UC Santa Cruz
UC Santa Cruz Previously Published Works

Title

Model-agnostic neural mean field with a data-driven transfer function

Permalink

Journal

ISSN

Authors

Spaeth, Alex
Haussler, David
Teodorescu, Mircea

Publication Date

DOI

Peer reviewed

NEUROMORPHIC
Computing and Engineering

CrossMark

PAPER

Model-agnostic neural mean field with a data-driven transfer function

Alex Spaeth[1,2,*] ⓘ, David Haussler[2,3] and Mircea Teodorescu[1,2,3] ⓘ

1 Electrical and Computer Engineering Department, University of California, Santa Cruz, Santa Cruz, CA, United States of America
2 Genomics Institute, University of California, Santa Cruz, Santa Cruz, CA, United States of America
3 Biomolecular Engineering Department, University of California, Santa Cruz, Santa Cruz, CA, United States of America
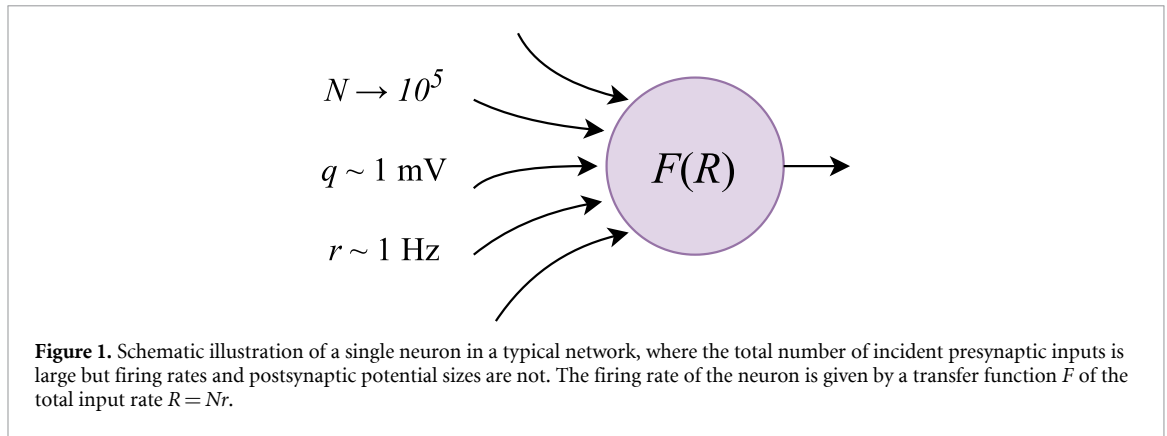* Author to whom any correspondence should be addressed.

E-mail: atspaeth@ucsc.edu

Abstract

As one of the most complex systems known to science, modeling brain behavior and function is both fascinating and extremely difficult. Empirical data is increasingly available from *ex vivo* human brain organoids and surgical samples, as well as *in vivo* animal models, so the problem of modeling the behavior of large-scale neuronal systems is more relevant than ever. The statistical physics concept of a mean-field model offers a tractable way to bridge the gap between single-neuron and population-level descriptions of neuronal activity, by modeling the behavior of a single representative neuron and extending this to the population. However, existing neural mean-field methods typically either take the limit of small interaction sizes, or are applicable only to the specific neuron models for which they were derived. This paper derives a mean-field model by fitting a transfer function called Refractory SoftPlus, which is simple yet applicable to a broad variety of neuron types. The transfer function is fitted numerically to simulated spike time data, and is entirely agnostic to the underlying neuronal dynamics. The resulting mean-field model predicts the response of a network of randomly connected neurons to a time-varying external stimulus with a high degree of accuracy. Furthermore, it enables an accurate approximate bifurcation analysis as a function of the level of recurrent input. This model does not assume large presynaptic rates or small postsynaptic potential size, allowing mean-field models to be developed even for populations with large interaction terms.

1. Introduction

The brain is one of the most complex systems known to science, which makes the problem of computationally modeling its behavior and function both fascinating and extremely difficult. The computational substrate of the brain consists of billions of neurons, each with tens of thousands of connections spanning neurons from the local region to distant areas [1]. Computational neuroscientists are thus faced with a problem of scale: while the biophysical behavior of the individual neuron is understood from first principles [2, 3], and large numbers of neurons have been modeled phenomenologically as a dynamical system [4, 5], unifying different scales of description remains a fundamental difficulty and the subject of ongoing research [6, 7]. Today, as neuromorphic systems are emerging for computation and robotic control [8–13] and human brain organoids and surgical samples are becoming key systems for studying disease states and neuronal computation [14–19], the question of modeling the behavior of large-scale neuronal systems is becoming all the more crucial.

One of the most popular approaches to unifying model scales is to make the mean-field assumption, namely that all the neurons in the population under consideration, as well as their inputs, are independent and statistically identical [20, 21]. This assumption is applicable to a variety of systems, ranging from randomly connected sparse networks [20] to modular networks of biological interest [22]. When

Figure 1. Schematic illustration of a single neuron in a typical network, where the total number of incident presynaptic inputs is large but firing rates and postsynaptic potential sizes are not. The firing rate of the neuron is given by a transfer function F of the total input rate $R = Nr$.

connections are not strongly selective, neurons within a population are essentially equivalent and can be represented by one population firing rate. Even networks with all-to-all connectivity can be treated, provided the neurons are stochastic enough to remain approximately independent [23]. Under the mean-field assumption, the members of a neuronal population act as independent samples of a shared statistical distribution; this reduces many problems to the study of the firing rate of a single representative neuron receiving N inputs at an average rate r as depicted in figure 1. Similar approaches have been used extensively in statistical physics, for example in the study of coupled oscillators [24, 25].

Analytical mean-field models can also be derived for some neuronal systems, such as networks of rate neurons [26] or of stochastic spiking neurons characterized by a renewal process [27]. For dynamical spiking neurons, on the other hand, no such analytical approach is known. Thus current mean-field models examine the limit of an infinite number of presynaptic neurons with infinitesimal postsynaptic potentials, and apply the central limit theorem to replace these inputs with a single white Gaussian noise process. This so-called *diffusion approximation* transforms the single-neuron dynamics into a drift-diffusion problem, for which many techniques have been developed [28, 29]. Perturbation analysis can model complex interactions, but is only valid for small effects [30, 31]. For more general study of effects such as spike frequency adaptation (SFA) or synaptic dynamics, the Fokker–Planck equation can be solved to get the time evolution of the probability distribution of the neuron's membrane voltage under the drift-diffusion process [23, 32]. However, each solution is particular to a single neuronal model, and calculating solutions is computationally expensive even when mathematical analysis has revealed an explicit model [33].

One popular basis for constructing mean-field models is a neuron's 'transfer function', which represents its firing rate as a function of the amount of input noise, analogous to the firing rate curve used to study the deterministic dynamics of an individual neuron [20, 21]. Analytical transfer functions are known for a variety of different neuron types and noise models [34–36]. However, we are not aware of any mean-field model making use of a more general noise assumption than the simple diffusion approximation.

Mean-field models based on the diffusion approximation and the approximate transfer function of [20] have been applied to a variety of different model neurons [37], have been studied in terms of bifurcation theory [38, 39], and have been compared to biological data [22]. However, the sigmoidal shape of this transfer function does not match analytical results [36, 40]. It also contains an arbitrary timescale T, which must be large enough for a quasistatic assumption to hold, but smaller than the minimum interspike interval (ISI) in the system [41]. Furthermore, it is explicitly assumed through moment closure that the membrane potential itself is normally distributed as well [42], which is inaccurate for realistic noise sources [43].

In this paper, we propose a novel, explicitly phenomenological approach based on numerically fitting a simple parametrized transfer function that can match both the shape of analytical and observed transfer functions and simulation results. It is also inexpensive to evaluate, as it has only four parameters, which are fitted with a good level of accuracy based on only a few simulations of single uncoupled neurons. Being based on a transfer function, this model is still limited by the assumption of a stationary firing rate distribution, but it makes no assumption about the form or dynamics of the membrane potential distribution. In particular, we do not require the diffusion approximation and can describe populations of neurons with finite numbers of connections and finitely large postsynaptic potentials. We numerically solve a consistency condition to find stable firing rates of a neuronal population, achieving good accuracy in a wide range of simulation examples. We demonstrate the effectiveness of our model by predicting a saddle-node bifurcation to bistability as a function of the connectivity parameter. Furthermore, we implement a popular first-order dynamical model and demonstrate its predictive power alongside its limitations.

2. Method

There exists a wide variety of dynamical spiking neural models. All are excitable, in the sense that increasing membrane voltage past a threshold causes a rapid upstroke followed by a return to equilibrium, called a spike [44]. However, their underlying dynamics vary significantly, and this variation can have important computational implications [45]. This paper includes results for three popular neuron models: the leaky integrate-and-fire (LIF), Izhikevich (Izh), and Hodgkin–Huxley (HH) models.

2.1. Leaky integrate-and-fire model

The dynamics of the LIF neuron are given in equation (1). The membrane potential of this neuron undergoes a drift back towards its resting state, while being driven by a time-dependent input $\xi(t)$. It is also common to introduce a refractory period T_{ref}, a duration for which the state of the neuron does not change after every spike.

$$
\begin{aligned}
&\text{When } v < V_{\text{thresh}} \\
&\dot{v} = \frac{1}{\tau}\left(V_{\text{rest}} - V\right) + \xi(t) \\
&\text{When } v = V_{\text{thresh}} \\
&v \leftarrow V_{\text{rest}} (\text{then wait } T_{\text{ref}} \text{ before continuing to integrate}).
\end{aligned}
\tag{1}
$$

Synaptic input to the neuron is represented by a random process $\xi(t)$. It is common in models to assume that each presynaptic neuron fires as a Poisson point process [46]. However, it is well known that individual neurons frequently deviate from Poisson statistics [47–50]. Fortunately, the superposition of N independent renewal processes with total rate R tends in the large-N limit to a Poisson process, without requiring that the input renewal processes be Poisson themselves, or even identically distributed [51]. Therefore, our model makes the somewhat more general assumption that each presynaptic neuron fires as an independent renewal process.

Each presynaptic firing event is modeled as a Dirac delta which, when integrated, instantaneously increases the membrane voltage by its fixed postsynaptic potential (PSP) value q. A typical order of magnitude for these parameters in the neocortex is 10^4 neurons [1] firing at a rate of several hertz [20], with PSP magnitudes of about one millivolt [52], as in figure 1.

2.2. Izhikevich and Hodgkin–Huxley models

Besides the LIF neuron, we also consider the Izhikevich and Hodgkin–Huxley models. The Izhikevich model is a quadratic integrate-and-fire neuron augmented with an adaptation variable u [53]. The result is a model not substantially more computationally expensive than LIF, but with much more realistic spike shapes. On the other hand, the HH model is based on biophysically modeling the dynamics of the membrane potential [2]. This results in a more realistic but much more complex model. See appendix D for dynamical equations as well as a more thorough comparison of the dynamical features of all three models.

The Izhikevich model uses Dirac delta synapses. However, delta synapses are not supported by the HH neuron available in the simulator. Instead, alpha postsynaptic currents are used, where the weight w of a synapse gives the peak value of the postsynaptic current $\alpha(t) = wte^{1-t/\tau_{\text{syn}}}/\tau_{\text{syn}}$, which is integrated by a membrane capacitance C_m. To keep inputs comparable between neuron models, we compute w as a function of an equivalent PSP size q by fixing its integral so that it does not depend on the synaptic time constant and the zero-τ_{syn} limit corresponds to a delta PSP of size q:

$$
C_m q = \int_0^\infty \alpha(t)\, dt = we\tau_{\text{syn}} \quad \implies \quad w = \frac{C_m q}{e\tau_{\text{syn}}}.
\tag{2}
$$

2.3. Excitatory-inhibitory balance

Although it is convenient to assume that the inputs to the neuron are as homogeneous as possible, models generally assume a distinction between excitatory and inhibitory inputs, which differ at least in the sign of their PSPs. In our simulations, a fraction η of the input connections are from excitatory neurons with PSP amplitude $q_e > 0$, and the remainder of the input is inhibitory, with a negative PSP with absolute value q_i. This applies both to the background input, where the input to each neuron is drawn from two independent Poisson sources with rates ηR and $(1 - \eta)R$ corresponding to its own excitatory and inhibitory background, as well as to simulations with recurrent connectivity, where each neuron is either excitatory or inhibitory and produces the corresponding PSP in all of its postsynaptic neurons.

An exact transfer function is still unknown for even the simple LIF neuron under Poisson inputs [54]. However, it has long been known that in the diffusion limit of infinitely large total input rate R and small PSPs, the membrane voltage of the LIF neuron follows an Ornstein–Uhlenbeck process with time constant equal to the membrane time constant τ [34]. For a given excitatory fraction η, the drift rate μ and the diffusion constant D of this Ornstein–Uhlenbeck process would be given by:

$$
\begin{aligned}
\mu &= \eta R q_e - (1-\eta) R q_i \\
D &= \eta R q_e^2 + (1-\eta) R q_i^2 .
\end{aligned}
\tag{3}
$$

We are interested in modeling neurons in the *asynchronous irregular* state observed in biological systems, which computational experiments have shown occurs when spiking is driven by fluctuations rather than drift [39, 47, 55, 56]. The condition where mean synaptic input $\mu = 0$, is called loose excitatory-inhibitory (EI) balance and is widely observed in experiments and commonly assumed in models [41, 57]. Setting $\mu = 0$ fixes the inhibitory PSP size q_i as a function of q_e:

$$
\mu = 0 \quad \implies \quad q_i = \frac{q_e \eta}{1-\eta} .
\tag{4}
$$

In the diffusion limit, any inputs that yield the same value of D will behave identically, regardless of the value of η. We do not expect this to hold exactly because we are using finite values of R and q, but it remains a convenient parametrization. Therefore, we use equation (4) and (3) to define an effective diffusion coefficient for the balanced input regime as follows:

$$
D = q^2 R \quad \text{where} \quad q = q_e \sqrt{\frac{\eta}{1-\eta}} = q_i \sqrt{\frac{1-\eta}{\eta}} .
\tag{5}
$$

Throughout this paper, we use the value $\eta = 0.8$, as is commonly assumed in computational models [53, 55]. This results in excitatory and inhibitory PSP sizes that differ by a factor of 4. For example, for $q = 1\,\mathrm{mV}$, we have $q_e = 0.5\,\mathrm{mV}$ and $q_i = 2\,\mathrm{mV}$.

2.4. Refractory SoftPlus

The statistics of interspike intervals (ISIs) in the diffusion limit of equation (3) are found by solving a first passage time problem [34], but these results are mathematically complex, and do not apply to other neuron models, or even outside the diffusion limit (see appendix A).

For this reason, it is appealing to derive a mean-field model from an approximate transfer function that can be fitted numerically to multiple neuron models. In this section, we propose one called 'Refractory SoftPlus' after a function occasionally used in machine learning [58]. SoftPlus is a smooth rectified linear function of one variable whose sharpness is controlled by a shape parameter β. It generalizes the ramp function $x \mapsto \max(x, 0)$, to which it converges in the infinite-β limit.
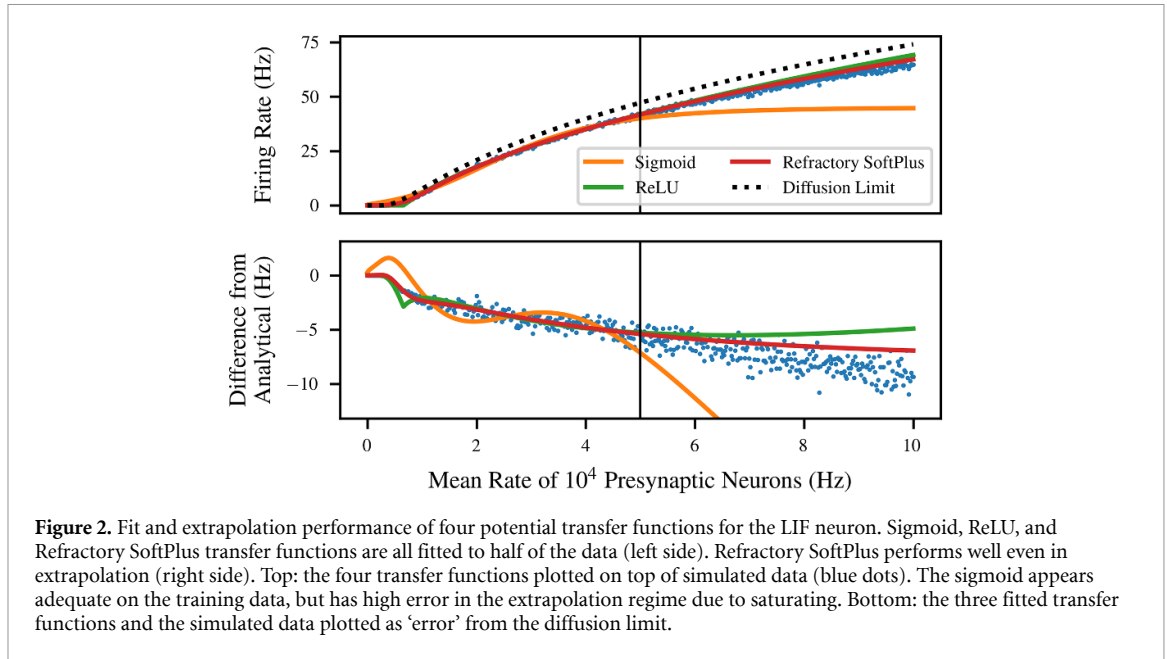
$$
\mathrm{SoftPlus}(x; \beta) = \beta^{-1} \ln\left(1 + e^{\beta x}\right) .
\tag{6}
$$

To define the Refractory SoftPlus transfer function, we augment SoftPlus with scale and shift parameters α and σ_0. The input to SoftPlus is taken to be $\sqrt{D} = q\sqrt{R}$ because the firing rate of LIF neuron in the diffusion limit scales as the square root of the diffusion coefficient [34]. If the value of q is not available, it can be taken equal to unity, which will change the fitted parameter values, but not the results. For a transfer function which handles variable q, see appendix B.

We additionally model an absolute refractory period of duration T_{ref} by inverting the output of SoftPlus to produce an ISI, adding the refractory period, then inverting again. This yields the following transfer function as a function of R with four shape parameters:

$$
S_{\mathrm{ref}}\left(R; \alpha, \beta, \sigma_0, T_{\mathrm{ref}}\right) = \left(T_{\mathrm{ref}} + \alpha\, \mathrm{SoftPlus}\left(q\sqrt{R} - \sigma_0\; ; \beta\right)^{-1}\right)^{-1} .
\tag{7}
$$

For any given neuron model, and for fixed values of q and η, we compute the firing rate of the neuron over 100 simulated seconds for a range of 500 different values of R. For maximum comparability between different values of q, we use evenly spaced R calculated so that D ranges from 0 to 100 $\mathrm{mV}^2\,\mathrm{s}^{-1}$. This cutoff point is arbitrary and our results do not depend on it; it is chosen only because it tends to drive our simulated neurons to fire in the tens of hertz, so that most of the shape of their transfer functions can be observed. The four parameters α, β, σ_0, and T_{ref} are then optimized to fit this simulated transfer function via nonlinear least squares. To make performance more comparable across models and conditions, we report

Figure 2. Fit and extrapolation performance of four potential transfer functions for the LIF neuron. Sigmoid, ReLU, and Refractory SoftPlus transfer functions are all fitted to half of the data (left side). Refractory SoftPlus performs well even in extrapolation (right side). Top: the four transfer functions plotted on top of simulated data (blue dots). The sigmoid appears adequate on the training data, but has high error in the extrapolation regime due to saturating. Bottom: the three fitted transfer functions and the simulated data plotted as 'error' from the diffusion limit.

fitting error normalized by the maximum firing rate observed in the simulation. For further details on numerical methods, see appendix C.

Figure 2 shows an example of the performance of this transfer function in comparison to transfer functions from the literature. The simulation data is from the LIF neuron of equation (1), with $V_{\text{thresh}} - V_{\text{rest}} = 15\,\text{mV}$, $\tau = 10\,\text{ms}$, and $T_{\text{ref}} = 2\,\text{ms}$, subject to Poisson inputs with $\eta = 0.8$ and $q = 1\,\text{mV}$ (giving $q_e = 0.5\,\text{mV}$ and $q_i = 2\,\text{mV}$). It is compared to the analytical solution for the diffusion limit [34] as well as ReLU and sigmoidal transfer functions. These are simply univariate functions of $\sqrt{D}$ augmented with scale and shift parameters (see appendix C); for the ReLU, this is the ramp function defined above [58], whereas for the sigmoidal transfer function it is the hyperbolic tangent [21]. In each case except the analytical solution, the transfer function is fitted to the left half of the data, i.e.the region where $R < 50\,\text{kHz}$. The full range of values is then used to show not only the performance of the original fit but also its ability to extrapolate.
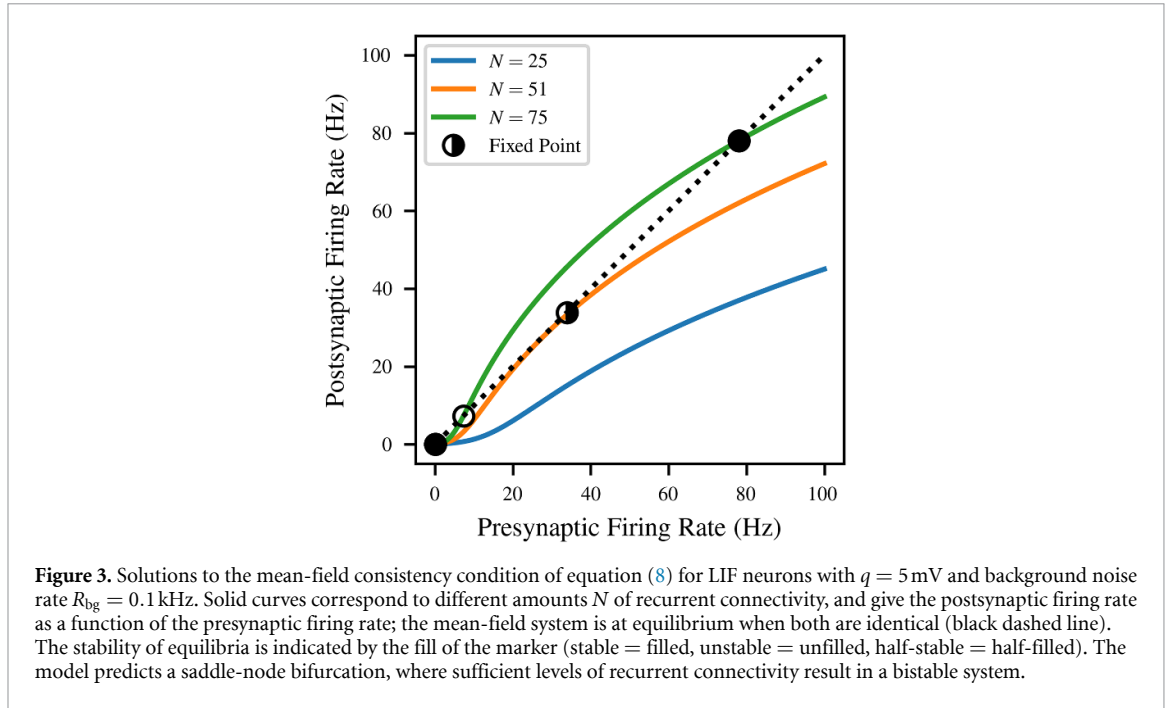
The curve for the diffusion limit (black dashed) represents an analytical solution in the limit of infinitesimal PSP size (with constant D) and simulation time step. However, in this simulation, where both quantities are finite, the analytical solution loses substantial accuracy, resulting in a noticeable bias in the reported 'error' in the lower panel. The fitted ReLU (green) performs better, but suffers from an overestimate in the extrapolation regime, as well as a small downward bump caused by its sharp elbow. On the other hand, Refractory SoftPlus (red) is capable of capturing the functional form of the simulated transfer function, with virtually no error in the interpolation regime, and lower error in extrapolation as well.

2.5. Mean-field modeling

Next, we apply the Refractory SoftPlus fit of a neuronal transfer function to derive a mean-field model for the equilibrium firing rate of a recurrently connected population of M neurons, each of which fires as a renewal process at the rate $F(R)$ when subject to a Poisson input with total rate R. We assume that each neuron is subject to background synaptic input at a total rate R_{bg}, in addition to receiving inputs from a random set of N other neurons within the population. We exclude multiple and self connections from the present analysis. Note that R_{bg} can be quite large compared to physiological firing rates (even in the tens of kHz) because it represents the combined activity of a large number of input neurons. The neurons within the population are grouped into an excitatory and an inhibitory subpopulation such that the excitatory fraction η is the same for the population as for the background input, and the PSP weights q_e and q_i are set as described in section 2.3 as well.

Under these conditions, the recurrent input and the background input will behave similarly, such that the neuron can be described by its transfer function using a total input rate $r_{\text{eff}} = R_{\text{bg}} + Nr$. This enables writing a consistency condition which can be solved to find the equilibrium firing rate of the population:

$$r = F(r_{\text{eff}}) = F\left(R_{\text{bg}} + Nr\right). \tag{8}$$

Figure 3. Solutions to the mean-field consistency condition of equation (8) for LIF neurons with $q = 5\,\text{mV}$ and background noise rate $R_{\text{bg}} = 0.1\,\text{kHz}$. Solid curves correspond to different amounts N of recurrent connectivity, and give the postsynaptic firing rate as a function of the presynaptic firing rate; the mean-field system is at equilibrium when both are identical (black dashed line). The stability of equilibria is indicated by the fill of the marker (stable = filled, unstable = unfilled, half-stable = half-filled). The model predicts a saddle-node bifurcation, where sufficient levels of recurrent connectivity result in a bistable system.

The solutions to this equation are visualized for several concrete values of N in figure 3. In this example, we numerically fitted the parameters of equation (7), as described in section 2.4, for LIF neurons with the large PSP size $q = 5\,\text{mV}$ and sparse Poisson background input $R_{\text{bg}} = 0.1\,\text{kHz}$. In each case, stability is defined by interpreting equation (8) as a map. If the absolute value of the derivative of the map is greater than 1 at the fixed point, iterated application will diverge and the fixed point is unstable [59]. Stable fixed points have derivatives with magnitude less than 1, and half-stable fixed points are those where the absolute value of the derivative is exactly 1.

As shown in the figure, the number and location of solutions to the consistency condition depend on the value of the connectivity parameter N. When there is very little recurrent connectivity, the only stable firing rate is the spontaneous rate caused by the background input, close to $r = 0$ (blue). As N increases, the consistency curve gradually slopes upwards to contact the $r = F(r_{\text{eff}})$ line. This causes a saddle-node bifurcation at $N = 51$, which creates a half-stable fixed point at $r = 33.9\,\text{Hz}$ (orange). As N continues to increase, the bifurcation continues, with the half-stable fixed point becoming one stable and one unstable fixed point (green).
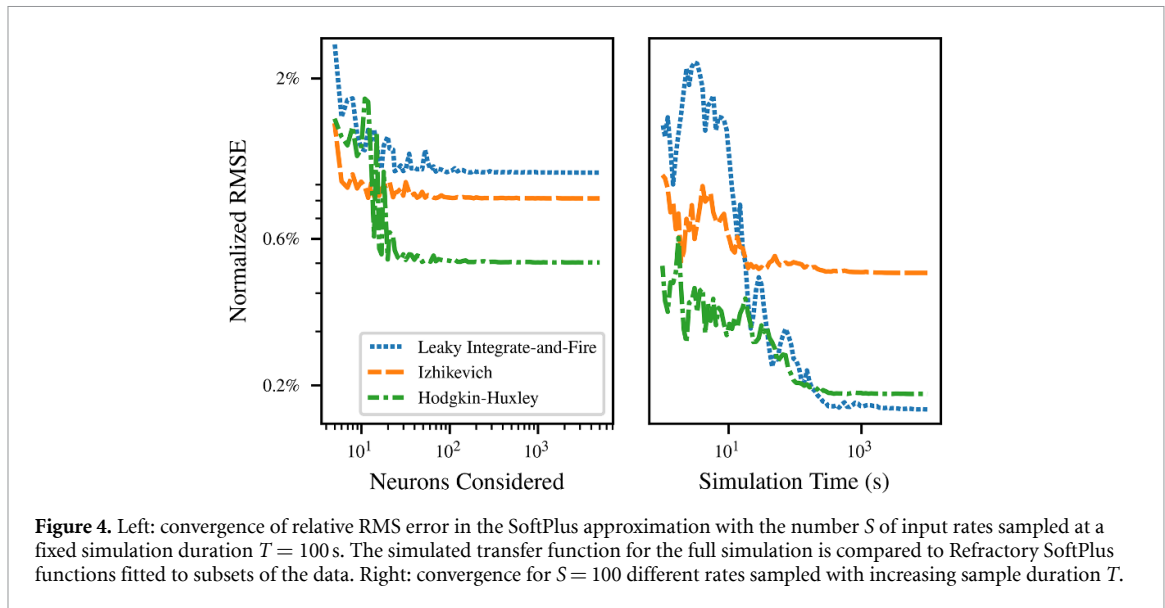
2.6. Dynamics model

So far we have described only the equilibrium firing rate of a population of neurons, but it is typically of greater interest to investigate the response of a population to a time-varying external input. We do this using a simplified version of the Master equation approach [41]. As in the prior work, we assume that the mean firing rate evolves according to linear dynamics with a fixed characteristic timescale τ [60]. This yields the following simple dynamical model:

$$\tau \frac{\mathrm{d}}{\mathrm{d}t} r = F\left(Nr + R_{\text{bg}}\right) - r. \tag{9}$$

The dynamics of equation (9) have stationary points at the zeros of its right hand side, which are exactly the fixed points of equation (8). In prior work on the master equation formalism, the timescale τ was considered somewhat arbitrary, but it was tightly coupled with the transfer function, and needed to be chosen carefully to yield accurate results [41, 60].

However, in our model, τ appears only as the timescale of convergence, so model performance is insensitive to its value. As a result, τ can be set approximately based on a simple experiment such as a step response (as we do in this paper), or, when studying slow phenomena, simply specified to be faster than the largest frequency of interest. The only limitation is that τ should be large enough to enable the quasistatic assumption. The fitted transfer function describes the stationary firing rate of the neuron, which means that higher-order dynamical features such as SFA can lead to errors in the transient response of the model if the dynamics of interest are not fast compared to τ.

Figure 4. Left: convergence of relative RMS error in the SoftPlus approximation with the number S of input rates sampled at a fixed simulation duration $T = 100$ s. The simulated transfer function for the full simulation is compared to Refractory SoftPlus functions fitted to subsets of the data. Right: convergence for $S = 100$ different rates sampled with increasing sample duration T.

3. Results and discussion

3.1. Transfer function fit performance

In order to assess the practicality of a data-driven mean-field model based on a Refractory SoftPlus fit, we first considered the amount of simulation data required to achieve an accurate fit. The computation time required to simulate a neuron's transfer function scales as $\mathcal{O}(ST)$, i.e. bilinearly with the number S of different input rates considered and the total simulated time T over which firing rates are calculated. Increasing either of these quantities increases the quality of the fit by increasing the precision or number of data points used. However, the discrepancy between Refractory SoftPlus and the true transfer function of the neuron also imposes a limit on the achievable accuracy. Therefore, we sought to minimize the computational costs of achieving the best possible quality of fit to the simulated transfer function.

Figure 4 demonstrates convergence of the fit for the LIF (blue), Izhikevich (orange), and the Hodgkin-Huxley (green) neurons. For each model, we simulated the transfer function once for large numbers S of input rates and once for large T, then considered subsets of the data. Fits were performed to subsets of the full simulation output, and the error of the fitted function was assessed on the final simulated transfer function.

The left side of figure 4 demonstrates convergence with the number S of input rates sampled. The PSP amplitude was fixed at $q = 1$ mV, and input rates ranged from 0 to $R = 100$ kHz. Fits were then calculated for evenly spaced subsets of this range. The fit error is already only a few percent at $S = 5$, and decreases to a final plateau value by $S = 100$. This can be attributed to the small number of parameters included in the fitted transfer function, which allows it to match observations without a large number of different data points.

The right side of figure 4 shows the convergence of the fit error as the duration of simulation considered for the calculation of firing rate for each neuron increases from $T = 1$ s to $T = 1 \times 10^4$ s. A single simulation was performed of $S = 100$ different rates for a duration $T = 1 \times 10^4$ s. Shorter simulation times were handled by considering only the beginning of the simulation data. This ensures that any startup transients and nonstationary effects, e.g. due to SFA, are included to the same extent that they would be in disjoint simulations. Again, fit error begins at only a few percent, and inclusion of more data in the fit soon reaches a point of diminishing returns.

Interestingly, for smaller T, the error of the fitted curve was substantially lower than the error of the short-time firing rate estimates to which the curve was fitted. The advantage continues for T up to 2.7×10^3 s for LIF neurons, 1.1×10^2 s for Izhikevich, and 7.4×10^2 s for Hodgkin–Huxley, with error below 0.5% in each case. (After this point, the error in the short-time firing rate estimates continue downwards as $T^{-1/2}$ as suggested by appendix E.) This demonstrates that the fitted Refractory SoftPlus transfer function is able to reject a significant amount of noise in the empirical transfer function by providing averaging between adjacent input rates.

3.2. Generalization to other neuron models

In the majority of this paper, we use the simulator's default parameters, but varying these does not affect our conclusions. Indeed, the results can be generalized not only to other parameter values, but also to other

Table 1. Distributions of random parameter values used to generate figure 5. The notation $\text{Exp}(\mu)$ represents an exponential distribution by its mean, $\mathcal{N}(\mu, \sigma)$ represents a normal distribution by its mean and standard deviation, and $\mathcal{U}_{[a,b]}$ represents a uniform distribution over an interval.

Neuron	Parameter	Default	Distribution
Leaky Integrate-and-Fire	T_{ref}	2 ms	$\text{Exp}(2)$
	C	250 pF	$\mathcal{N}(250, 50)$
Izhikevich	a	0.02	$\mathcal{U}_{[0.02, 0.1]}$
	b	0.2	$\mathcal{U}_{[0.2, 0.25]}$
	c	-65	$\mathcal{U}_{[-65, -50]}$
	d	8	$\mathcal{U}_{[2,8]}$
Hodgkin–Huxley	T_{ref}	2 ms	$\text{Exp}(2)$
	C	100 pF	$\mathcal{N}(100, 10)$
	$\tau_{\text{syn,exc}}$	0.2 ms	$\text{Exp}(1)$
	$\tau_{\text{syn,inh}}$	2 ms	$\text{Exp}(1)$
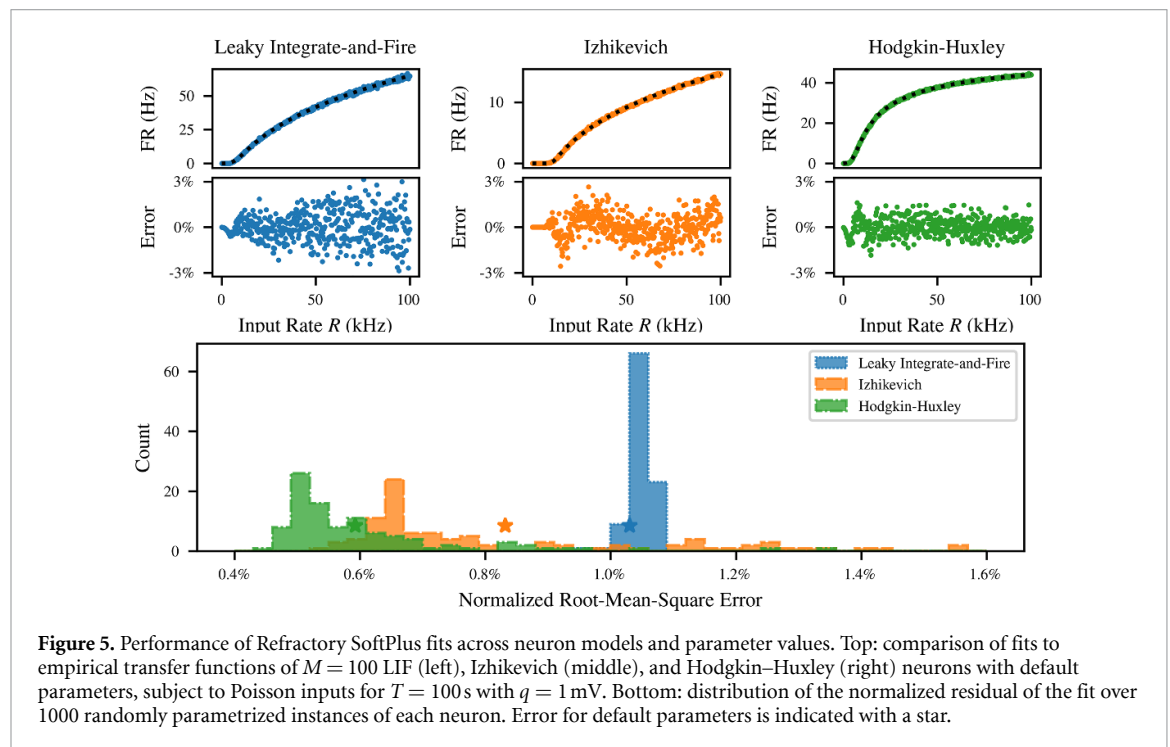
Figure 5. Performance of Refractory SoftPlus fits across neuron models and parameter values. Top: comparison of fits to empirical transfer functions of $M = 100$ LIF (left), Izhikevich (middle), and Hodgkin–Huxley (right) neurons with default parameters, subject to Poisson inputs for $T = 100$ s with $q = 1$ mV. Bottom: distribution of the normalized residual of the fit over 1000 randomly parametrized instances of each neuron. Error for default parameters is indicated with a star.

neuron models. To demonstrate this, we assessed the quality of the fit as the parameters of each neuron model were varied randomly according to the distributions given in table 1. In each case, dynamics of a single postsynaptic neuron were numerically simulated for 100 seconds subject to balanced Poisson input with $\eta = 0.8$ and $q = 1.0$ mV at 500 different total input rates R from 0 to 100 kHz. Performance is evaluated on the LIF neuron discussed above, together with the efficient quadratic integrate-and-fire neuron of Izhikevich [53] and the biophysical Hodgkin–Huxley neuron [2]. The results of this comparison are shown in figure 5. Crucially, Refractory SoftPlus is a good fit for the transfer function of all these models, despite their different dynamics and qualitative behavior.

For the LIF and Hodgkin–Huxley neurons, the distributions were derived from the simulator defaults by giving each time constant an exponential distribution with its default value as the mean, while the membrane capacitance was normally distributed about its default. For the Izhikevich neuron, the distribution of each parameter was made uniform over the range of values considered in the original description of the model [53]. These choices of distribution are arbitrary, but they cover a broad parameter range for each model. For each of the models shown in figure 5, the fitting error under default parameters is indicated with a colored star in the histogram. This point is not an outlier, indicating that our results generalize well across neuron models for a broad range of parameters.

Note, however, that these results should be expected to depend on the dynamical regimes considered. In particular, bursting and strong adaptation are incompatible with any transfer function approach. The
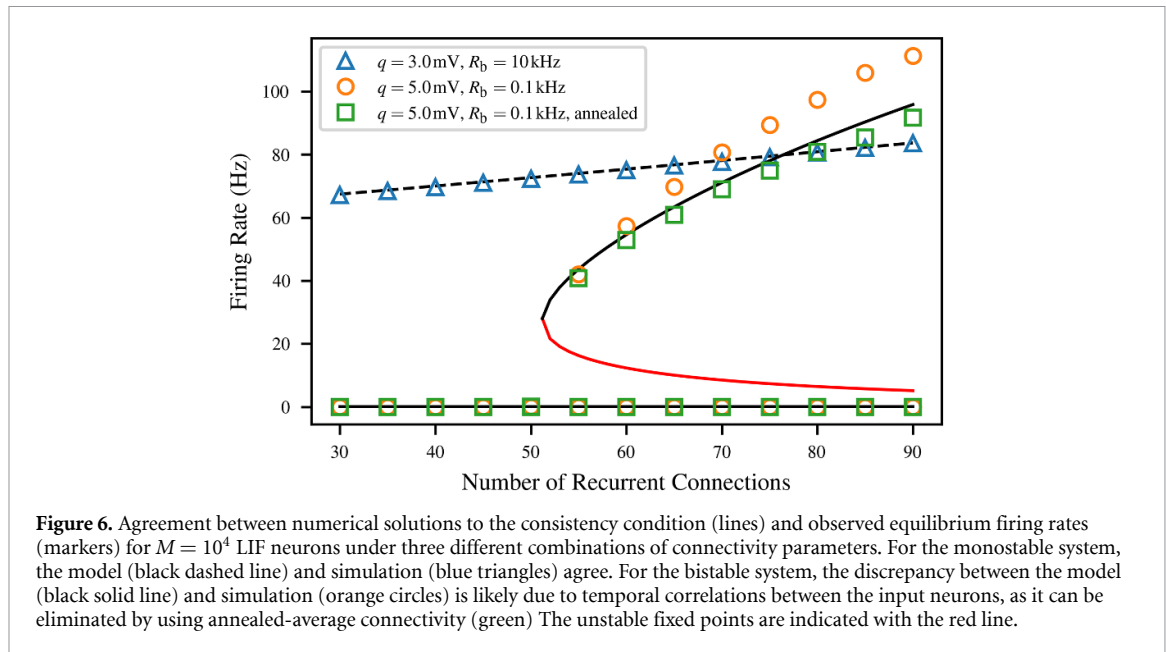
Figure 6. Agreement between numerical solutions to the consistency condition (lines) and observed equilibrium firing rates (markers) for $M = 10^4$ LIF neurons under three different combinations of connectivity parameters. For the monostable system, the model (black dashed line) and simulation (blue triangles) agree. For the bistable system, the discrepancy between the model (black solid line) and simulation (orange circles) is likely due to temporal correlations between the input neurons, as it can be eliminated by using annealed-average connectivity (green) The unstable fixed points are indicated with the red line.

qualitative dynamics of the LIF neuron are unaffected by its parameters, so the performance of the model is quite consistent in this case. On the other hand, the Izhikevich and HH neurons experience multiple dynamical regimes throughout their parameter space. In their default parametrization, all neurons considered are in a dynamical regime with weak or no adaptation, and do not experience bursts. This is addressed in more detail in appendix D.

Also note that for simplicity, the notation with which equation (2) describes the parametrization of synaptic weights to match results between PSP-based and PSC-based neurons ignores the fact that excitatory and inhibitory PSCs may have different parameters. In our simulations, not only were the PSP amplitudes q_e and q_i different because $\eta \neq \frac{1}{2}$, but also the synaptic time constants in the Hodgkin-Huxley neuron differ for excitatory and inhibitory synapses. However, despite these asymmetries and significant deviations from the input regime in which the approximations were derived, the flexibility of our assumptions has allowed fits to the simulated transfer function of the Hodgkin-Huxley neuron to be quite good as well.

3.3. Bifurcation analysis

Next we compared the bifurcation and equilibrium locations between the mean-field model and simulation. Activity fixed points were calculated for three different networks of $M = 10^5$ LIF neurons, as shown in figure 6. Simulations were performed at a range of values of N, each corresponding to a different network where each neuron receives connections from N presynaptic neurons sampled without replacement from the remainder of the population. In this and all other simulations where recurrent connectivity was included, recurrent synapses included an axonal delay chosen uniformly at random between 1 and 10 ms to represent the possibility of neurons spread out in space [1, 61], but results were similar for small fixed delays as well. The effect of delays on population dynamics is discussed in appendix F.

Each displayed fixed point is the average firing rate across 10 simulations, each 2.5 s in duration. The lower stable population rate is computed by simply simulating the network from a zero initial condition subject to its background Poisson input. To find the higher stable population rate in the bistable network, we transiently raised the firing rate of the population using a 'warm-up' input intended to shift it into the basin of attraction of the upper fixed point. We did this by simulating the population without counting spikes for a 'warm-up time' of 1 s while the population was subject to additional excitatory Poisson input. To minimize the effects of this drive on the firing rate, we linearly ramped it down from R_{bg} to 0 over that time. This was done for all values of N, not only those expected to be bistable.

The first condition pictured is a high-noise case which is monostable for the full range of N. Here, $q = 3$ mV, and $R_{bg} = 10$ kHz. In this situation, the single firing rate fixed point increases slightly from the baseline spontaneous activity level at $N = 0$, but no interesting dynamics are observed. The model agrees well with simulation for the full simulated range, which is expected because recurrent input contributes under half of the total (ranging from 17% at $N = 30$ to 42% at $N = 90$). This means that there is enough background input to ensure that the assumptions of the model are satisfied even if the recurrent input deviates from Poisson statistics.

The second condition (orange) uses $q = 5\,\text{mV}$ and $R_{\text{bg}} = 0.1\,\text{kHz}$, the combination of parameters which produced a saddle-node bifurcation in figure 3. In this case, the accuracy of the mean-field model suffers somewhat. The network transitions from monostable to bistable between $N = 50$ and $N = 55$ as predicted, but the equilibrium firing rate in the active state is typically noticeably underestimated. This is likely caused by deviations in the input statistics from those used in calibrating the mean-field model. Across the upper branch of this condition, the fraction of input to the neuron which comes from the rest of the network rather than the Poisson background ranges from 97% to 99%. At the same time, firing rates increase, meaning that the absolute refractory period will make up a larger part of the interspike interval, and deviations from Poisson behavior will be more significant.

Indeed, the overestimation disappears entirely for annealed average connectivity. In this condition, rather than generating a fixed network where each neuron has N fixed presynaptic neurons, all-to-all connectivity is combined with probabilistic synapses, whose transmission probability N/M is chosen to match the input rate due to N presynaptic neurons on average [62]. This breaks up temporal correlations and noticeably improves model predictions of equilibrium firing rates.

An additional caveat must be made specifically for the bistable networks with N just above the bifurcation. At $N = 55$, the unstable fixed point, which marks the boundary between the basins of attraction of the two equilibria, is at a relatively high firing rate. As a result, population activity fluctuations are often sufficient to move the network into the basin of attraction of the low-frequency equilibrium. Therefore, for $N = 55$ in particular, we rejected the simulations where this occurred, so the upper fixed point is based on only 1 data point for the second condition and 5 for the third. We did not observe the reverse basin hopping effect, where the network spontaneously jumps to the higher fixed point, because fluctuations near zero firing rate are much smaller.

3.4. Finite size effects

Although our mean-field model is based on finite connectivity, it does not explicitly depend on the total number of neurons M in the population. However, a sufficiently large M is implicitly assumed by requiring the neurons to be independent of each other. Two deterministic neurons will only fire independently if their inputs are independent. This can happen in two ways: (1) when M is very large, the inputs to each neuron will be disjoint, and therefore conditionally independent given the population firing rate; and (2) when $R_{\text{bg}} \gg Nr$, the input to each neuron is predominantly from the background, which is independent by construction. In this section, we investigate the effect of varying M on the dynamical and steady-state behavior of the model.

Figure 7 explores a moderate case with R_{bg} and Nr on the same order of magnitude, by setting $q = 3\,\text{mV}$, $R_{\text{bg}} = 10\,\text{kHz}$, and $N = 75$. For each of a range of values of M, ten different networks were simulated, each with the same overall parameters but different randomly generated connectivity. The binned firing rate was calculated as the number of spikes in each 1 ms bin of the 2 s recording, divided by M as well as the bin size. The result is a time-varying approximation of the firing rate in hertz, which follows a random walk with regression to the mean predicted by the model, as pictured. The standard deviation of the binned firing rate decreases proportionally to $M^{-1/2}$, as might be expected from the central limit theorem. This is illustrated using the function $M \to kM^{-1/2}$ with k fitted to the data by least squares (black dashed). The variation in the mean between simulations also decreases, but does not follow such a clear functional form.

Particularly surprising in light of past results showing substantial finite-size effects in population bursts [63] is the accuracy of the fixed point even when $M = 100$. For small M, input is highly correlated between neurons because the 75 synapses are chosen from only 100 neurons. This highly correlated input makes up about 37% of the total input to the neuron, but the location of the firing rate fixed point remains quite close to the model predictions. Evidently, the simple presence of background input on the same order of magnitude as the recurrent input adequately rescues the independence assumption even when the recurrent input itself is not close to independent.

Past mean-field models have assumed a specific form for the variance in the population rate, based on the binomial distribution [41, 64]. However, this assumption only applies to the fluctuating regime, where population activity is primarily input-driven; the results of these works consistently used less than 10% recurrent input. On the other hand, in the biologically common 'reverberating' regime, where population activity is driven substantially by recurrent inputs, correlated fluctuations tend to appear, which cannot be modeled in this way [62, 65]. Indeed, although at $N = 0$ the binomial prediction is recovered (not pictured), in this parameter regime, we observe an order of magnitude greater variance than would be predicted by such models.

Despite the insensitivity of the equilibrium location to M, an important finite size effect appears in the bifurcation analysis of figure 6. The simulations with $N = 55$ are close to a bifurcation, which means that the unstable fixed point marking the watershed between the lower and upper fixed points is 16 Hz. As a result, at
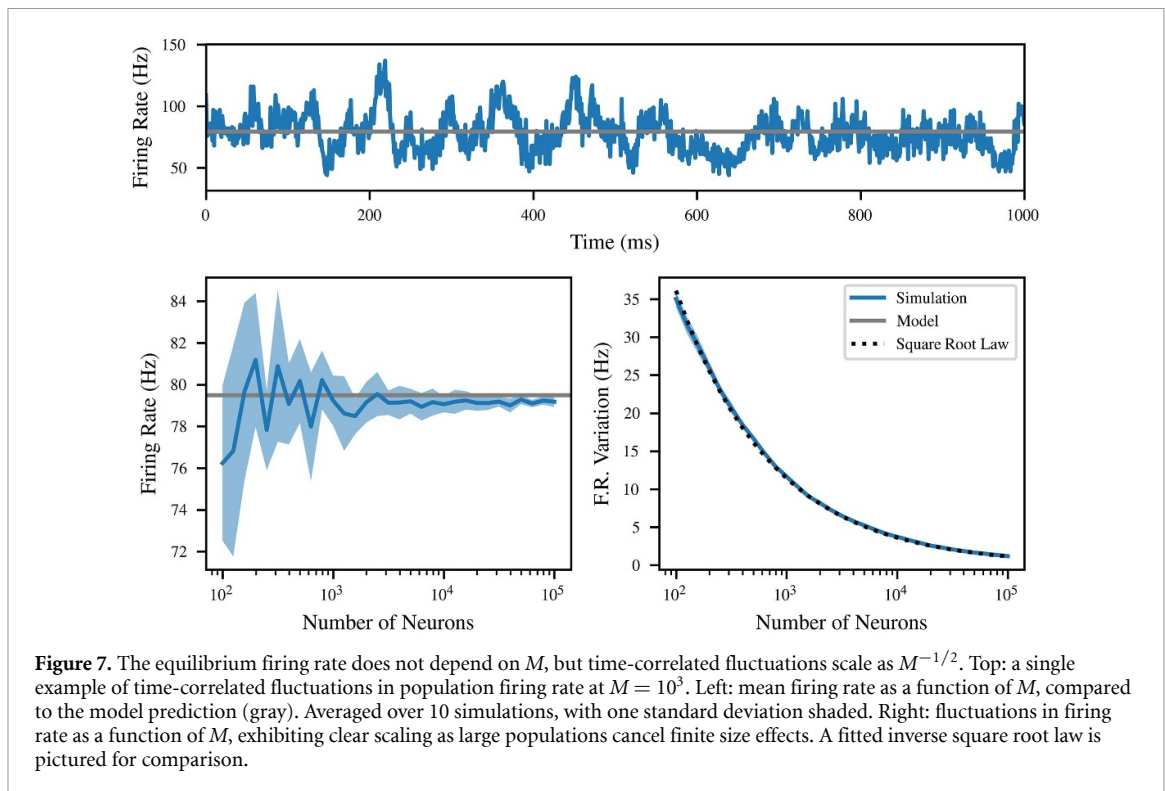
Figure 7. The equilibrium firing rate does not depend on M, but time-correlated fluctuations scale as $M^{-1/2}$. Top: a single example of time-correlated fluctuations in population firing rate at $M = 10^3$. Left: mean firing rate as a function of M, compared to the model prediction (gray). Averaged over 10 simulations, with one standard deviation shaded. Right: fluctuations in firing rate as a function of M, exhibiting clear scaling as large populations cancel finite size effects. A fitted inverse square root law is pictured for comparison.

Table 2. Practical stability of the upper equilibrium at $q = 5\,\mathrm{mV}$, $R_{\mathrm{bg}} = 0.1\,\mathrm{kHz}$, and $N = 55$ in figure 6, expressed as the percentage of 2.5 s simulations which avoid dropping to the other equilibrium as a function of M. Results are given for static random networks as well as annealed average connectivity.

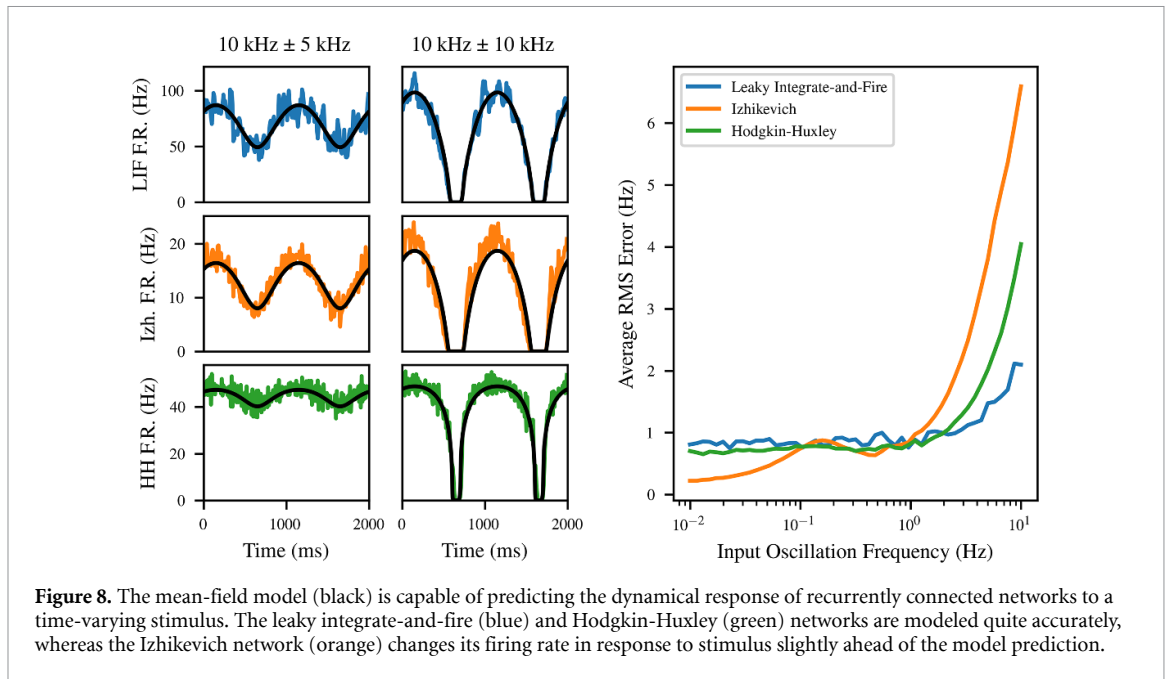| Network size	Stability	
	Static	Annealed
$M = 5000$	0%	0%
$M = 10000$	2%	10%
$M = 20000$	10%	72%
$M = 50000$	80%	100%
$M = 100000$	100%	100%

$M = 10^5$, the fluctuations in population rate are large enough that most simulations 'fall off' the fixed point within a few seconds of simulation. In other words, stable equilibria predicted by equation (9) may be only metastable for finite M. Interestingly, it was recently proved that all-to-all coupled probabilistic integrate-and-fire neurons driven by Poisson noise have a unique stationary firing rate distribution [66]. Thus, the appearance of metastability in our results suggests that finite-size effects are exacerbated by either hard spiking thresholds or sparse connectivity.

To investigate further, we define practical stability as the probability that a network will stay near the upper fixed point for the entire duration of the 2.5 s simulation used to calculate the firing rate in the previous section. This finite-size effect was recently investigated in detail in networks of quadratic integrate-and-fire neurons [67]. Table 2 presents the practical stability of this particular case across a range of values of M. There is a clear trend towards higher practical stability as population rate fluctuations are reduced, with better practical stability for annealed-average than for static random connectivity. 50 simulations were performed for each case in the table.

3.5. Dynamical response
Finally, we investigated the ability of the simple dynamics described in section 2.6 to model the sinusoid-following behavior of a recurrently connected network. In each case, a network of $M = 10^3$ neurons with $q = 3\,\mathrm{mV}$, $N = 50$, and $R_{\mathrm{bg}} = 10\,\mathrm{kHz}$ was simulated subject to independent per-neuron Poisson background input as above, but in this case it is not stationary. Instead, the rate of the background input varies sinusoidally with time:

$$R_{\mathrm{in}}(t) = R_{\mathrm{bg}} + A_{\mathrm{bg}} \sin\left(2\pi f_{\mathrm{bg}} t\right) \qquad (10)$$

Figure 8. The mean-field model (black) is capable of predicting the dynamical response of recurrently connected networks to a time-varying stimulus. The leaky integrate-and-fire (blue) and Hodgkin-Huxley (green) networks are modeled quite accurately, whereas the Izhikevich network (orange) changes its firing rate in response to stimulus slightly ahead of the model prediction.

We keep the frequency $f_{bg} = 1$ Hz fixed and consider two different oscillation amplitudes $A_{bg} = 5$ kHz and 10 kHz. The left part of figure 8 considers these two amplitudes (columns) across three different neuron models (rows). In the first case, the population firing rate also oscillates roughly sinusoidally, whereas when the background rate drops to zero, the effects appear much less linear.

The right side of figure 8 shows the RMS error of the left experimental condition as a function of the frequency with which the firing rate of the background input varies over time. The mean-field prediction is compared to the average of 50 independent simulations of the same experimental condition in order to avoid confounding by the substantial variation between independent simulations. The result is an error which remains relatively flat with frequency below 1 Hz, but which becomes substantially larger as frequency increases. This is to be expected, as the mean-field model is based on stationarity assumptions which cannot be expected to hold as the rate of population fluctuations approaches the slowest time constants of the individual neuronal dynamics.

It has been observed that the response of the population firing rate to fluctuations in the input is much faster than the response of any individual neuron to those changes [68]. Indeed, we find that although the autonomous population rate fluctuations in figure 7 have a timescale in the tens of milliseconds, the population rate responds to a step input virtually instantaneously (not pictured). Indeed, past work suggests that the perturbation response of the population rate is limited by synaptic, not membrane, time constants [68]. Therefore, we set $\tau = 1$ ms in these simulations. Appendix F shows that changing τ has minimal effect for slowly fluctuating inputs.

Interestingly, the firing rate of the Izhikevich neurons appears to lead the prediction of the mean-field model, as if a derivative term is missing from the true dynamics. The derivative effect disappears under a set of parameters where SFA is absent (not pictured). This suggests that the discrepancy is caused by a failure of the model's stationarity assumption on short timescales. Since the recovery variable u is directly driven by firing, there can be a significant difference between the steady-state and transient response of the neuron. Indeed, a recent detailed mean-field model of all-to-all coupled Izhikevich neurons observed a similar effect which disappeared when SFA terms were replaced with their population mean [69]. This limitation is not absolute, however. First, the small-deviation example is less problematic, because the population continues firing throughout the simulation in this case, so the adaptation level varies less. Also, the Hodgkin-Huxley neuron has larger-timescale effects due to dynamical variables which persist across multiple firings, and yet is modeled quite accurately in both examples.

4. Conclusion

We have developed an approximate mean-field model capable of finding the activity fixed points of a randomly connected network of neurons under the influence of synaptic inputs from other neurons in the network together with background Poisson input. This method is applicable to a wide variety of different neuronal models, and parameter values for the transfer function used in the mean-field approximation can

be derived from relatively small-scale and short-duration simulations. The resulting model predicts equilibrium and dynamical population firing rates in a randomly connected network as well as the location of a bifurcation from monostability to bistability.

The method of fitting a transfer function to simulations of the neuron under Poissonian input allows relaxing two key assumptions of past mean-field models. The diffusion limit is avoided by explicitly using the rate of finite-magnitude presynaptic events as input to the transfer function. Furthermore, since the transfer function is fitted numerically using only the neuron's spike times, our model does not utilize any information about the neuron's subthreshold dynamics. This is not entirely hypothesis-free, as some subthreshold dynamics may result in transfer functions that cannot be approximated well by Refractory SoftPlus, but our approach can just as easily incorporate a different transfer function in such cases.

Mean-field models often assume a particular distribution of the postsynaptic membrane potential, generally that it is normally distributed [20, 41, 70]. However, realistic membrane potential distributions can be far from normal [43], making this so-called 'moment closure' assumption hard to justify. Another approach instead assumes a Lorentzian distribution, but this is incompatible with master-equation-type models due to the pathological behavior of the moments of the Cauchy distribution [26, 69]. Our approach avoids moment closure by instead assuming that the *input* to the neuron is sufficiently similar to the input used in calculating the original transfer function.

We observed an improvement in model accuracy in the most difficult parameter regime when the recurrent input within the population was made more Poisson-like by using annealed-average connectivity instead of a static random network. However, annealed-average connectivity is a strong assumption not typical in modeling; if accuracy in this regime is essential, it would be more useful to fit the transfer function under inputs more similar to those which will later be observed. Although the present work assumes Poisson inputs with loose excitatory-inhibitory balance, this assumption is not fundamental, and simulated transfer functions can in principle be calculated for other input models. This possibility is a key advantage of our approach.

The main limitations of our model are in the simplicity of the dynamics. Past mean-field models of firing rate variance have approached only the fluctuating regime [41, 64], and generalization to the reverberating dynamics considered here is challenging. For this reason, we used only a first-order master equation, which does not quantify fluctuations from the mean, and so requires somewhat larger populations for full accuracy in some parameter regimes [60]. Furthermore, the use of a transfer function requires assuming a stationary distribution of neuronal firing rate [60], but by making a quasistatic assumption, such models can also describe the dynamics of the firing rate [41]. Furthermore, mean-field models can be augmented to represent other population dynamics such as adaptation variables or population heterogeneity, either by replacing dynamical terms with explicit population averages [71, 72], or by expanding a model's existing usage of distributional models of single-neuron quantities [73]. Incorporating these effects into the dynamics is a key direction of future work.

Like the results of previous authors, this approximation can be viewed as a bridge between computational neuroscience and machine learning perspectives on neural networks. Given that a similar functional form appears to be conserved across multiple spiking neural models, it is conceivable that it serves a function in biological neural networks as well as in artificial neural networks. In any case, theoretical developments may benefit from the existence of such a simple functional form which is robust across such a wide variety of neuronal dynamics.

Data availability statement

Acknowledgment

Conflict of interest

had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

ORCID iDs

Alex Spaeth ● https://orcid.org/0000-0003-0702-3945
Mircea Teodorescu ● https://orcid.org/0000-0001-7085-5248

References

[1] Braitenberg V and Schüz A 1991 *Anatomy of the Cortex: Statistics and Geometry* (*Studies of Brain Function*) (Springer)
[2] Hodgkin A L and Huxley A F 1952 A quantitative description of membrane current and its application to conduction and excitation in nerve *J. Physiol.* **117** 500–44
[3] Izhikevich E M 2006 *Dynamical Systems in Neuroscience* (MIT Press)
[4] Wilson H R and Cowan J D 1973 A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue *Kybernetik* **13** 55–80
[5] Chaudhuri R, Knoblauch K, Gariel M-A, Kennedy H and Wang X-J 2015 A large-scale circuit mechanism for hierarchical dynamical processing in the primate cortex *Neuron* **88** 419–31
[6] Vladimirski B B, Tabak J, O'Donovan M J and Rinzel J 2008 Episodic activity in a heterogeneous excitatory network, from spiking neurons to mean field *J. Comput. Neurosci.* **25** 39–63
[7] Shine J M, Müller E J, Munn B, Cabral J, Moran R J and Breakspear M 2021 Computational models link cellular mechanisms of neuromodulation to large-scale neural dynamics *Nat. Neurosci.* **24** 765–76
[8] Vandesompele A, Urbain G, wyffels F and Dambre J 2020 Populations of spiking neurons for reservoir computing: closed loop control of a compliant quadruped *Cogn. Syst. Res.* **58** 317–23
[9] Spaeth A, Tebyani M, Haussler D and Teodorescu M 2020 Spiking neural state machine for gait frequency entrainment in a flexible modular robot *PLoS One* **15** e0240267
[10] Jaeger H 2021 Towards a generalized theory comprising digital, neuromorphic and unconventional computing *Neuromorph. Comput. Eng.* **1** 012002
[11] Tayarani-Najaran M H and Schmuker M 2021 Event-based sensing and signal processing in the visual, auditory and olfactory domain: a review *Front. Neural Circuits* **15** 610446
[12] Davies M *et al* 2021 Advancing neuromorphic computing with Loihi: a survey of results and outlook *Proc. IEEE* **109** 911–34
[13] Polykretis I, Supic L and Danielescu A 2023 Bioinspired smooth neuromorphic control for robotic arms *Neuromorph. Comput. Eng.* **3** 014013
[14] Trujillo C A *et al* 2019 Complex oscillatory waves emerging from cortical organoids model early human brain network development *Cell Stem Cell* **25** 558–69
[15] Samarasinghe R A *et al* 2021 Identification of neural oscillations and epileptiform changes in human brain organoids *Nat. Neurosci.* **24** 1–13
[16] Sharf T, van der Molen T, Glasauer S M K, Guzman E, Buccino A P and Luna G *et al* 2022 Functional neuronal circuitry and oscillatory dynamics in human brain organoids *Nat. Commun.* **13** 4403
[17] Cai H *et al* 2023 Brain organoid reservoir computing for artificial intelligence *Nat. Electron.* **6** 1032–9
[18] van der Molen T *et al* 2023 Protosequences in human cortical organoids model intrinsic states in the developing cortex *bioRxiv* (https://doi.org/10.1101/2023.12.29.573646) (posted online 30 December 2023, accessed 1 July 2024)
[19] Andrews J P *et al* 2024 Multimodal evaluation of network activity and optogenetic interventions in human hippocampus *Nat. Neurosci.* accepted
[20] Amit D J and Brunel N 1997 Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex *Cereb. Cortex* **7** 237–52
[21] Gerstner W, Kistler W M, Naud R and Paninski L 2014 *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition* (Cambridge University Press)
[22] Chemla S *et al* 2019 Suppressive traveling waves shape representations of illusory motion in primary visual cortex of awake primate *J. Neurosci.* **39** 4282–98
[23] Chen L and Campbell S A 2022 Exact mean-field models for spiking neural networks with adaptation *J. Comput. Neurosci.* **50** 445–69
[24] Pietras B and Daffertshofer A 2019 Network dynamics of coupled oscillators and phase reduction techniques *Phys. Rep.* **819** 1–105
[25] Bick C, Goodfellow M, Laing C R and Martens E A 2020 Understanding the dynamics of biological and neural oscillator networks through exact mean-field reductions: a review *J. Math. Neurosci.* **10** 9
[26] Montbrió E, Pazó D and Roxin A 2015 Macroscopic description for networks of spiking neurons *Phys. Rev.* X **5** 021028
[27] Pietras B, Gallice N and Schwalger T 2020 Low-dimensional firing-rate dynamics for populations of renewal-type spiking neurons *Phys. Rev.* E **102** 022407
[28] Renart A, Brunel N and Wang X J 2004 Mean-field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks *Computational Neuroscience: A Comprehensive Approach* (*Mathematical Biology and Medicine Series*) (Chapman & Hall/CRC) pp 431–90
[29] Augustin M, Ladenbauer J, Baumann F and Obermayer K 2017 Low-dimensional spike rate models derived from networks of adaptive integrate-and-fire neurons: comparison and implementation *PLOS Comput. Biol.* **13** e1005545
[30] Bos H, Diesmann M and Helias M 2016 Identifying anatomical origins of coexisting oscillations in the cortical microcircuit *PLOS Comput. Biol.* **12** e1005132
[31] Xiao Z-C, Lin K K, Young L-S and Rubin J 2021 A data-informed mean-field approach to mapping of cortical parameter landscapes *PLOS Comput. Biol.* **17** e1009718
[32] Hertäg L, Durstewitz D and Brunel N 2014 Analytical approximations of the firing rate of an adaptive exponential integrate-and-fire neuron in the presence of synaptic noise *Front. Comput. Neurosci.* **8** e116
[33] Rosenbaum R 2016 A diffusion approximation and numerical methods for adaptive neuron models with stochastic inputs *Front. Comput. Neurosci.* **10** e39

[34] Ricciardi L M and Sacerdote L 1979 The Ornstein-Uhlenbeck process as a model for neuronal activity *Biol. Cybern.* **35** 1–9
[35] Brunel N and Latham P E 2003 Firing rate of the noisy quadratic integrate-and-fire neuron *Neural Comput.* **15** 2281–306
[36] Droste F and Lindner B 2014 Integrate-and-fire neurons driven by asymmetric dichotomous noise *Biol. Cybern.* **108** 825–43
[37] Carlu M *et al* 2020 A mean-field approach to the dynamics of networks of complex neurons, from nonlinear integrate-and-fire to Hodgkin-Huxley models *J. Neurophysiol.* **123** 1042–51
[38] Liang J, Zhou T and Zhou C 2020 Hopf bifurcation in mean field explains critical avalanches in excitation-inhibition balanced neuronal networks: a mechanism for multiscale variability *Front. Syst. Neurosci.* **14** 580011
[39] Ehsani M and Jost J 2023 Scale free avalanches in excitatory-inhibitory populations of spiking neurons with conductance based synaptic currents *J. Comput. Neurosci.* **51** 149–72
[40] Huang C-H and Lin C C K 2021 A novel density-based neural mass model for simulating neuronal network dynamics with conductance-based synapses and membrane current adaptation *Neural Netw.* **143** 183–97
[41] El Boustani S and Destexhe A 2009 A master equation formalism for macroscopic modeling of asynchronous irregular activity states *Neural Comput.* **21** 46–100
[42] Touboul J, Hermann G and Faugeras O 2012 Noise-induced behaviors in neural mean field dynamics *SIAM J. Appl. Dyn. Syst.* **11** 49–81
[43] Richardson M J E and Swarbrick R 2010 Firing-rate response of a neuron receiving excitatory and inhibitory synaptic shot noise *Phys. Rev. Lett.* **105** 178102
[44] Izhikevich E M 2000 Neural excitability, spiking and bursting *Int. J. Bifurcation Chaos* **10** 1171–266
[45] Manna D L, Vicente-Sola A, Kirkland P, Bihl T J and Caterina G D 2022 Simple and complex spiking neurons: perspectives and analysis in a simple STDP scenario *Neuromorph. Comput. Eng.* **2** 044009
[46] Dayan P and Abbott L F 2001 Theoretical neuroscience: computational and mathematical modeling of neural systems *Computational Neuroscience* (Massachusetts Institute of Technology Press)
[47] Softky W R and Koch C 1993 The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs *J. Neurosci.* **13** 334–50
[48] Stevens C and Zador A 1995 When is an integrate-and-fire neuron like a Poisson neuron? *Advances in Neural Information Processing Systems* vol 8 (MIT Press) pp 103–9
[49] Ramlow L and Lindner B 2021 Interspike interval correlations in neuron models with adaptation and correlated noise *PLOS Comput. Biol.* **17** e1009261
[50] Gallinaro J V, Gašparović N and Rotter S 2022 Homeostatic control of synaptic rewiring in recurrent networks induces the formation of stable memory engrams *PLOS Comput. Biol.* **18** e1009836
[51] Khinchin A Y 1955 *Mathematical Methods in the Theory of Queuing* (Akademia Nauka) (English translation published by Dover, Mineola, New York, 2013)
[52] Song S, Sjöström P J, Reigl M, Nelson S and Chklovskii D B 2005 Highly nonrandom features of synaptic connectivity in local cortical circuits *PLOS Biol.* **3** e68
[53] Izhikevich E M 2003 Simple model of spiking neurons *IEEE Trans. Neural Netw.* **14** 1569–72
[54] Droste F and Lindner B 2017 Exact analytical results for integrate-and-fire neurons driven by excitatory shot noise *J. Comput. Neurosci.* **43** 81–91
[55] Brunel N 2000 Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons *J. Comput. Neurosci.* **8** 183–208
[56] Destexhe A, Rudolph M, Fellous J-M and Sejnowski T J 2001 Fluctuating synaptic conductances recreate *in vivo*-like activity in neocortical neurons *Neuroscience* **107** 13–24
[57] Denéve S and Machens C K 2016 Efficient codes and balanced networks *Nat. Neurosci.* **19** 375–82
[58] Glorot X, Bordes A and Bengio Y 2011 Deep sparse rectifier neural networks *Proc. 14th Int. Conf. on Artificial Intelligence and Statistics* (*JMLR Workshop and Conf. Proc.*) pp 315–23
[59] Wiggins S 2003 *Introduction to Applied Nonlinear Dynamical Systems and Chaos* (*Texts in Applied Mathematics*) 2nd edn (Springer)
[60] Zerlaut Y, Chemla S, Chavane F and Destexhe A 2018 Modeling mesoscopic cortical dynamics using a mean-field model of conductance-based networks of adaptive exponential integrate-and-fire neurons *J. Comput. Neurosci.* **44** 45–61
[61] Lemaréchal J-D *et al* 2022 A brain atlas of axonal and synaptic delays based on modelling of cortico-cortical evoked potentials *Brain* **145** 1653–67
[62] Zierenberg J, Wilting J and Priesemann V 2018 Homeostatic plasticity and external input shape neural network dynamics *Phys. Rev. X* **8** 031018
[63] Gast R, Schmidt H and Knösche T R 2020 A mean-field description of bursting dynamics in spiking neural networks with short-term adaptation *Neural Comput.* **32** 1615–34
[64] Soula H and Chow C C 2007 Stochastic dynamics of a finite-size spiking neural network *Neural Comput.* **19** 3262–92
[65] Nolte M, Reimann M W, King J G, Markram H and Muller E B 2019 Cortical reliability amid noise and chaos *Nat. Commun.* **10** 1–15
[66] Schmutz V, Löcherbach E and Schwalger T 2023 On a finite-size neuronal population equation *SIAM J. Appl. Dyn. Syst.* **22** 996–1029
[67] Klinshov V V, Smelov P S and Kirillov S 2023 Constructive role of shot noise in the collective dynamics of neural networks *Chaos* **33** 061101
[68] Gerstner W 2000 Population dynamics of spiking neurons: fast transients, asynchronous states and locking *Neural Comput.* **12** 43–89
[69] Gast R, Solla S A and Kennedy A 2023 Macroscopic dynamics of neural networks with heterogeneous spiking thresholds *Phys. Rev. E* **107** 024306
[70] Vellmer S and Lindner B 2021 Fokker-Planck approach to neural networks and to decision problems *Eur. Phys. J. Spec. Top.* **230** 2929–49
[71] Di Volo M and Destexhe A 2021 Optimal responsiveness and information flow in networks of heterogeneous neurons *Sci. Rep.* **11** 17611
[72] Lorenzi R M *et al* 2023 A multi-layer mean-field model of the cerebellum embedding microstructure and population-specific dynamics *PLOS Comput. Biol.* **19** e1011434
[73] Gast R, Solla S A and Kennedy A 2024 Neural heterogeneity controls computations in spiking neural networks *Proc. Natl Acad. Sci.* **121** e2311885121