

ON STATIONARY DISTRIBUTIONS OF STOCHASTIC NEURAL NETWORKS

K. BOROVKOV * ** AND

G. DECROUEZ,* *The University of Melbourne*

M. GILSON,*** *RIKEN Brain Science Institute and The University of Melbourne*

Abstract

The paper deals with nonlinear Poisson neuron network models with bounded memory dynamics, which can include both Hebbian learning mechanisms and refractory periods. The state of the network is described by the times elapsed since its neurons fired within the post-synaptic transfer kernel memory span, and the current strengths of synaptic connections, the state spaces of our models being hierarchies of finite-dimensional components. We prove the ergodicity of the stochastic processes describing the behaviour of the networks, establish the existence of continuously differentiable stationary distribution densities (with respect to the Lebesgue measures of corresponding dimensionality) on the components of the state space, and find upper bounds for them. For the density components, we derive a system of differential equations that can be solved in a few simplest cases only. Approaches to approximate computation of the stationary density are discussed. One approach is to reduce the dimensionality of the problem by modifying the network so that each neuron cannot fire if the number of spikes it emitted within the post-synaptic transfer kernel memory span reaches a given threshold. We show that the stationary distribution of this ‘truncated’ network converges to that of the unrestricted network as the threshold increases, and that the convergence is at a superexponential rate. A complementary approach uses discrete Markov chain approximations to the network process.

Keywords: Neural network; nonlinear Poisson neuron; Markov process; ergodicity; stationary distribution

2010 Mathematics Subject Classification: Primary 60J25

Secondary 92B20; 60J99

1. Introduction

Neurons are electrically excitable cells whose main function is to process and transmit information. They connect to each other to form neural networks that constitute core components of the nervous system, and so building and studying mathematical models of such networks is of key interest. To justify the modelling approach used in this paper (as described in detail in Section 2), we will briefly describe the mechanism enabling neurons to communicate with each other.

Received 29 June 2012; revision received 17 October 2013.

* Postal address: Department of Mathematics and Statistics, The University of Melbourne, Parkville, VIC 3010, Australia.

** Email address: borovkov@unimelb.edu.au

*** Current address: Departament de Tecnologies de la Informació i les Comunicacions, Universitat Pompeu Fabra, Barcelona 08018, Spain.

The anatomy of a neuron involves three distinct parts with different electrical activity functions: dendrites that form a tree and contain post-synaptic receptors (inputs), the cell body (soma) that integrates the input currents coming from the dendrites, and a long-limbed axon that terminates with pre-synaptic buttons (outputs). A typical feature of the neuronal electrical activity is the propagation of membrane depolarisation. The membrane of a resting neuron is polarised. Brief high-amplitude depolarisations that propagate from the soma along the axon are called action potentials (or spikes) and have a characteristic shape. When a spike reaches an axonal termination that ‘connects’ to a post-synaptic neuron, neurotransmitters are released into the extra cellular space and excite receptors on the post-synaptic neuron (usually on dendrites). This generates a local variation of the membrane potential in that neuron, which propagates towards the soma. The soma can be seen as a spatiotemporal integrator of these post-synaptic potentials (PSPs) to generate an output spike. The soma potential often remains close to the resting value for a few milliseconds after firing an action potential, which is referred to as the refractory period (during which the neuron cannot fire again). These basic elements of the neuronal information processing actually depend upon many different mechanisms at the molecular level, such as ionic concentrations, density of ion channels, axonal myelination, and types of neurotransmitters (for a review, we refer the reader to [1]).

There exists extensive literature on mathematical modelling of both individual neurons and neural networks. We refer the reader interested in neurophysiological principles of neuron and brain operations to [1] and [27], and to the more advanced expositions in [31] and [32] of the circuitry of the brain. A detailed (but accessible and rather nontechnical) discussion of brain networks, covering structural, functional, and effective connectivity and their respective dynamics, is presented in [33] (the book also contains an extensive bibliography of the relevant research work in the complex network theory). A detailed overview of the computational modelling of nervous systems from the molecular and cellular level, including mathematical modelling of adaptation and learning, is given in [14]. Monograph [22] is a systematic study of the relationship of electrophysiology, nonlinear dynamics, and computational properties of neurons. One can also mention here [15] and [25], and refer the reader to [16] for a recent review of the literature in the area.

As the duration of an action potential is relatively short (usually less than 1 ms for sodium-based action potentials), for modelling purposes, spikes are often considered to be instantaneous. Hence, a neuron can be modelled using a point process whose intensity depends on the past activity of the neuron, its incoming synaptic stimulation, or other mechanisms. The use of point processes, such as Hawkes processes, for modelling the spiking activity of neurons dates back to the papers [10], [11], and [13], and made it possible to study analytically the neuronal response to various input stimulations (for a review, see [15]). Applications of particular neural nonlinear point process models to real data can be found in [28], [29], [30], [34], and [35].

In Section 2 we present the description of the network model we deal with in this paper. The network is ‘assembled’ of nonlinear Poisson neurons, which can be viewed as extended versions of self-exciting Hawkes point processes. To model Hebbian learning, we allow the strengths of synaptic connections in the model to change depending on the order in which the connected neurons are firing. Instead of using the formalism of point processes (as in, e.g. [7] and [9]), we choose an alternative description in terms of multivariate Markov processes whose states represent the networks’ spiking histories, with the state space being the product of a hierarchy of simplices and finite sets representing synapses’ strengths. This approach proves to be rather convenient and allows one to demonstrate ergodicity of the network processes under rather general conditions (Section 3), and, moreover, to study the stationary distributions thereof. We

show that the stationary joint distribution of the neurons’ spiking histories has a smooth density with respect to a natural measure on the respective factor of the state space, and give upper bounds for the components of that density on different components of the state space of the process.

In Section 4 we discuss a way to reduce the dimensionality of the model and approximate its stationary distribution with more tractable objects. The approach is based on ‘truncating’ the original process by ‘forbidding’ neurons to fire once they have fired a given number n of spikes recently (within the ‘memory window’ of the neuron). This modified process will still be Markovian and ergodic, its stationary distribution confined to a space of lower dimensionality and approximating that of the original process at a superexponential rate in n . In fact, such dynamics do make physical sense when the existence of the refractory period is taken into account, but one can further simplify the model by choosing an even lower threshold n .

In Section 5 we deal with the problem of computing the stationary distributions of our networks. We derive systems of differential equations for the stationary distributions (unfortunately, they seem to be tractable in the simplest cases only, which are discussed as examples and accompanied by numerical illustrations). We also briefly discuss the approximation of the stationary distribution of our network process by those of discrete Markov chains constructed as discretised (in both time and space) versions of the process. Computing the stationary distributions for the chains is more feasible, as it requires only solving systems of linear algebraic equations.

2. Network dynamics and their description by Markov processes

2.1. Model description

Our neural network model consists of N neurons and M external sources. Both external sources and neurons can fire spikes, which are assumed to be generated by a random mechanism. External sources are assumed to fire according to independent Poisson processes with constant rates $\hat{\rho}_k$, $k \in \{1, 2, \dots, M\}$ (all the quantities related to external sources will be labelled with hats, and all the processes in the paper will be assumed to be right-continuous), whereas a neuron’s instantaneous firing rate is determined by the value of the activation function of the so-called synaptic influx. For neuron i , the latter is the sum of all PSPs generated by spikes arriving to the i th neuron’s synapses from external sources and other neurons, and also background activity.

More precisely, let $\{\hat{T}_{k,n}\}_{n \in \mathbb{Z}}$ be the times at which external source $k \in \{1, \dots, M\}$ fired, and let $\{T_{j,n}\}_{n \in \mathbb{Z}}$ be the times at which neuron $j \in \{1, \dots, N\}$ fired. Let $\hat{\epsilon}_{ik}(t)$ and $\epsilon_{ij}(t)$ denote the post-synaptic response kernel functions describing the effects on the neuron i potential from accepting spikes through synapses connecting source k to neuron i and neuron j to neuron i , respectively. Then the total time t synaptic influx for neuron i is given by

$$J_i(t) := v_i + \sum_{k,m} \hat{W}_{ik}(\hat{T}_{k,m}) \hat{\epsilon}_{ik}(t - \hat{T}_{k,m}) + \sum_{j,n} W_{ij}(T_{j,n}) \epsilon_{ij}(t - T_{j,n}),$$

where v_i is a constant representing the background activity for neuron i , the synaptic weights $\hat{W}_{ik}(t)$ and $W_{ij}(t)$ can be positive (excitatory synapse) or negative (inhibitory synapse), and, to reflect brain plasticity (e.g. to model Hebbian learning), they can depend on time as well. If, at time t , there is no synaptic connection of external source k to neuron i , we simply have $\hat{W}_{ik}(t) = 0$, and likewise for network’s neuron connections.

The kernels $\hat{\epsilon}_{ik}(t) \geq 0$ and $\epsilon_{ij}(t) \geq 0$ are assumed to vanish outside a compact interval: for any $k \leq M$ and $i, j \leq N$,

$$\hat{\epsilon}_{ik}(t) = \epsilon_{ij}(t) = 0 \quad \text{for } t \notin [0, \Theta], \quad \Theta = \text{constant} > 0,$$

which ensures causality and also means that the *direct* effect of any given spike on a neuron completely disappears within a finite time Θ (for real-life neurons, the order of magnitude of Θ is 10^2 ms).

The previous papers [17] and [19] by one of the authors dealt with the case where the shapes of all kernels were identical to some function ϵ , but could incorporate individual synaptic delays \hat{d}_{ik} and d_{ij} :

$$\hat{\epsilon}_{ik}(t) = \epsilon(t - \hat{d}_{ik}) \quad \text{and} \quad \epsilon_{ij}(t) = \epsilon(t - d_{ij}).$$

The delays account for both the axonal propagation of action potential up to the synaptic site, and for the diffusion time of the neurotransmitters in the synaptic cleft. In the present paper we allow each synapse to have individual properties.

The effect of the synaptic influx on the behaviour of neuron i is expressed via an activation function $\zeta_i(\cdot)$, which is assumed to be continuous nondecreasing (and usually ‘S-shaped’), with

$$0 < \underline{\zeta}_i \leq \zeta_i(x) \leq \bar{\zeta}_i \leq \bar{\zeta} := \max_j \bar{\zeta}_j < \infty, \quad x \in \mathbb{R}, \quad i \leq N.$$

Namely, denoting by $T_i(t)$ the time of the last spike fired by neuron i prior to time t and by \mathcal{F}_t the σ -algebra generated by the evolution of our system up to time t , and setting $\Delta(t) := (t, t + \Delta)$ for $\Delta > 0$, we have, as $\Delta \rightarrow 0$,

$$\mathbb{P}(\text{neuron } i \text{ fires during } \Delta(t) \mid \mathcal{F}_t) = \zeta_i(J_i(t))r(t - T_i(t))\Delta + o(\Delta), \tag{1}$$

where we used a left-continuous function $r(\cdot) \in [0, 1]$ to model the existence of the so-called absolute refractory period, i.e. the time period during which a just fired neuron is unable to fire again. One can take, e.g.

$$r(s) := \mathbf{1}(s \notin (0, \delta_{AR})), \quad \delta_{AR} = \text{constant} > 0, \tag{2}$$

the indicator function of the complement of the interval $(0, \delta_{AR}]$ (for real-life neurons, δ_{AR} is about 1 ms). Whatever the shape of r , we always assume that $r(s) = 1$ for $s \geq \Theta$.

In addition to (1), we assume that

$$\mathbb{P}(\text{more than one neuron fires during } \Delta(t) \mid \mathcal{F}_t) = o(\Delta) \quad \text{as } \Delta \rightarrow 0,$$

which basically means that, given the past history \mathcal{F}_t , the instantaneous firing of different neurons is driven by independent random mechanisms.

Note that the widely studied classical Hawkes process (see [6], [8], [9], [21], and [24]) corresponds to the identity activation function ζ_i in (1), and that the positivity of ζ_i means that neurons can fire spikes in the absence of any external stimulation. The use of nonlinear bounded functions ζ_i is motivated by the experimentally observed saturation of the neuronal firing rate when its excitation increases. Observe also that the temporal spread of the synaptic responses (modelled by ϵ_{ij}) induces specific temporal correlations between the neuronal spike trains, which can be evaluated for the case of linear activation functions ζ_i [18], [19], [21].

Furthermore, our model assumes a Hebbian learning mechanism in the form of spike-timing-dependent plasticity (STDP): if, within a short enough time interval, there are spikes at both pre-synaptic and post-synaptic sides of a connection, this can change the weight of the connection.

The weight increases if the post-synaptic spike follows the pre-synaptic spike (reinforcement of the synapse), and decreases otherwise (depression of the synapse).

For simplicity, we assume that, for each of the connections, the synaptic weight can assume only finitely many values: for a common finite L ,

$$\widehat{W}_{ik}(t) \in \widehat{G}_{ik} := \{\hat{g}_{ik}(1) \leq \hat{g}_{ik}(2) \leq \dots \leq \hat{g}_{ik}(L)\},$$

$$W_{ij}(t) \in G_{ij} := \{g_{ij}(1) \leq g_{ij}(2) \leq \dots \leq g_{ij}(L)\},$$

and the following discrete approximation of the STDP mechanisms discussed in, e.g. [12]. For any $i, j \in \{1, \dots, N\}$ and $m \in \{1, \dots, L\}$, we have the collection of points

$$-\infty < u_{ij}(m, m + 1) < u_{ij}(m, m + 2) < \dots < u_{ij}(m, L + 1) = 0$$

$$= u_{ij}(m, 1) < u_{ij}(m, 2) < \dots < u_{ij}(m, m) < \infty.$$

In real-life situations, the length $\delta_{LW} := \max\{u_{ij}(m, m), |u_{ij}(m, m + 1)|\}$ of the ‘learning window’ is about 10^2 ms. We assume that $\delta_{LW} < \Theta$.

Now suppose that, for a given time t , we have $W_{ij}(t-) = g_{ij}(m)$ and either $t = T_i(t)$ or $t = T_j(t)$ (i.e. one of the neurons i, j fired at time t). Then we put

$$W_{ij}(t) := g_{ij}(d) \quad \text{if} \quad \begin{cases} t = T_i(t) & \text{and} & T_j(t) - t \in (u_{ij}(m, d), u_{ij}(m, d + 1)], \\ t = T_j(t) & \text{and} & t - T_i(t) \in (u_{ij}(m, d), u_{ij}(m, d + 1)], \end{cases}$$

for $d \in \{1, \dots, L\}$, whereas otherwise the value of the weight remains unchanged. The dynamics of the weights \widehat{W}_{ik} have the same nature.

Note that the above mechanism allows one to model the emergence of new synaptic connections as well. Altogether, our network model provides a certain degree of biological realism together with a mathematical framework that allows a tractable analysis.

2.2. A Markovian representation for the network model

Note that, at time t , the knowledge of all the current synaptic weights and the times of all the spikes fired in the network within the time interval $(t - \Theta, t]$ is all the information from the past and present that one needs to uniquely specify the probability distribution of the future evolution of the system. Therefore, to obtain a Markovian description of the network, we denote by $\hat{v}_k(t)$ the number of spikes fired by external source k in the time window $(t - \Theta, t]$, $k \leq M$. If $\hat{v}_k(t) = 0$ then we say that source k was at the state

$$\widehat{X}_k(t) \equiv (\widehat{X}_{k,1}(t), \widehat{X}_{k,2}(t), \widehat{X}_{k,3}(t), \dots) = (0, 0, 0, \dots) \in \mathbb{R}_+^N$$

at time t . If $\hat{v}_k(t) \geq 1$, we set $\widehat{X}_{k,1}(t) := \widehat{T}_k(t) - t + \Theta \in (0, \Theta]$, which is the time till the last spike fired by k prior to the ‘present’ time t disappears from the moving window $(s - \Theta, s]$, $s \geq t$. Furthermore, we denote by $\widehat{X}_{k,2}(t) := \widehat{T}_k(\widehat{T}_k(t)-) - t + \Theta \in (0, \Theta]$ the time till the second to last spike fired by k prior to time t disappears from the moving window $(s - \Theta, s]$, and so on. This way, we obtain the vector

$$\widehat{X}_k(t) = (\widehat{X}_{k,1}(t), \widehat{X}_{k,2}(t), \dots, \widehat{X}_{k,\hat{v}_k(t)}(t), 0, 0, \dots), \quad k = 1, \dots, M,$$

with $\hat{v}_k(t) < \infty$ and $\widehat{X}_{k,1}(t) > \widehat{X}_{k,2}(t) > \dots > \widehat{X}_{k,\hat{v}_k(t)}(t) > 0$ almost surely (the activation functions are bounded and having two spikes at exactly the same time is a zero probability event).

Likewise, the state of neuron i is described by the vector

$$X_i(t) = (X_{i,1}(t), X_{i,2}(t), \dots, X_{i,v_i(t)}(t), 0, 0, \dots), \quad i = 1, \dots, N,$$

with $X_{i,1}(t) > X_{i,2}(t) > \dots > X_{i,v_i(t)}(t) > 0$ almost surely, where $v_i(t)$ is the number of spikes fired by i during $(t - \Theta, t]$. Now the complete history of spikes within the time window $(t - \Theta, t]$ is described by the vector

$$Z(t) := (\widehat{X}(t); X(t)) := (\widehat{X}_1(t), \widehat{X}_2(t), \dots, \widehat{X}_M(t); X_1(t), X_2(t), \dots, X_N(t)).$$

We also need to specify the state of the synaptic connections. This can be done by using the matrices

$$\widehat{W}(t) = (\widehat{W}_{ik}(t))_{i \leq N, k \leq N}, \quad W(t) = (W_{ij}(t))_{i, j \leq N},$$

in which nonexistent connections will be represented by zero entries. Now the value

$$U(t) := (Z(t); \widehat{W}(t); W(t))$$

will completely specify the time t state of the network. The state space for the process U will be taken to be

$$H := S \times \left(\prod_{i \leq N, k \leq M} \widehat{G}_{ik} \right) \times \left(\prod_{i, j \leq N} G_{ij} \right)$$

with $S := E^{M+N}$, where $E := \bigcup_{n \geq 0} E^{(n)}$ is the union of the simplices

$$E^{(n)} := \{(x_1, x_2, \dots) \in \mathbb{R}_+^N : \Theta \geq x_1 > x_2 > \dots > x_n > 0; x_{n+m} = 0, m > 0\}$$

for $n = 0, 1, 2, \dots$. Note that the n -dimensional simplex $E^{(n)}$ is a face of the $(n + 1)$ -dimensional simplex $E^{(n+1)}$, $n \geq 0$. We will endow H with the product σ -algebra

$$\mathcal{H} := \mathcal{I} \otimes \left(\bigotimes_{i \leq N, k \leq M} 2^{\widehat{G}_{ik}} \right) \otimes \left(\bigotimes_{i, j \leq N} 2^{G_{ij}} \right),$$

where $\mathcal{I} := \mathcal{C}^{\otimes(M+N)}$ and \mathcal{C} is the trace of the cylindrical σ -algebra on the space \mathbb{R}^N on E .

2.3. The model dynamics

The process U is a piecewise deterministic (linear in its component Z) Markov process which evolves as follows.

(i) Inside jump-free time intervals, we have, for any $k \leq M, i \leq N$, and $n \geq 1$,

$$\frac{d\widehat{X}_{k,n}(t)}{dt} = -\mathbf{1}(\widehat{X}_{k,n}(t) > 0), \quad \frac{dX_{i,n}(t)}{dt} = -\mathbf{1}(X_{i,n}(t) > 0),$$

whereas $\widehat{W}(t)$ and $W(t)$ remain unchanged. Thus, in the absence of jumps, all the nonzero components of the process decay at the unit rate, and when the ‘first visible in the window’ spike of, say, neuron i that occurred at time $T_{i,n}$ ‘disappears’ from the moving time window $(t - \Theta, t]$ at time $t' = T_{i,n} + \Theta$, the number of positive components of X_i drops by 1: $v_i(t') = v_i(t' -) - 1$.

(ii) Given that the current state of the process is $\mathbf{u} = (\hat{\mathbf{x}}; \mathbf{x}; \hat{\mathbf{w}}; \mathbf{w}) \in H$, where $\hat{\mathbf{x}} = (\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_M)$ has components $\hat{\mathbf{x}}_k = (\hat{x}_{k,1}, \dots, \hat{x}_{k,m_k}, 0, 0, \dots) \in E^{(m_k)}$ with $m_k \geq 0, k \leq M$, and likewise $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ has $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,n_i}, 0, 0, \dots) \in E^{(n_i)}, i \leq N, \hat{\mathbf{w}} = (\hat{w}_{ik})$

and $\mathbf{w} = (w_{ij})$ take values in the respective finite product set, the instantaneous firing rate for source k is $\hat{\rho}_k$, and, for neuron i , it is given by

$$R_i(\mathbf{u}) := \varsigma_i \left(v_i + \sum_{k \leq M} \hat{w}_{ik} \sum_{m \geq 1} \hat{\epsilon}_{ik}(\Theta - \hat{x}_{k,m}) + \sum_{j \leq N} w_{ij} \sum_{n \geq 1} \epsilon_{ij}(\Theta - x_{j,n}) \right) r(\Theta - x_{i,1}). \quad (3)$$

(iii) When source k fires (say, at time $t' = \hat{T}_{k,m}$), the only change in the state of Z is in the component \hat{X}_k :

$$\hat{v}_k(t') = \hat{v}_k(t'-) + 1$$

almost surely (as it is impossible to simultaneously ‘lose’ a spike in the time window and acquire a new spike). The new values of the components are

$$\begin{aligned} \hat{X}_{k,1}(t') &= \Theta, \\ \hat{X}_{k,2}(t') &= \hat{X}_{k,1}(t'-), \\ \hat{X}_{k,3}(t') &= \hat{X}_{k,2}(t'-), \\ &\vdots \\ \hat{X}_{k,v_k(t')}(t') &= \hat{X}_{k,v_k(t')-1}(t'-). \end{aligned}$$

Likewise, a spike fired by neuron i will mean similar changes in the component X_j .

The synaptic weight \hat{W}_{ik} can change when the spike was fired either by source k or by neuron i . As $\delta_{LW} < \Theta$, the state of U just prior to the spike completely specifies to what value the synaptic weight should change, according to the learning rules listed in the description of the model. Similarly for the weights W_{ij} that can change when the spike is fired by either of the neurons i and j .

It is quite straightforward to write down the generator of the process U , following the above description. The vector field specifying the dynamics of the process Z between jumps is piecewise linear, and it changes its direction when, for one of the components $\hat{X}_k \in E$ or $X_i \in E$, the respective integral curve running inside $E^{(n)}$, $n > 1$, hits the face $E^{(n-1)}$ of that simplex and then continues inside that lower-dimensional simplex. The domain of the generator will consist of all bounded functions $U \mapsto \mathbb{R}$ that are path continuous and differentiable in z for that vector field (cf. [23]).

3. Ergodicity and the properties of stationary distributions

In this section we establish strong ergodicity of the Markov process U . In the case of constant synaptic weights, the ergodicity of the network immediately follows from Theorem 5 of [7] on the stability of multivariate point processes with bounded memory dynamics (see also Theorem 6 of [7] for the stability of a nonlinear multivariate Hawkes process with PSP transfer kernels having unbounded supports, and [26]). However, even in that special case, our Markov process framework allows us to come up with a much shorter and simpler proof of stability.

Theorem 1. *Under the stated assumptions, the process U is strongly ergodic: it has a unique stationary distribution π on (H, \mathcal{H}) such that*

$$\sup_{\mathbf{u} \in H} \sup_{B \in \mathcal{H}} |\mathbb{P}(U(t) \in B \mid U(0) = \mathbf{u}) - \pi(B)| \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

Moreover, the convergence is exponentially fast.

It would be most interesting to know the properties of π . One basic fact that we can easily establish is that the stationary distribution π_S of the component $Z(t) = (\widehat{X}(t); X(t))$ on (S, \mathcal{S}) has a density with respect to some naturally chosen measure. Moreover, we can obtain upper bounds for the density.

More precisely, the natural measure on (S, \mathcal{S}) is taken to be the product measure μ^{M+N} , where

$$\mu(B) = \sum_{n \geq 0} \mu_n(B_n) \quad \text{for } B = \bigcup_{n \geq 0} (B_n \times \{\mathbf{0}\}), \quad \mathbf{0} = (0, 0, \dots) \in \mathbb{R}_+^N,$$

with μ_n the n -dimensional Lebesgue measure and B_n the Borel subsets of the respective n -dimensional simplices

$$E_0^{(n)} := \{(x_1, \dots, x_n) \in \mathbb{R}_+^n : \Theta \geq x_1 > x_2 > \dots > x_n > 0\} \tag{4}$$

that can be identified with $E^{(n)} = E_0^{(n)} \times \{\mathbf{0}\}$. We use the convention that μ_0 is just the unit mass at 0.

For $(\mathbf{m}; \mathbf{n}) := (m_1, \dots, m_M; n_1, \dots, n_N) \in \mathbb{Z}_+^{M+N}$, set

$$E^{(\mathbf{m}; \mathbf{n})} := \prod_{k \leq M} E^{(m_k)} \times \prod_{i \leq N} E^{(n_i)}.$$

Similarly, $E_0^{(\mathbf{m}; \mathbf{n})}$ is the product of the respective E_0 sets.

To simplify the formulation of the next theorem, we will slightly abuse notation by identifying the sets $E^{(\mathbf{m}; \mathbf{n})}$ with $E_0^{(\mathbf{m}; \mathbf{n})}$, and so considering the latter as the components of the space S (so that the components of μ are actually given on finite-dimensional spaces).

Theorem 2. *Under the stated assumptions, if all the functions $\hat{\epsilon}_{ik}, \epsilon_{ij}, \zeta_i$, and r are continuously differentiable, then, for any $(\mathbf{m}; \mathbf{n}) \in \mathbb{Z}_+^{M+N}$, the restriction of π_S to $E_0^{(\mathbf{m}; \mathbf{n})}$ has a density $\psi_{\mathbf{m}, \mathbf{n}}$ with respect to μ admitting the upper bound*

$$\left(\prod_{k \leq M} \hat{\rho}_k^{m_k} \right) \left(\prod_{i \leq N} \bar{\zeta}_i^{n_i} \right) \exp\{-\Theta \Sigma_{\hat{\rho}}\} \leq \Lambda^{\Sigma_{\mathbf{m}} + \Sigma_{\mathbf{n}}} \exp\{-\Theta \Sigma_{\hat{\rho}}\},$$

where $\Sigma_{\hat{\rho}} := \sum_k \hat{\rho}_k$, $\Lambda := \max\{\max_k \hat{\rho}_k, \max_i \bar{\zeta}_i\}$, $\Sigma_{\mathbf{m}} := \sum_{k \leq M} m_k$, and $\Sigma_{\mathbf{n}} := \sum_{i \leq N} n_i$. The density function $\psi_{\mathbf{m}, \mathbf{n}}$ is continuously differentiable in the interior of $E_0^{(\mathbf{m}, \mathbf{n})}$ and has finite limits on its boundary.

Remark 1. It is easy to see from the proof of Theorem 2 that, if we assume that the function r has form (2) (and so is not continuously differentiable), then the assertion of the theorem will remain true with the only amendment that the density components $\psi_{\mathbf{m}, \mathbf{n}}$ will be continuously differentiable inside their supports in the spaces of the respective dimensionalities.

Proof of Theorem 1. It is obvious that U is aperiodic and stochastically continuous. Therefore, to establish the ergodicity of U , it suffices to show that the Markov ‘skeleton’ chain $Y = \{Y_n := U(\Theta n), n = 0, 1, 2, \dots\}$ has a recurrent state whose first hitting time distribution tail decays exponentially fast uniformly in the chain’s initial state $U(0)$ (see, e.g. Theorem 18.1 of [2]).

First we will use the standard argument to show that the tail of $\tau := \inf\{n > 0 : Z(\Theta n) = (\mathbf{0}; \mathbf{0})\}$ (i.e. the first value n such that there were no spikes from either the external sources or

the network neurons in the time interval $(\Theta(n - 1), \Theta n]$ admits such a bound. Indeed, setting, for convenience, $\mathbb{P}_u(\cdot) := \mathbb{P}(\cdot \mid U(0) = \mathbf{u})$, we have, for any $\mathbf{u} \in H$ and $t > 0$,

$$\begin{aligned} \mathbb{P}_u(\text{no spikes in } (0, t]) &= \exp\left\{-\int_0^t \left(\sum_k \hat{\rho}_k + \sum_i R_i(\cdot)\right) ds\right\} \\ &\geq \exp\{-(\Sigma_{\hat{\rho}} + \Sigma_{\bar{\zeta}})t\} \\ &=: e^{-\gamma t}, \end{aligned} \tag{5}$$

where $\Sigma_{\bar{\zeta}} := \sum_i \bar{\zeta}_i$, and (\cdot) represents the values of the argument of R_i along the trajectory of U on $[0, t]$ that started at \mathbf{u} and experienced no jumps.

Now setting $A_n := \{\text{no spikes in } (\Theta(n - 1), \Theta n]\}$ we find, recursively using the Markov property and bound (5), that, for $n \geq 0$,

$$\begin{aligned} \mathbb{P}_u(\tau > n) &= \mathbb{P}_u\left(\bigcap_{m=1}^n A_m^c\right) \\ &= \mathbb{E}\mathbb{P}_u\left(\bigcap_{m=1}^n A_m^c \mid Y_{n-1}\right) \\ &= \mathbb{E}\mathbb{P}_u\left(\bigcap_{m=1}^{n-1} A_m^c \mid Y_{n-1}\right) \mathbb{P}_{Y_{n-1}}(A_n^c) \\ &\leq (1 - e^{-\gamma\Theta}) \mathbb{E}\mathbb{P}_u\left(\bigcap_{m=1}^{n-1} A_m^c \mid Y_{n-1}\right) \\ &= (1 - e^{-\gamma\Theta}) \mathbb{P}_u\left(\bigcap_{m=1}^{n-1} A_m^c\right) \\ &\leq \dots \\ &\leq (1 - e^{-\gamma\Theta})^n. \end{aligned}$$

Next we observe that $(\widehat{W}(\Theta n); W(\Theta n))$ is clearly an indecomposable aperiodic finite Markov chain, and hence it is ergodic. Take any fixed state $(\widehat{\mathbf{w}}'; \mathbf{w}')$ of this chain; as is well known, for any initial condition, the first hitting time of that state has an exponentially fast decaying distribution tail. Hence, it is obvious that the state $(\mathbf{0}; \mathbf{0}; \widehat{\mathbf{w}}'; \mathbf{w}') \in H$ will be positive recurrent for the chain Y , and that the tail of the first hitting time of that state by Y will admit a geometrically fast vanishing upper bound uniform in the initial condition of the state. This completes the proof.

Proof of Theorem 2. First we will establish existence of the density for transition probabilities, and then infer the desired result from that fact.

Suppose that our process started at point $U(0) = \mathbf{u}_0$ and that, at time Θ , was at a point $\mathbf{u}_1 = (\hat{\mathbf{x}}; \mathbf{x}; \widehat{\mathbf{w}}; \mathbf{w}) \in H$ with $\hat{\mathbf{x}}_k \in E^{(m_k)}$, $k \leq M$, $\mathbf{x}_i \in E^{(n_i)}$, $i \leq N$. It is clear that the states \mathbf{u}_0 and $\mathbf{z} = (\hat{\mathbf{x}}; \mathbf{x}) \in S$ completely specify the trajectory of $U(t)$ on the time interval $[0, \Theta]$. Denote that trajectory by $\mathbf{u}(t)$, $t \in [0, \Theta]$ (so that $\mathbf{u}(0) = \mathbf{u}_0$ and $\mathbf{u}(\Theta) = \mathbf{u}_1$). Then, observing that $x_{i,1}, \dots, x_{i,n_i}$ are the firing times for neuron i in the time interval $[0, \Theta]$, we use the standard

argument to show that

$$\begin{aligned}
 & \mathbb{P}_{\mathbf{u}_0}(Z(\Theta) \in d\hat{\mathbf{x}}_1 \times \dots \times d\hat{\mathbf{x}}_M \times d\mathbf{x}_1 \times \dots \times d\mathbf{x}_N) \\
 &= \left(\prod_{k \leq M} \hat{\rho}_k^{m_k} e^{-\hat{\rho}_k \Theta} \right) \mu_{m_1}(d\hat{\mathbf{x}}_1) \dots \mu_{m_M}(d\hat{\mathbf{x}}_M) \\
 & \quad \times \left[\prod_{i \leq N} \left(\prod_{l_i \leq n_i} R_i(\mathbf{u}(x_{i,l_i} -)) \right) \exp \left\{ - \int_0^\Theta R_i(\mathbf{u}(t)) dt \right\} \right] \\
 & \quad \times \mu_{n_1}(d\mathbf{x}_1) \dots \mu_{n_N}(d\mathbf{x}_N) \\
 &=: p(\mathbf{u}_0, \mathbf{z}) \mu_{m_1}(d\hat{\mathbf{x}}_1) \dots \mu_{m_M}(d\hat{\mathbf{x}}_M) \mu_{n_1}(d\mathbf{x}_1) \dots \mu_{n_N}(d\mathbf{x}_N) \\
 &= p(\mathbf{u}_0, \mathbf{z}) \mu(d\mathbf{z}). \tag{6}
 \end{aligned}$$

Clearly, the function $p(\mathbf{u}_0, \mathbf{z})$ is continuously differentiable in \mathbf{z} in the interior of $E^{(m,n)}$, has finite limits on its boundary, and admits an upper bound of the form

$$p(\mathbf{u}_0, \mathbf{z}) \leq \left(\prod_{k \leq M} \hat{\rho}_k^{m_k} \right) \left(\prod_{i \leq N} \bar{\varsigma}_i^{n_i} \right) \exp\{-\Theta \Sigma_{\hat{\rho}}\}. \tag{7}$$

Next, in view of (6), for any $B \in \mathcal{F}$, we can use Fubini’s theorem to write

$$\begin{aligned}
 \pi_S(B) &= \int_H \pi(d\mathbf{u}_0) \int_B p(\mathbf{u}_0, \mathbf{z}) \mu(d\mathbf{z}) \\
 &= \int_B \left[\int_H \pi(d\mathbf{u}_0) p(\mathbf{u}_0, \mathbf{z}) \right] \mu(d\mathbf{z}) \\
 &=: \int_B \psi(\mathbf{z}) \mu(d\mathbf{z}). \tag{8}
 \end{aligned}$$

This means that π_S does have density ψ with respect to μ , and (7) implies that ψ admits the desired upper bound.

That ψ is continuously differentiable in the relative interiors of the components of its supporting space follows from (6), the last relation in (8), and the assumption that ς_i and the kernel functions ϵ_{ij} are all continuously differentiable; the ϵ vanish outside $[0, \Theta]$. This completes the proof.

4. Approximating π by finite-dimensional distributions

In this section we will assume constant synaptic weights in our model, so that the time t state of the network will now be completely specified by the value $Z(t)$ only. Even in this simple case, the state space is an infinite hierarchy of multidimensional simplices, so working with nontrivial distributions on it is not easy. The natural question in such a situation is whether one can find an appropriate approximation to the distribution in question, together with an approximation error bound.

For our model, a tempting approach to finding such approximations is to consider ‘truncated’ processes $Z^{(n)}$ in which none of the neurons is ‘allowed’ to fire more than the fixed number $n \geq 1$ times within any given time interval of length Θ . In fact, if the model assumes the existence of absolute refractory periods of positive length by stipulating, say, that (2) holds, then that condition will automatically be satisfied (note, however, that one can still apply truncation with $n < \Theta/\delta_{AR}$ to reduce dimensionality).

The only difference in the dynamics of the process $Z^{(n)}$ compared to those of Z is that the neurons' firing intensities will now be given by

$$R_i^{(n)}(z) := R_i(z)\mathbf{1}(x_{i,n} = 0), \quad i = 1, \dots, N$$

(cf. (3)). It is obvious that $Z^{(n)}$ will also be an ergodic Markov process. Denote its stationary distribution on (S, \mathcal{B}) by $\pi^{(n)}$, while, for the stationary distribution of Z (on the same measurable space), we will reuse the notation π .

Theorem 3. *Under the assumption of constant synaptic weights, we have*

$$\sup_{B \in \mathcal{B}} |\pi(B) - \pi^{(n)}(B)| \leq Cn^{-(n+1)/2} e^{\alpha n},$$

where $C := 2N \exp\{\Theta(\Sigma_{\hat{\rho}} + \Sigma_{\bar{\zeta}})\}/\sqrt{\pi}$ and $\alpha := (1 + \ln(\Theta\bar{\zeta}))/2$.

Proof. We will use a coupling argument. Assume that Π is a Poisson random field of unit intensity on $\mathbb{R}_+ \times \mathbb{R}$, given on some probability space, and construct a process $\{(Z(t), Z^{(n)}(t))\}_{t \geq 0}$ with the state space $S \times S$, whose components will follow the original and 'truncated' dynamics, respectively, start at a common state $Z(0) = Z^{(n)}(0) \in S$, and will be driven by the field Π via the following simple mechanism.

Introduce the intervals

$$\widehat{I}_k := \left(-\sum_{m=1}^k \hat{\rho}_m, -\sum_{m=1}^{k-1} \hat{\rho}_m \right], \quad k = 1, \dots, M,$$

$$I_i := (\vartheta_{i-1}, \vartheta_i], \quad i = 1, \dots, N, \quad \vartheta_i := \sum_{j=1}^i \bar{\zeta}_j,$$

and stipulate that, in both processes Z and $Z^{(n)}$, external source k fires at time t if and only if $\Pi(\{t\} \times \widehat{I}_k) > 0$ (note that $\{\Pi([0, t] \times \widehat{I}_k)\}_{t \geq 0}$ are independent Poisson processes with constant intensities $\hat{\rho}_k$, $k = 1, \dots, M$).

Likewise, in the process Z neuron i fires at time t if

$$\Pi(\{t\} \times (\vartheta_{i-1}, \vartheta_{i-1} + R_i(Z(t-)))) > 0,$$

and that happens in the process $Z^{(n)}$ at time t if

$$\Pi(\{t\} \times (\vartheta_{i-1}, \vartheta_{i-1} + R_i^{(n)}(Z^{(n)}(t-)))) > 0.$$

Clearly, $(Z, Z^{(n)})$ is a well-defined Markov process, and its components follow the desired dynamics. Note also that the process will be ergodic, like each of its components (this is obvious from, e.g. Theorem 1).

Now denote by $(Z(\infty), Z^{(n)}(\infty))$ a random element of $S \times S$ whose distribution coincides with the stationary distribution of $(Z, Z^{(n)})$. Using the standard argument, it is easily seen that

$$|\pi(B) - \pi^{(n)}(B)| \leq \mathbb{P}(Z(\infty) \neq Z^{(n)}(\infty)) =: P_n, \quad B \in \mathcal{B}.$$

To bound P_n , denote by $|T_K|$, $K \in \mathbb{N}$, the total length of the set

$$T_K := \{t \in [0, \Theta K]: Z(t) \neq Z^{(n)}(t)\}$$

and observe that, from the ergodicity of $(Z, Z^{(n)})$, we have

$$P_n = \lim_{K \rightarrow \infty} \frac{|T_K|}{\Theta K}.$$

As we are interested in the stationary distribution, we can assume without loss of generality that the common starting point of Z and $Z^{(n)}$ has no components x_i in $E^{(m)}$, $m \geq n$. Then the trajectories $Z(t)$ and $Z^{(n)}(t)$, having originated at the same point, will coincide with each other till the time T' when one of the values $X_i(t)$, $i = 1, \dots, N$, enters $E^{(n)}$. Then the respective neuron i will stay silent in $Z^{(n)}$ at least till the time when the number of spikes produced by i and ‘visible’ in the time window $(t - \Theta, t]$ drops below n , while in Z the respective neuron will still be able to fire. Thus, past the time point T' , the trajectories $Z(t)$ and $Z^{(n)}(t)$ can diverge. They will have to meet again, though, and the latest that will occur is at the end of the next ‘silent interval’ of length Θ , which, in its turn, occurs no later than the time

$$\inf\{t > T' : \Pi((t - \Theta, t] \times (-\Sigma_{\hat{\rho}}, \Sigma_{\bar{\sigma}}] = 0)\}.$$

To make use of the above argument to obtain an upper bound for P_n , introduce two random sequences, $\{\varkappa_m\}_{m \geq 0}$ and $\{\gamma_m\}_{m \geq 0}$, as follows. Letting, for brevity, $\theta_j := \Theta_j$, set

$$V_{i,j-} := \Pi((\theta_{j-1}, \theta_{j-1} + \frac{1}{2}\Theta] \times I_i), \quad V_{i,j+} := \Pi((\theta_{j-1} + \frac{1}{2}\Theta, \theta_j] \times I_i),$$

and then put $\varkappa_0 := \gamma_0 := 0$, and, for $m \geq 1$,

$$\begin{aligned} \varkappa_{m+1} &:= \inf\left\{j > \gamma_m : \max_{i \leq N} \max\{V_{i,j-}, V_{i,j+}\} \geq \frac{1}{2}n\right\}, \\ \gamma_{m+1} &:= \inf\{j > \varkappa_{m+1} : \Pi((\theta_{j-1}, \theta_j] \times (-\Sigma_{\hat{\rho}}, \Sigma_{\bar{\sigma}}] = 0)\}. \end{aligned}$$

Clearly, both $\{\varkappa_m\}$ and $\{\gamma_m\}$ are well-defined, almost surely infinite increasing sequences of proper random variables.

Now if $Z(\theta_{j-1}) = Z^{(n)}(\theta_{j-1})$ but, for some $t \in (\theta_{j-1}, \theta_j]$, we have $Z(t) \neq Z^{(n)}(t)$, then, for some $i \leq N$, at least one of the following two relations must hold:

$$\max\{V_{i,(j-1)-}, V_{i,(j-1)+}\} \geq \frac{n}{2}, \quad \max\{V_{i,j-}, V_{i,j+}\} \geq \frac{n}{2}$$

(if none of the two holds then, in any time interval of length Θ within $(\theta_{j-2}, \theta_j]$, neuron i will have fewer than n spikes). Thus, the values \varkappa_m ‘mark’ time intervals where Z and $Z^{(n)}$ may split, whereas the γ_m ‘mark’ those intervals following \varkappa_m where Z and $Z^{(n)}$ must merge (provided that they have indeed split).

Set

$$\tau_K := \inf\{m \geq 1 : \varkappa_m > K\} - 1.$$

Clearly, for $t > 0$ and an arbitrarily fixed $\varepsilon > 0$,

$$\mathbb{P}(|T_K| > t) \leq \mathbb{P}(|T_K| > t, \tau_K \leq \varepsilon K) + \mathbb{P}(\tau_K > \varepsilon K). \tag{9}$$

First we will bound the last term. Observe that

$$\tau_K \leq \sum_{j=1}^K \chi_j, \quad \text{where} \quad \chi_j := \mathbf{1}\left(\max_{i \leq N} \max\{V_{i,j-}, V_{i,j+}\} \geq \frac{1}{2}n\right)$$

are independent and identically distributed (i.i.d.) Bernoulli random variables with success probability

$$\mathbb{P}(\chi_j = 1) = \mathbb{P}\left(\max_{i \leq N} \max\{V_{i,j-}, V_{i,j+}\} \geq \frac{n}{2}\right) \leq 2 \sum_{i \leq N} \mathbb{P}\left(V_{i,j-} \geq \frac{n}{2}\right).$$

Since $V_{i,j-}$ has the Poisson distribution with parameter $\lambda_i := \Theta \bar{\zeta}_i / 2$, we can use Taylor’s formula for the exponential series (with remainder in Lagrange form) and then Stirling’s formula to write

$$\begin{aligned} \mathbb{P}\left(V_{i,j-} \geq \frac{n}{2}\right) &\leq \frac{\lambda_i^{n/2}}{(n/2)!} \\ &\leq \frac{1}{\sqrt{2\pi}} \left(\frac{n}{2}\right)^{-(n+1)/2} \exp\left\{\frac{n}{2}(1 + \ln \lambda_i)\right\} \\ &\leq \frac{1}{\sqrt{\pi}} n^{-(n+1)/2} e^{\alpha n}. \end{aligned}$$

Therefore,

$$\mathbb{P}(\chi_j = 1) \leq \frac{2N}{\sqrt{\pi}} n^{-(n+1)/2} e^{\alpha n} =: p_n.$$

Assuming that $p_n < 1$ (otherwise, the bound in the theorem will become trivial), and that $\delta := \varepsilon - p_n \equiv \varepsilon - \mathbb{E}\chi_j > 0$, we obtain

$$\begin{aligned} \mathbb{P}(\tau_K > \varepsilon K) &= \mathbb{P}(\tau_K - K \mathbb{E}\chi_1 > \delta K) \\ &\leq \mathbb{P}\left(\sum_{j \leq K} (\chi_j - \mathbb{E}\chi_j) > \delta K\right) \\ &\leq e^{-2\delta^2 K} \end{aligned} \tag{10}$$

by virtue of Theorem 5.5.1 of [3].

Now we will turn to the first term on the right-hand side of (9). From the definitions of our random variables, it is obvious that $|T_K| \leq \Theta \sum_{m \leq \tau_K} (\gamma_m - \varkappa_m)$, and so

$$\mathbb{P}(|T_K| > t, \tau_K \leq \varepsilon K) \leq \mathbb{P}\left(\sum_{m \leq \varepsilon K} (\gamma_m - \varkappa_m) > \frac{t}{\Theta}\right) =: Q.$$

From the strong Markov property, it follows that $\eta_m := \gamma_m - \varkappa_m$ are i.i.d. geometric random variables, with $\mathbb{P}(\eta_1 = k) = q(1 - q)^{k-1}$, $k = 1, 2, \dots$, where

$$q := \mathbb{P}(\Pi((t - \Theta, t] \times (-\Sigma_{\hat{\rho}}, \Sigma_{\bar{\zeta}}]) = 0) = \exp\{-\Theta(\Sigma_{\hat{\rho}} + \Sigma_{\bar{\zeta}})\}.$$

Clearly, $\mathbb{E}\eta_1 = 1/q$ and $\varphi(a) := \mathbb{E}e^{a\eta_1} < \infty$ for $a < -\ln(1 - q)$, and

$$\varphi(a) = 1 + \frac{a}{q} + o(1) \quad \text{as } a \rightarrow 0. \tag{11}$$

Therefore, assuming without loss of generality that εK is an integer, we have, by the exponential Chebyshev’s inequality,

$$Q \leq (\varphi(a))^{\varepsilon K} e^{-at/\Theta} = \exp\left\{-\varepsilon K \left(\frac{at}{\Theta \varepsilon K} - \ln \varphi(a)\right)\right\}.$$

We can see from (11) that, choosing $t = t_K := \Theta \varepsilon K(1 + h)/q$ for an arbitrarily fixed $h > 0$, we will have, for small enough a , the bound

$$Q \leq e^{-\varepsilon c K} \quad \text{for some } c = c(a, h).$$

From here, (9), and (10), we obtain the bound

$$\mathbb{P}(|T_K| > t_K) \leq e^{-\varepsilon c K} + e^{-2\delta^2 K}.$$

Clearly, $\sum_K \mathbb{P}(|T_K| > t_K) < \infty$, and so, by the Borel–Cantelli lemma, with probability 1, we have $|T_K| \leq t_K$ for all large enough K . Therefore,

$$P_n \leq \limsup_{K \rightarrow \infty} \frac{t_K}{K\Theta} = \frac{\varepsilon(1 + h)}{q}.$$

As this holds for any $\varepsilon > p_n$ and $h > 0$, we conclude that $P_n \leq p_n/q$, which completes the proof of the theorem.

5. Computing the stationary distribution

In the case of constant synaptic weights, it is not difficult to derive differential equations (and boundary conditions for them) for the components $\psi_{m,n}$ of the stationary density of Z . They can be derived from the general relation

$$\mathbb{E}Af(Z(\infty)) = 0, \tag{12}$$

where A is the infinitesimal generator of the process, f is a function from a suitable subset of the domain of A , and, as before, $Z(\infty) \sim \pi$. It may be easier, however, to obtain them via a direct argument, making use of our Theorem 2 (of which the conditions will be assumed to be satisfied in this section unless we explicitly state otherwise).

To show how to do this, we will first consider the simple case of a network with one external source and one neuron (with feedback). Suppose that the neuron has an absolute refractory period (so that the state space is actually finite dimensional). For simplicity, we assume throughout this section that $\Theta = 1$ (which clearly does not restrict generality).

In this case, the state space of the process is just $E \times E$, so that each state $(\hat{x}; \mathbf{x}) = (\hat{x}_1, \hat{x}_2, \dots, \hat{x}_m, 0, 0, \dots; x_1, x_2, \dots, x_n, 0, 0, \dots)$ (note that here we suppress the unnecessary first subscript indicating the number of the external sources or neurons; likewise, $\hat{\rho}$ will denote here $\hat{\rho}_1$, etc.) belongs to one of the components $E^{(m)} \times E^{(n)}$, $m, n \geq 0$. We will denote the respective density components by $\psi_{m,n}$.

The first density component, $\psi_{0,0}$, is just the stationary probability of the silent state, for which we have, as $\delta \searrow 0$,

$$\begin{aligned} \psi_{0,0} &= \mathbb{P}(Z(\delta) = (0, 0)) \\ &= \mathbb{P}(Z(\delta) = (0, 0) \mid Z(0) = (0, 0))\mathbb{P}(Z(0) = (0, 0)) \\ &\quad + \int_0^\delta \mathbb{P}(Z(\delta) = (0, 0) \mid Z(0) = (y, 0))\psi_{1,0}(y) \, dy \\ &\quad + \int_0^\delta \mathbb{P}(Z(\delta) = (0, 0) \mid Z(0) = (0, y))\psi_{0,1}(y) \, dy + O(\delta^2) \\ &= e^{-(\hat{\rho} + \zeta(v))\delta} \psi_{0,0} + \int_0^\delta (1 + o(1))\psi_{1,0}(y) \, dy + \int_0^\delta (1 + o(1))\psi_{0,1}(y) \, dy + o(\delta), \end{aligned}$$

where the $O(\delta^2)$ term corresponds to the possibility that $Z(0) \in E^{(m)} \times E^{(n)}$ with $m + n > 1$. From the above representation we obtain

$$(\hat{\rho} + \zeta(v))\psi_{0,0} = \psi_{1,0}(0) + \psi_{0,1}(0), \tag{13}$$

where, using Theorem 2, we put $\psi_{1,0}(0) := \psi_{1,0}(0+)$ and $\psi_{0,1}(0) := \psi_{0,1}(0+)$.

In the case where $mn > 0$, we fix a point $\mathbf{z} = (\hat{\mathbf{x}}; \mathbf{x})$ in the interior of $E_0^{(m)} \times E_0^{(n)}$ (see (4)) and set $I_{\mathbf{z}}(\delta) := I_{\hat{\mathbf{x}}}(\delta) \times I_{\mathbf{x}}(\delta)$, where

$$I_{\hat{\mathbf{x}}}(\delta) := (\hat{x}_1, \hat{x}_1 + \delta) \times \cdots \times (\hat{x}_m, \hat{x}_m + \delta),$$

$$I_{\mathbf{x}}(\delta) := (x_1, x_1 + \delta) \times \cdots \times (x_n, x_n + \delta),$$

and $\delta > 0$ is small enough so that $I_{\mathbf{z}}(\delta) \subset E_0^{(m)} \times E_0^{(n)}$. Using the notation $\mathbf{z} + \theta$ for shifting all the components of the vector \mathbf{z} by the same amount $\theta \in \mathbb{R}$, and, as before, slightly abusing notation by identifying $E^{(m)} \times E^{(n)}$ with $E_0^{(m)} \times E_0^{(n)}$, we have

$$\begin{aligned} \mathbb{P}(Z(\delta) \in I_{\mathbf{z}}(\delta)) &= \mathbb{P}(Z(\delta) \in I_{\mathbf{z}}(\delta) \mid Z(0) \in I_{\mathbf{z}}(\delta) + \delta)\mathbb{P}(Z(0) \in I_{\mathbf{z}}(\delta) + \delta) \\ &\quad + \mathbb{P}(Z(\delta) \in I_{\mathbf{z}}(\delta), Z(0) \in [(I_{\hat{\mathbf{x}}}(\delta) + \delta) \times (0, \delta)] \times (I_{\mathbf{x}}(\delta) + \delta)) \\ &\quad + \mathbb{P}(Z(\delta) \in I_{\mathbf{z}}(\delta), Z(0) \in (I_{\hat{\mathbf{x}}}(\delta) + \delta) \times [(I_{\mathbf{x}}(\delta) + \delta) \times (0, \delta)]) \\ &\quad + O(\delta^{m+n+2}) \quad \text{as } \delta \searrow 0, \end{aligned}$$

where the last term corresponds to the possibility of $Z(0)$ being in a space of dimensionality higher than $n + m + 1$. Expressing the probabilities above as integrals of the respective density components and using Theorem 2, we obtain the relation

$$\begin{aligned} &\int_{I_{\mathbf{z}}(\delta)} \psi_{m,n}(\mathbf{y})(\mu_m \otimes \mu_n)(d\mathbf{y}) \\ &= \int_{I_{\mathbf{z}}(\delta)} [1 - \delta(\hat{\rho} + R(\mathbf{y} + \delta))] \psi_{m,n}(\mathbf{y} + \delta)(\mu_m \otimes \mu_n)(d\mathbf{y}) \\ &\quad + (1 + o(1))\delta^{m+n+1}[\psi_{m+1,n}(\hat{\mathbf{y}}; 0; \mathbf{y}) + \psi_{m,n+1}(\hat{\mathbf{y}}; \mathbf{y}; 0)] + O(\delta^{m+n+2}). \end{aligned}$$

Subtracting from both sides the integral of $\psi_{m,n}(\mathbf{y} + \delta)$ over $I_{\mathbf{z}}$ we can then easily verify that the relation implies that the following differential equation must be satisfied: for $\theta \in (0, 1 - \max\{\hat{x}_1, x_1\})$,

$$\begin{aligned} \frac{\partial}{\partial \theta} \psi_{m,n}(\mathbf{z} + \theta) &= (\hat{\rho} + R(\mathbf{z} + \theta))\psi_{m,n}(\mathbf{z} + \theta) - \psi_{m+1,n}(\hat{\mathbf{x}} + \theta, 0; \mathbf{x} + \theta) \\ &\quad - \psi_{m,n+1}(\hat{\mathbf{x}} + \theta; \mathbf{x} + \theta, 0). \end{aligned} \tag{14}$$

Of course, the equation will hold along the whole interval formed by the intersection of $E_0^{(m)} \times E_0^{(n)}$ with the straight line passing through the point \mathbf{z} and having the directional vector $\mathbf{e}_{m+n} := (1, \dots, 1) \in \mathbb{R}^{m+n}$, the boundary condition at its right point being specified by the rates of transition to $E_0^{(m)} \times E_0^{(n)}$ from the state space components of lower dimensionalities. For example, if $\hat{x}_1 < x_1$ then the right endpoint for the interval of validity of (14) corresponds to the point where the ray $\mathbf{x} + \theta \mathbf{e}_n$, $\theta > 0$, hits the ‘right’ face of $E_0^{(n)}$ (the point $\hat{\mathbf{x}} + \theta \mathbf{e}_m$ still being in the interior of $E_0^{(m)}$). At that location, the system can only enter the component $E_0^{(m)} \times E_0^{(n)}$ by a jump from

$$(\hat{\mathbf{x}} + 1 - x_1; \mathbf{x}^* + 1 - x_1) \in E_0^{(m)} \times E_0^{(n-1)}, \quad \text{where } \mathbf{x}^* := (x_2, x_3, \dots, x_n),$$

caused by a new spike fired by the neuron. Using a probabilistic argument similar to that above, it is easy to see that the following must hold:

$$\psi_{m,n}(z + 1 - x_1) = R(\hat{x} + 1 - x_1; \mathbf{x}^* + 1 - x_1)\psi_{m,n-1}(\hat{x} + 1 - x_1; \mathbf{x}^* + 1 - x_1). \tag{15}$$

A similar equation will hold in the case where $\hat{x}_1 > x_1$, but then the coefficient of $\psi_{m-1,n}$ on the right-hand side of the respective relation will simply be $\hat{\rho}$. The case where only one of m, n is 0 is treated similarly.

Solving equations of the form (14) with boundary conditions (15), complemented by (13) and the condition that $\sum_{m,n} \int \psi_{m,n} d(\mu_m \otimes \mu_n) = 1$, is hardly possible except for the simplest cases. One such case is considered in the following example.

Example 1. Consider the case of a single neuron with feedback and no external sources. Moreover, assume that the firing rate function has the property

$$R(x) = 0 \quad \text{for all } x \in E^{(n)}, n \geq 2, \tag{16}$$

so that there cannot be more than two spikes in any given time interval of length $\Theta = 1$ (say, due to the length of the absolutely refractory period exceeding $\frac{1}{2}$). Thus, the state space of the system is just $E^{(0)} \times E^{(1)} \times E^{(2)}$ (which we again can and will identify with $E_0^{(0)} \times E_0^{(1)} \times E_0^{(2)}$); the density components being $\psi_n, n = 0, 1, 2$ (for $n > 2$, all $\psi_n \equiv 0$).

Using an obvious notational convention, we see that an analog of (13) in this case has the form

$$R(0)\psi_0 = \psi_1(0), \tag{17}$$

while an analog of (14) is, in the case $n = 1$,

$$\frac{d\psi_1(\theta)}{d\theta} = R(\theta)\psi_1(\theta) - \psi_2(\theta, 0), \quad \theta \in (0, 1), \tag{18}$$

with the boundary condition (an analog of (15))

$$\psi_1(1) = R(0)\psi_0. \tag{19}$$

When $n = 2$, an analog of (14) has the form

$$\frac{\partial \psi_2(y + \theta, \theta)}{\partial \theta} = R(y + \theta, \theta)\psi_2(y + \theta, \theta) - \psi_3(y + \theta, \theta, 0) \equiv 0, \quad \theta \in (0, 1 - y),$$

for any $y \in (0, 1)$, the right-hand side of the equation being zero owing to (16), with the boundary condition (again an analog of (15))

$$\psi_2(1, 1 - y) = R(1 - y)\psi_1(1 - y). \tag{20}$$

The last two relations immediately imply that, for any $y \in (0, 1)$,

$$\psi_2(y + \theta, \theta) = R(1 - y)\psi_1(1 - y), \quad \theta \in (0, 1 - y).$$

Therefore, $\psi_2(\theta, 0) = R(1 - \theta)\psi_1(1 - \theta)$, so that (18) becomes

$$\frac{d\psi_1(\theta)}{d\theta} = R(\theta)\psi_1(\theta) - R(1 - \theta)\psi_1(1 - \theta), \quad \theta \in (0, 1).$$

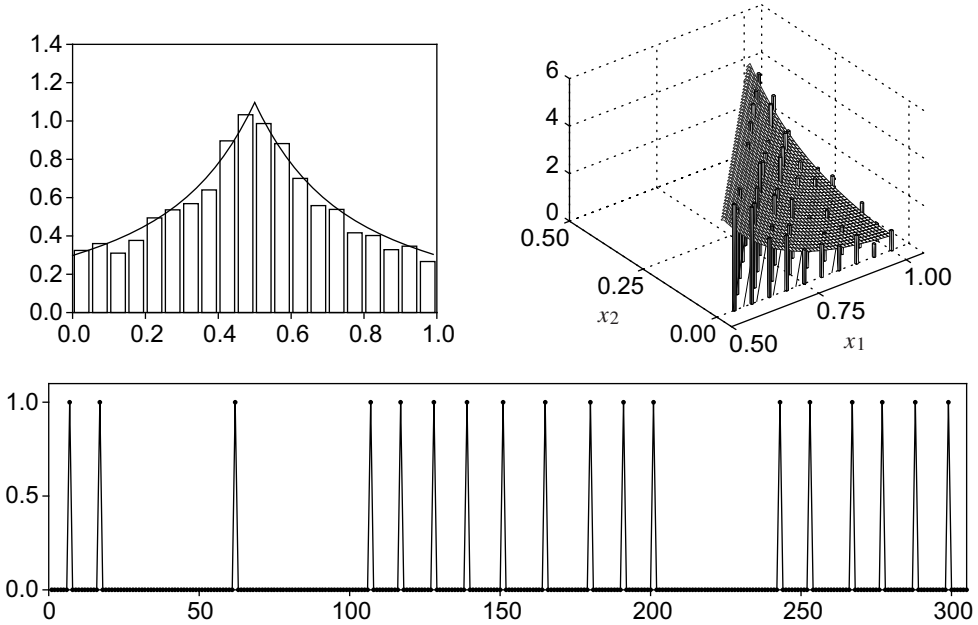


FIGURE 1: Numerical illustration to Example 1. The top row displays the stationary densities $\psi_1(x_1)$, $x_1 \in (0, 1)$ (left), and $\psi_2(x_1, x_2)$, $1 \geq x_1 > x_2 > 0$ (right), together with their histogram estimates from a simulation. The bottom pane depicts a simulation of the neuron’s activity during a time interval of length 15 (comprising 300 discrete time steps of length 0.05).

This means that the function ψ_1 is symmetric about the point $\theta = \frac{1}{2}$, so that $\psi_1(\theta) = \psi_1(1 - \theta)$, $\theta \in (0, 1)$ (hence, conditions (17) and (19) are consistent), and the last differential equation can be rewritten as

$$\frac{d\psi_1(\theta)}{d\theta} = (R(\theta) - R(1 - \theta))\psi_1(\theta), \quad \theta \in (0, 1). \tag{21}$$

Setting $\varphi(\theta) := \exp\{\int_0^\theta (R(y) - R(1 - y)) dy\}$, we derive from (17) and (21) that

$$\psi_1(\theta) = R(0)\psi_0\varphi(\theta), \quad \theta \in (0, 1).$$

Together with (20) this completely specifies the density function ψ (computing ψ_0 is trivial).

To numerically illustrate this example, take the PSP kernel of the form

$$\epsilon(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)}, \quad x \in (0, 1),$$

with $\alpha = 1.5$ and $\beta = 3$, B being the beta function, so that the direct effect of a spike vanishes after $\Theta = 1$ time units, and the activation function $\zeta(x) = 6/(1 + \exp(1 - x))$, $x > 0$. With parameter values $v_0 = 0.3$ and $W = 1$, the instantaneous firing rate of the neuron now has the form

$$R(x) = \zeta(0.3 + \epsilon(1 - x))r(1 - x), \quad x \in (0, 1),$$

where $r(\cdot)$ is defined in (2) with $\delta_{AR} = \frac{1}{2}$, which is consistent with the restriction that the neuron can fire at most twice during a unit time interval. Figure 1 depicts the density components ψ_1

and ψ_2 together with their histogram estimators from a simulated run of the network, and a typical simulated realisation of the neuron’s activity. For simulation purposes, we discretised time using a step size of 0.05 (the bottom pane in Figure 1 shows the first 300 time steps). The neuron’s activity was simulated for 2×10^5 time steps, and the histograms were derived from the 10^4 disjoint ‘time windows’ (20 time steps each) into which the total time interval was partitioned. The histograms are in very good agreement with the theoretical stationary densities. The estimated stationary probability to be in the ‘silent state’ is $\hat{\psi}_0 = 0.136$, whereas the true probability thereof is $\psi_0 = 0.149$.

Example 2. One can also obtain a closed-form solution in the case of a single neuron with feedback and no absolutely refractory period, but under the special assumption that the neuron’s PSP kernel is exponential: $\epsilon(t) = e^{-\alpha t} \mathbf{1}(t \geq 0)$ for some $\alpha > 0$. Of course, this violates the basic assumption that ϵ has a finite support and is smooth, but such a simplified mathematical model could still provide a useful approximation to more realistic models.

It is not hard to see that in this case the dynamics of the system can be described by a univariate Markov process $Y(t) := \sum_n \epsilon(t - T_n) > 0$. Assuming without loss of generality that $\alpha = 1$, we can see that the process Y is driven by the Ornstein–Uhlenbeck-type equation

$$dY(t) = -Y(t) dt + dJ(t), \quad t > 0,$$

$J(t)$ being a pure jump process with unit jumps and instantaneous jump rate $\gamma(Y(t))$, where $\gamma(y) := \zeta(v + Wy)$ and W is the weight of the ‘self-connection’ of our neuron. The infinitesimal generator A of the process Y is clearly

$$Af(x) = -xf'(x) + \gamma(x)(f(x + 1) - f(x)), \quad x > 0; \tag{22}$$

its domain’s description can be found in, e.g. Proposition 4 of [20]. It is not hard to see that Y is ergodic (see, e.g. [4]) and so has a unique stationary distribution that we will again denote by π . Substituting (22) into (12) (with $Z(\infty)$ replaced by $Y(\infty)$) yields

$$\int_0^\infty yf'(y)\pi(dy) = \int_0^\infty \gamma(x) \left(\int_x^{x+1} f'(y) dy \right) \pi(dx).$$

A routine calculation now leads to

$$\int_0^\infty yf'(y)\pi(dy) = \int_0^\infty f'(y)b(y) dy, \quad b(y) := \int_{(y-1)^+}^y \gamma(x)\pi(dx),$$

where we have used the notation $x^+ := \max\{x, 0\}$ for the positive part of x . As this equation holds for a large enough class of functions f (see, e.g. [20]), and b is continuous and locally bounded, we conclude that π has a locally bounded and continuous density ψ on $(0, \infty)$, satisfying

$$y\psi(y) = \int_{(y-1)^+}^y \gamma(x)\psi(x) dx, \quad y > 0.$$

This equation can be solved recursively, on intervals $I_n := (n, n + 1)$, $n \geq 0$. Straightforward calculations show that the stationary density is given by

$$\psi(y) := \varphi_n(y), \quad y \in I_n, \quad n = 0, 1, 2, \dots,$$




FIGURE 2: Numerical illustration to Example 2. The top row displays the stationary densities ψ in the cases $v = 2$ (left) and $v = 0.1$ (right), together with their histogram estimates from simulations. The bottom pane depicts typical simulated trajectories of the process Y in the respective cases.

where the φ_n are found recursively as

$$\begin{aligned} \varphi_0(y) &= \psi(1) \exp\left(\int_1^y \frac{\gamma(x) - 1}{x} dx\right), \quad y \in I_0, \\ \varphi_n(y) &= \Phi_n[\varphi_{n-1}](y), \quad y \in I_n, \quad n \geq 1, \end{aligned}$$

where, for f defined on I_{n-1} with a finite $f(n-)$, we denote by $\Phi_n[f](y)$, $y \in I_n$, the solution ϕ of the problem

$$\phi'(y) = \frac{\gamma(y) - 1}{y} \phi(y) - \frac{\gamma(y - 1)}{y} f(y), \quad y \in I_n, \quad \phi(n) = f(n-).$$

The only unknown constant $\psi(1)$ is just the normalizing factor that is to be determined from $\int_0^\infty \psi(y) dy = 1$. At the integer points the density ψ can be defined by continuity.

To numerically illustrate Example 2, we chose $\zeta(y) = 2/(1 + \exp\{1 - y\})$ and $W = 1$, and considered two different values for v : 2 and 0.1. The top panes in Figure 2 depict the plots of the stationary densities of $Y(t)$ in these two cases, together with their histogram estimators. The histograms were constructed using the values of 10^4 simulations of the discrete-time versions of the processes after 5×10^3 discrete time steps of length 0.02. The bottom panes show typical realisations of the simulated processes. Note that, in the case of low background activity $v = 0.1$, the stationary density ‘explodes’ at 0 (as $\zeta(0.1) \approx 0.58 < 1$; cf. the expression for φ_0), reflecting the fact that the process can stay relatively long in a ‘quiet state’ in that case. When $v = 2$, we have $\zeta(2) \approx 1.46 > 1$, implying that the stationary density of the process $Y(t)$ does not often visit the vicinity of 0.

In the general case, we can only hope to compute approximations to the stationary distribution of the network. One way to do that is to discretise the state space and approximate the differential equations for the density components $\psi_{m,n}$ discussed at the beginning of this section with respective difference equations, and then to solve the latter. However, although the existence of a

solution to the original system of differential equations follows from Theorem 2, establishing its uniqueness and also the convergence of the solutions to the systems of approximating difference equations presents a challenge.

An alternative approach is to first approximate the original stochastic process with a sequence of finite ergodic Markov chains in discrete time, and then to prove convergence of their stationary distributions to the desired stationary distribution of U . For a detailed description of the approximating Markov chains, the systems of equations their stationary distributions satisfy, and the proof of (weak) convergence of those distributions to the stationary distribution of U , the interested reader is referred to Section 5 of [5].

Acknowledgements

This research was supported by the ARC Discovery Grant DP120102398 and the ARC Centre of Excellence for Mathematics and Statistics of Complex Systems. The authors are grateful to the anonymous referee for useful comments that helped to improve the exposition of the paper.

References

- [1] BEAR, M. F., CONNORS, B. W. AND PARADISO, M. A. (2007). *Neuroscience: Exploring the Brain*, 3rd edn. Lippincott Williams & Wilkins, Philadelphia, PA.
- [2] BOROVKOV, A. A. (1998). *Ergodicity and Stability of Stochastic Processes*. John Wiley, Chichester.
- [3] BOROVKOV, A. A. (2013). *Probability Theory*. Springer, London.
- [4] BOROVKOV, K. AND LAST, G. (2008). On level crossings for a general class of piecewise-deterministic Markov processes. *Adv. Appl. Prob.* **40**, 815–834.
- [5] BOROVKOV, K., DECROUEZ, G. AND GILSON, M. (2012). On stationary distributions of stochastic neural networks. Preprint. Available at <http://uk.arxiv.org/abs/1206.4489>.
- [6] BRÉMAUD, P. (1981). *Point Processes and Queues*. Springer, New York.
- [7] BRÉMAUD, P. AND MASSOULIÉ, L. (1996). Stability of nonlinear Hawkes processes. *Ann. Prob.* **24**, 1563–1588.
- [8] BRÉMAUD, P. AND MASSOULIÉ, L. (2001). Hawkes branching point processes without ancestors. *J. Appl. Prob.* **38**, 122–135.
- [9] BRÉMAUD, P. AND MASSOULIÉ, L. (2002). Power spectra of general shot noises and Hawkes point processes with a random excitation. *Adv. Appl. Prob.* **34**, 205–222.
- [10] BRILLINGER, D. R. (1975). The identification of point process systems. *Ann. Prob.* **3**, 909–929.
- [11] BRILLINGER, D. R. (1988). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biol. Cybernetics* **59**, 189–200.
- [12] BURKITT, A. N., GILSON, M. AND VAN HEMMEN, J. L. (2007). Spike-timing-dependent plasticity for neurons with recurrent connections. *Biol. Cybernetics* **96**, 533–546.
- [13] CHORNOBOY, E. S., SCHRAMM, L. P. AND KARR, A. F. (1988). Maximum likelihood identification of neural point process systems. *Biol. Cybernetics* **59**, 265–275.
- [14] DAYAN, P. AND ABBOTT, L. F. (2001). *Theoretical Neuroscience. Computational and Mathematical Modeling of Neural Systems*. MIT Press, Cambridge, MA.
- [15] GERSTNER, W. AND KISTLER, W. M. (2002). *Spiking Neuron Models. Single Neurons, Populations, Plasticity*. Cambridge University Press.
- [16] GILSON, M. (2009). Biological learning mechanisms in spiking neuronal networks. Doctoral Thesis, The University of Melbourne. Available at <http://www.t35.ph.tum.de/addons/publications/Gilson-2009.pdf>.
- [17] GILSON, M. *et al.* (2009). Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks III: Partially connected neurons driven by spontaneous activity. *Biol. Cybernetics* **101**, 411–426.
- [18] GILSON, M. *et al.* (2009). Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks IV. *Biol. Cybernetics* **101**, 427–444.
- [19] GILSON, M., BURKITT, A. AND VAN HEMMEN, J. L. (2010). STDP in recurrent neuronal networks. *Front. Comput. Neurosci.* **4**, 23.
- [20] HARRISON, J. M. AND RESNICK, S. I. (1976). The stationary distribution and first exit probabilities of a storage process with general release rule. *Math. Operat. Res.* **1**, 347–358.
- [21] HAWKES, A. G. (1971). Point spectra of some mutually exciting point processes. *J. R. Statist. Soc. B* **33**, 438–443.
- [22] IZHIKEVICH, E. M. (2007). *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. MIT Press, Cambridge, MA.

- [23] JACOBSEN, M. (2006). *Point Process Theory and Applications: Marked Point and Piecewise Deterministic Processes*. Birkhäuser, Boston, MA.
- [24] KEMPTER, R., GERSTNER, W. AND VAN HEMMEN, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* **59**, 4498–4514.
- [25] KOCH, C. AND IDAN, S. (eds) (1998). *Methods in Neuronal Modeling: From Ions to Networks*, 2nd edn. MIT Press, Cambridge, MA.
- [26] MASSOULIÉ, L. (1998). Stability results for a general class of interacting point processes dynamics, and applications. *Stoch. Process. Appl.* **75**, 1–30.
- [27] NICHOLLS, J. G. *et al.* (2012). *From Neuron to Brain*, 5th edn. Sinauer Associates, Sunderland, MA.
- [28] PANINSKI, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Comput. Neural Systems* **15**, 243–262.
- [29] PILLOW, J. W., AHMADIAN, Y. AND PANINSKI, L. (2011). Model-based decoding, information estimation, and change-point detection techniques for multineuron spike trains. *Neural Comput.* **23**, 1–45.
- [30] PILLOW, J. W. *et al.* (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999.
- [31] SHEPHERD, G. M. (ed.) (2004). *The Synaptic Organization of the Brain*, 5th edn. Oxford University Press.
- [32] SHEPHERD, G. M. AND GRILLNER, S. (eds) (2010). *Handbook of Brain Microcircuits*. Oxford University Press.
- [33] SPORNS, O. (2011). *Networks of the Brain*. MIT Press, Cambridge, MA.
- [34] STEVENSON, I. H. *et al.* (2009). Bayesian inference of functional connectivity and network structure from spikes. *IEEE Trans. Neural Systems Rehabil. Eng.* **17**, 203–213.
- [35] TRUCCOLO, W. *et al.* (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.* **93**, 1074–1089.