the threshold, and consequently the downstream firing rate. This variance is easy to compute for a passive membrane (i.e., neglecting spike-reset dynamics). First, note that

$$
V(t) = V_r + \int_{-\infty}^{t} I_{ds}(s) L(t - s)\, ds
$$

where L is a one-sided exponential filter

$$
\begin{aligned}
L(t) &= \frac{1}{\tau}\exp(-t/\tau_{LIF}), \qquad t \geq 0\\
&= 0, \qquad t < 0.
\end{aligned}
$$

We compute the variance as follows, using the causality of L to take each

upper limit of integration to infinity:

$$
\begin{aligned}
\mathbf{E}[V(t)^2] &= \mathbf{E}\left[\int_{-\infty}^{t}(I_{ds}(s)-\langle I_{ds}\rangle)L(t-s)\,ds\int_{-\infty}^{t}(I_{ds}(r)-\langle I_{ds}\rangle)L(t-r)\,dr\right] \\
&= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}ds\,dr\,L(t-s)L(t-r)\,\mathbf{E}\left[I(s)I(r)\right] \\
&= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}ds\,dr\,L(t-s)L(t-r)A_{ds}(s-r) \\
&= \int_{-\infty}^{\infty}dz\,A_{ds}(z)(L*\tilde{L})(-z) \qquad\qquad\qquad (18)
\end{aligned}
$$

where $\tilde{L}(t)\equiv L(-t)$; the last step involved the substitution $z=s-r$ and switching the order of integration. This final interior integral can be evaluated in the Fourier domain: using the properties that $\mathcal{F}[\tilde{L}](\omega)=\mathcal{F}[L](-\omega)$ and the fact that for real functions $\mathcal{F}[f](-\omega)=\overline{\mathcal{F}[f](\omega)}$, we find

$$
\begin{aligned}
\mathcal{F}\left[(L*\tilde{L})\right](\omega) &= |\mathcal{F}[L](\omega)|^2 \\
&= \frac{1}{1+\omega^2\tau_{LIF}^2}\,.
\end{aligned}
$$

Therefore (for example by consulting a transform table)

$$
(L*\tilde{L})(t) = \frac{1}{2\tau_{LIF}}\exp(-|t|/\tau_{LIF})\,.
$$

Substituting into Equation (18), we see that the variance of the downstream cell's voltage is given by a formula similar to that for the spike count covariances (Equation 3): both involve integrating the cross-covariance function against a (roughly) triangular-shaped kernel, with time scale τ_{LIF} in the former case and T in the latter.

Figure 10(c) shows the by-now familiar trends that this predicts. For short membrane time scales τ_{LIF}, Type II populations drive greater voltage variance; this is precisely analogous to the finding that spike-count correlations are greater for Type II cells over short time scales T. For long τ_{LIF}, Type I populations drive greater voltage variance, just as Type I spike trains are more correlated over long time scales T. In panel (d), we compare this trend with actual firing rates elicited in the downstream cell (from numerical simulation). The general trends match, validating our simple prediction.

In sum, downstream neurons with short membrane time scales ($\tau_{LIF} \lesssim 5$ ms) are preferentially driven by Type II cells upstream; for longer time scales, the preference shifts to Type I cells. Some implications of this finding are noted in the Discussion.

DISCUSSION

Diverging connections, leading to overlapped input shared across multiple neurons, are a ubiquitous feature of neural anatomy. We study the interplay between this connectivity pattern and basic properties of spike generation in creating collective spiking across multiple neurons. We range spike generation over the fundamental categories of Type I to Type II excitability (Rinzel and Ermentrout, 1998; Hodgkin, 1948). The transition in excitability is produced by varying the A-current conductance g_A within the well-studied Connor-Stevens neuron model.

Our principal finding is that excitability type plays a major role in how shared — i.e., correlated — input currents are transformed into correlated output spikes. Moreover, these differences depend strongly on the time scale T over which correlations are assessed. At short time scales T, Type-II neurons tend to produce relatively stronger spike correlations for comparable input currents (Marella and Ermentrout, 2008; Galán et al., 2007). At longer time scales, the opposite is generally true: for a broad range of input currents, Type-I neurons transfer most of the shared variance in their inputs ($\sim 80\%$) into shared variance in output spikes, while Type-II neurons transfer less than half ($\sim 40\%$).

We show that these results have direct implications for how downstream neurons with different membrane time constants will respond to Type I vs. Type II populations. Specifically, downstream neurons preferentially respond to populations that are strongly correlated on time scales similar to their membrane time constant. Interestingly, for the case we study, we find that the breakpoint between selectivity to Type I vs. Type II populations was for downstream membrane time constants of ≈ 5 ms, easily within the ranges found experimentally.

This raises interesting possibilities for neuromodulation. The membrane time constant of the downstream cell could be changed by shunting affects of additional background inputs — leading to a switch in its sensitivity to different upstream populations. Alternatively, modulators applied to the

upstream populations themselves could change their excitability from Type I to Type II (Stiefel et al., 2008; Stiefel et al., 2009), adjusting their impact on a downstream cell with a fixed membrane time constant.

Overall, we demonstrate and apply a general principle: the presence and balance among different membrane currents controls not only single-cell dynamics but also the strength and time scales of spike correlations in cell groups receiving common inputs. We show how this relationship can be understood. As a membrane current (here, g_A) is adjusted, firing rate - current curves progressively transition (here, from Type I to Type II). At the same time, there is a transition in periodic orbit types that neural trajectories visit (here, ranging from orbits "near" a fixed point to relatively "isolated" orbits (Rush and Rinzel, 1995)). In turn, this produces a steady progression of spike-triggered averages — and hence the filters that neurons apply to shared input signals (here, from primarily integrating to primarily differentiating modes (Mato and Samengo, 2008), cf. (Agüera y Arcas et al., 2003)). Basic formulas can then be used to translate these filtering properties into predictions for correlated spiking in neural pairs and populations (Ostojic et al., 2009) as well as the downstream impact of this cooperative activity. We anticipate that this approach will bear fruit in studies of the collective activity of a wide variety of neuron types.

Relationship with prior work

A number of prior studies have considered the problem of how spike generating dynamics affect the transfer of incoming current correlations into outgoing spike correlations (Binder and Powers, 2001; de la Rocha et al., 2007; Shea-Brown et al., 2008; Rosenbaum and Josić, 2011; Tchumatchenko et al., 2010; Marella and Ermentrout, 2008; Vilela and Lindner, 2009; Barreiro et al., 2010; Ostojic et al., 2009; Tchumatchenko et al., 2010; Hong and De Schutter, 2008). In particular, (de la Rocha et al., 2007; Shea-Brown et al., 2008; Rosenbaum and Josić, 2011) show that leaky integrate-and-fire (LIF) neurons can transfer up to 100% of current correlations into spike count correlations. The level transferred increases with the firing rate at which single neurons are operating, and the time scale T. These findings are simpler to state compared with the present results for conductance-based neuron models, for which 100% correlation transfer is never obtained, and trends with T differ depending on g_A.

Other works (Hong and De Schutter, 2008; Shea-Brown et al., 2008;

Vilela and Lindner, 2009) investigate correlation transfer in more complex spiking models. In particular, Shea-Brown et al. (2008) and Vilela and Lindner (2009) explore the full parameter space of input currents for the quadratic integrate and fire model — arguably that with the next level of complexity beyond the LIF model. These authors find similar trends in correlation transfer as the neurons' operating points change, but a limitation to 66% rather than 100% in correlation transfer. Meanwhile, Hong and De Schutter (2008) show complex dependencies on neural operating point for the Hodgkin-Huxley model. Taken together, these studies suggested that correlation transfer depends on spike generating dynamics in a rich and diverse ways.

This opened the door to a broader study, but exploring correlation transfer for the full space of possible spike generating dynamics in neural models is a daunting task. The axis that spans from Type I to Type II excitability provides a natural focus. This has been explored using sinusoidal "normal form," phase-reduced models (Marella and Ermentrout, 2008; Galán et al., 2007; Barreiro et al., 2010; Abouzeid and Ermentrout, 2011). These studies used simulations in the superthreshold regime, together with analysis in the limits of very short or very long time scales T, to show the same trend in correlation transfer over short vs. long T that we find here for conductance-based models. A greater frequency of instantaneous (small T) spikes for Type II vs. Type I neurons was predicted using these simplified models (Marella and Ermentrout, 2008; Galán et al., 2007); later, Barreiro et al. (2010) predicted the switch in relative correlation transfer efficiency from Type II to Type I models as T increases.

The present contribution is to test the resulting predictions using biophysical, conductance-based models valid in a wider range of firing regimes, to explain the origins of variable correlation transfer via filtering properties of Type I vs. Type II cells, and to demonstrate the impact on downstream neurons.

Scope, limitations, and open questions

The circuit model that we have studied, as illustrated in Fig. 1, is limited to a single, idealized feature of feedforward connectivity: overlapping inputs to multiple recipient cells. More realistic architecture could include delays in incoming inhibitory vs. excitatory inputs (Gabernet et al., 2005). Interactions of shared-input circuitry with recurrent connectivity also pose

important questions (Ly and Ermentrout, 2009). This is especially so given the distinct properties of Type I vs. Type II cells in synchronization due to reciprocal coupling (Rinzel and Ermentrout, 1998; Ermentrout and Terman, 2010).

Other aspects of our biophysical and circuit dynamics are also idealized. For one, individual input currents fluctuated on arbitrarily fast time scales (i.e., as white noise processes). Relaxing this would be an interesting extension. While prior studies (de la Rocha et al., 2007) suggest that trends will persist for inputs with fast (but finite) time scales, new effects could arise for slower-time scale inputs representative of slower synapses or even network-level oscillations. Another addition would be for inputs to arrive via excitatory and inhibitory conductances, rather than currents; while previous studies with integrate and fire cells (de la Rocha et al., 2007) have found that this yields qualitatively similar results, there could be interesting interactions with underling filtering properties in biophysical models. The same holds true for inputs that arrive at dendrites in multi-compartment models.

Likewise, the circuitry of Fig. 10(a) that we used to investigate the impact of correlated spiking on downstream neurons was highly idealized. An especially appealing extension would be to note that inhibitory and excitatory neurons often have different excitability types. Thus, downstream cells could receive input from both excitatory Type I and inhibitory Type II populations. Our results suggest that sensitivity to excitatory vs. inhibitory afferents would vary with membrane time constants downstream, possibly amplifying the modulatory effects identified here.

References

Abouzeid A, Ermentrout G (2011) Correlation transfer in stochastically driven oscillators over long and short time scales. *ArXiv preprint* p. arXiv:1101.1919v1.

Agüera y Arcas B, Fairhall A, Bialek W (2003) Computation in a single neuron: Hodgkin-Huxley revisted. *Neural Computation* 15:1715–1749.

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

Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366.

Bair W, Zohary E, Newsome W (2001) Correlated firing in macaque visual area mt: time scales and relationship to behavior. *Journal of Neuroscience* 21:1676.

Barreiro A (2009) Transfer of correlations in neural oscillators. Computational and Systems Neuroscience (COSYNE), Salt Lake City, UT.

Barreiro A, Shea-Brown E, Thilo E (2010) Timescales of spike-train correlation for neural oscillators with common drive. *Physical Review E* 81:011916.

Binder MD, Powers RK (2001) Relationship Between Simulated Common Synaptic Input and Discharge Synchrony in Cat Spinal Motoneurons. *J Neurophysiol* 86:2266–2275.

Brown E, Moehlis J, Holmes P (2004) On the phase reduction and response dynamics of neural oscillator populations. *Neural Computation* 16:673–715.

Connor J, Stevens C (1971) Prediction of repetitive firing behavior from voltage clamp data on an isolated neuron soma. *J Physiol* 213:31–54.

Connor J, Walter D, McKown R (1977) Neural repetitive firing: Modifications of the Hodgkin-Huxley axon suggested by experimental results from crustacean axons. *Biophysical Journal* 18:81–102.

Cox DR, Lewis PAW (1966) *The Statistical Analysis of a Series of Events.* John Wiley, London.

Dayan P, Abbott L (2001) *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* MIT Press, 1st edition.

de la Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. *Nature* 448:802–806.

Dhooge A, Govaerts W, Kuznetsov Y (2003) MATCONT: A MATLAB package for numerical bifurcation analysis of ODEs. *ACM Transactions on Mathematical Software* 29:141–164.

Ermentrout GB (1996) Type I membranes, phase resetting curves, and synchrony. *Neural Computation* 8:979–1001.

Ermentrout GB, Kopell N (1984) Frequency Plateaus in a Chain of Weakly Coupled Oscillators, I. *SIAM J Math Anal* 15:215–237.

Ermentrout G (2002) *Simulating, Analyzing, and Animating Dynamical Systems: a guide to XPPAUT for researchers and students.* SIAM.

Ermentrout G, Galan R, Urban N (2007) Relating neural dynamics to neural coding. *Physical Review Letters* 99:248103.

Ermentrout G, Terman D (2010) *Foundations of Mathematical Neuroscience.* Springer, New York.

Gabbiani F, Koch C (1998) Principles of spike train analysis. In Koch C, Segev I, editors, *Methods in neuronal modeling*, pp. 313–360. Cambridge, MA: MIT.

Gabernet L, Jadhav SP, Feldman DE, Carandini M, Scanziani M (2005) Somatosensory integration controlled by dynamic thalamocortical feedforward inhibition. *Neuron* 48:315–327.

Galán RF, Ermentrout G, Urban NN (2007) Stochastic dynamics of uncoupled neural oscillators: Fokker-Planck studies with the finite element method. *Phys Rev E* 76:056110.

Guckenheimer J, Holmes P (1983) *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields.* Springer-Verlag, New York.

Hansel D, Mato G, Meunier C (1995) Synchrony in excitatory neural networks. *Neural Computation* 7:307–337.

Hodgkin A (1948) The local electric changes associated with repetitive action in a non-medullated axon. *J Physiol* 107:165–181.

Hong S, De Schutter E (2008) Correlation susceptibility and single neuron computation. *BMC Neurosci.* 9:P141.

Izhikevich E (2007) *Dynamical Systems in Neuroscience: the Geometry of Excitability and Bursting.* MIT Press.

Lindner B, Doiron B, Longtin A (2005) Theory of oscillatory firing induced by spatially correlated noise and delayed inhibitory feedback. *Phys. Rev. E* 72:061919.

Ly C, Ermentrout G (2009) Synchronization dynamics of two coupled neural oscillators receiving shared and unshared noisy stimuli. *Journal of Computational Neuroscience* 26:425–443.

Marella S, Ermentrout G (2008) Class-II neurons display a higher degree of stochastic synchronization than class-I neurons. *Physical Review E* 77:041918.

Marsaglia G, Zaman A (1994) Some portable very-long-period random number generators. *Computers in Physics* 8:117–121.

Mato G, Samengo I (2008) Type I and type II neuron models are selectively driven by differential stimulus features. *Neural Computation* 20:2418–2440.

Ostojic S, Brunel N, Hakim V (2009) How connectivity, background activity, and synaptic properties shape the cross-correlation between spike trains. *Journal of Neuroscience* 29:10234–10253.

Ostojic S, Brunel N (2011) From spiking neuron models to linear-nonlinear models. *PLoS Comput Biol* 7:e1001056.

Perkel D, Gerstein G, Moore G (1967) Neuronal spike trains and stochastic point process: I. The single spike train. *Biophys J* 7:391–418.

Reyes A, Fetz E (1993) Effects of transient depolarizing potentials on the firing rate of cat neocortical neurons. *J. Neurophysiol* 60:1673–1683.

Rinzel J, Ermentrout GB (1998) Analysis of neural excitability and oscillations. In Koch C, Segev I, editors, *Methods in Neuronal Modeling*, pp. 251–291. MIT Press.

Rosenbaum R, Josić K (2011) Mechanisms that modulate the transfer of spiking correlations. *Neural Computation* 23:1261–1305.

Rush M, Rinzel J (1995) The potassium A-current, low firing rates and rebound excitation in Hodgkin-Huxley models. *Bulletin of Mathematical Biology* 57:899–929.

Salinas E, Sejnowski TJ (2000) Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *J Neurosci* 20:6193–6209.

Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.

Shea-Brown E, Josić K, de La Rocha J, Doiron B (2008) Correlation and synchrony transfer in integrate-and-fire neurons: Basic properties and consequences for coding. *Physical Review Letters* 100:108102.

Shea-Brown E, Thilo E, Barreiro A, Josić K, de la Rocha J, Johnson T, Doiron B (2009) How does spike generation control how correlated inputs become correlated spikes? Abstract 321.19, Society for Neuroscience.

Shlens J, Rieke F, Chichilnisky E (2008) Synchronized firing in the retina. *Curr Opinion Neurobiology* 18:396–402.

Stiefel KM, Gutkin BS, Sejnowski TJ (2008) Cholinergic Neuromodulation Changes Phase Response Curve Shape and Type in Cortical Pyramidal Neurons. *PLoS ONE* 3:e3947.

Stiefel KM, Gutkin BS, Sejnowski TJ (2009) The effects of cholinergic neuromodulation on neuronal phase-response curves of modeled cortical neurons. *J Computational Neuroscience* 26:289–301.

Tchumatchenko T, Malyshev A, Geisel T, Volgushev M, Wolf F (2010) Correlations and synchrony in threshold neuron models. *Phys Rev Lett* 104:58102.

Trong P, Rieke F (2008) Origin of correlated activity between parasol retinal ganglion cells. *Nature Neuroscience* 11:1341–1351.

Vilela RD, Lindner B (2009) Comparative study of different integrate-and-fire neurons: spontaneous activity, dynamical response, and stimulus-induced correlation. *Physical Review E* 80:031909.

Winfree A (2001) *The Geometry of Biological Time.* Springer, New York.

Zohary E, Shadlen M, Newsome W (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143.