Neural Field Models for Latent State Inference: Application to Large-Scale Neuronal Recordings

M. E. Rule, D. Schnoerr, M. H. Hennig, G. Sanguinetti

February 7, 2019

Abstract

Large-scale neural recordings are becoming increasingly better at providing a window into functional neural networks in the living organism. Interpreting such rich data sets, however, poses fundamental statistical challenges. The neural field models of Wilson, Cowan and colleagues remain the mainstay of mathematical population modeling owing to their interpretable, mechanistic parameters and amenability to mathematical analysis. We developed a method based on moment closure to interpret neural field models as latent state-space point-process models, making mean field models amenable to statistical inference. We demonstrate that this approach can infer latent neural states, such as active and refractory neurons, in large populations. After validating this approach with synthetic data, we apply it to high-density recordings of spiking activity in the developing mouse retina. This confirms the essential role of a long lasting refractory state in shaping spatio-temporal properties of neonatal retinal waves. This conceptual and methodological advance opens up new theoretical connections between mathematical theory and point-process state-space models in neural data analysis.

Significance Developing statistical tools to connect single-neuron activity to emergent collective dynamics is vital for building interpretable models of neural activity. Neural field models relate single-neuron activity to emergent collective dynamics in neural populations, but integrating them with data remains challenging. Recently, latent state-space models have emerged as a powerful tool for constructing phenomenological models of neural population activity. The advent of high-density multi-electrode array recordings now enables us to examine large-scale collective neural activity. We show that classical neural field approaches can yield latent state-space equations and demonstrate inference for a neural field model of excitatory spatiotemporal waves that emerge in the developing retina.

1 Introduction

Neurons communicate with each other using electrical impulses, or spikes. Understanding the dynamics and physiology of collective spiking in large networks of neurons is a central challenge in modern neuroscience, with immense translational and clinical potential. Modern technologies such as high-density multi-electrode arrays (HDMEA) enable the simultaneous recording of the electrical activity of thousands of interconnected neurons, promising to offer invaluable insights

into neural dynamics at the network level. However, the resulting data is high-dimensional and frequently exhibits complex, non-linear dynamics, presenting formidable statistical challenges.

Due to the high complexity of the data, most analyses of neuronal population activity take a descriptive approach, adopting methods from statistical signal processing such as state space models (SSM; Paninski et al. 2010; Zhao and Park 2016, 2017; Sussillo et al. 2016; Aghagolzadeh and Truccolo 2016a; Linderman et al. 2016; Gao et al. 2016) or autoregressive generalized-linear point-process models (PP-GLM; Paninski 2004; Pillow et al. 2008; Truccolo et al. 2005; Truccolo 2016). Such methods are effective in capturing the population statistics of the system, but fail to provide mechanistic explanations of the underlying neural dynamics. While this phenomenological description is valuable and can aid many investigations, the inability to relate microscopic single-neuron activity to emergent collective dynamics severely limits the scope of these models to extract biological insights from these large population recordings.

Connecting single-neuron dynamics with population behavior has been the central focus of research within the theoretical neuroscience community over the last four decades. Neural field models (Amari, 1977; Wilson et al., 1972; Cowan, 2014; Bressloff, 2012) have been crucial in understanding how macroscopic firing dynamics in populations of neurons emerge from the microscopic state of individual neurons. Such models have found diverse applications including working memory (numerous studies, see Durstewitz et al. 2000 for a review), epilepsy (e.g. Zhang and Xiao 2018; Proix et al. 2018; González-Ramírez et al. 2015; Martinet et al. 2017), and hallucinations (e.g. Ermentrout and Cowan 1979; Bressloff et al. 2001; Rule et al. 2011), and have been successfully related to neuroimaging data such as Electroencepelography (EEG; Moran et al. 2013; Bojak et al. 2010; Pinotsis et al. 2012), Magnetoencephelography (MEG; Moran et al. 2013), electromyography (EMG; Nazarpour et al. 2012), and Functional Magnetic Resonance Imaging (fMRI; Bojak et al. 2010), which measure average signals from millions of neurons. Nevertheless, using neural field models to directly model HDMEA spiking data remains an open statistical problem: HDMEA recordings provide sufficient detail to allow modeling of individual neurons, yet the large number of neurons present prevents the adoption of standard approaches to non-linear data assimilation such as likelihood free inference.

In this paper, we bridge the data-model divide by developing a statistical framework for Bayesian modeling in neural field models. We build on recent advances in stochastic spatiotemporal modeling, in particular a recent result by Schnoerr et al. (2016) which showed that a spatiotemporal agent-based model of reaction-diffusion type, similar to the ones underpinning many neural field models, can be statistically approximated by a spatiotemporal point process with a defined intensity evolution equation. Subsequently, Rule and Sanguinetti (2018) illustrated a moment-closure approach for mapping stochastic models of neuronal spiking onto latent state-space models, preserving the essential coarse-timescale dynamics. Here, we demonstrate that a similar approach can yield state-space models for neural fields derived directly from a mechanistic microscopic description. This enables us to leverage large-scale spatiotemporal inference techniques (Cseke et al., 2016; Zammit-Mangion et al., 2012) to efficiently estimate an approximate likelihood, providing a measure of fit of the model to the data that can be exploited for data assimilation. Our approach is in spirit similar to latent variable models such as the Poisson Linear Dynamical System (PLDS; Macke et al. 2011; Aghagolzadeh and Truccolo 2016b; Smith and Brown 2003; Zammit-Mangion et al. 2011), with the important difference that the dynamics of the latent variables is non-linear and emerges directly from a stochastic description of single-neuron activity.

We apply this approach to HDMEA recordings of spontaneous activity from ganglion cells in the developing mouse retina (Maccione et al., 2014), showing that the calibrated model effectively captures the non-linear excitable phenomenon of coordinated, wave-like patterns of spiking (Meister et al., 1991) that have been considered in both discrete (Hennig et al., 2009a) and field-theoretic models before (Lansdell et al., 2014).

2 Results

2.1 Neural field models for refractoriness-mediated retinal waves

Most classical neural field models (Wilson and Cowan, 1972, 1973) consider two neuron states: neurons may be either actively spiking (A state), or quiescent (Q state), and any refractory effects are absorbed into the coarse-grained dynamics. In contrast, developmental retinal waves exhibit slow afterhyperpolarization (sAHP) potentials that mediate the long-timescale refractory effects underlying developmental retinal waves (Hennig et al., 2009b). To address this, we consider models that explicitly incorporate additional refractory (R) states (e.g. Buice and Cowan 2007a, 2009; Figure 1) into our neural field model.

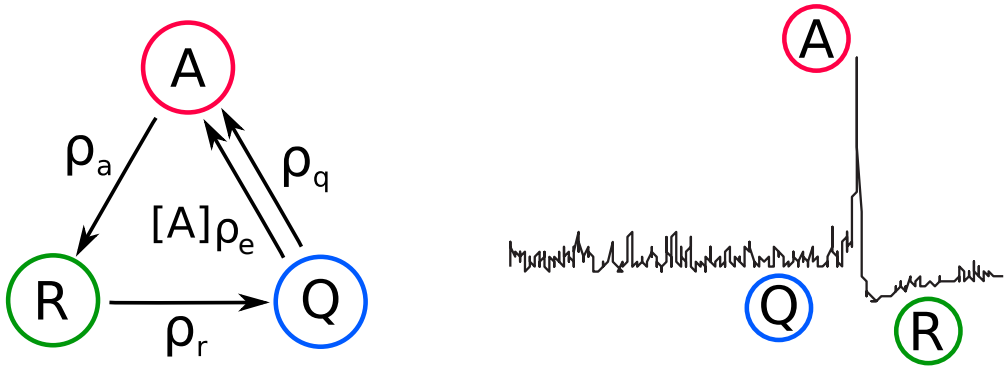

Figure 1: **3-state Quiescent-Active-Refractory (QAR) neural-field model.** Cells in the developing retina are modeled as having three activity states. Active cells (A; red) fire bursts of action potentials, before becoming refractory (R; green) for an extended period of time. Quiescent (Q; blue) cells may burst spontaneously, or may be recruited into a wave by other active cells. These three states are proposed to underlie critical multi-scale wave dynamics (Hennig et al., 2009b).

To develop a state-space formalism for inference and data assimilation in neural field models of retinal waves, we first use a master-equation approach (Buice and Cowan, 2007a; Ohira and Cowan, 1993; Bressloff, 2009) to define a three-state stochastic neural field model. We then outline a moment-closure approach (Schnoerr et al., 2017; Rule and Sanguinetti, 2018) to close a series expansion of network interactions in terms of higher moments (Buice et al., 2010), and obtain a second-order neural field model with field equations for both mean and covariance. We illustrate that a Langevin approximation (Riedler and Buckwar, 2013) of this model recapitulates spatiotemporal wave phenomena when sampled. Finally, to integrate this with spiking observations, we interpret the second-order neural field states as moments of a Gaussian process estimate

of latent neural activity driving conditionally Poisson spiking.

2.2 A stochastic three-state neural mass model

We consider a neural field model with three states as a generic model of a spiking neuron (Figure 1). In this model, the neuron is either actively spiking (A state), refractory (R state), or quiescent (Q state). We assume spontaneous, Poisson transitions between neural states, with a single quadratic pairwise interaction wherein active (A) cells excite nearby quiescent (Q) cells. Such a quadratic interaction can be viewed more generally as a locally-quadratic approximation of pairwise nonlinear excitatory interaction (Rule and Sanguinetti, 2018; Ale et al., 2013). Consider the following four state transitions of neurons:

$$\begin{aligned} Q &\xrightarrow{\rho_q} A \qquad & Q + A &\xrightarrow{\rho_e} A + A \\ A &\xrightarrow{\rho_a} R & R &\xrightarrow{\rho_r} Q \end{aligned} \tag{1}$$

Quiescent neurons may spike spontaneously with rate ρ_q, entering the active state; active neurons excite quiescent neurons with rate ρ_e, and active neurons become refractory with rate ρ_a. Refractory neurons become quiescent with rate ρ_r.

For illustration, consider the dynamics of a local (as opposed to spatially-extended) population of neurons. We may describe the fraction of neurons in each state in terms of a probability distribution $\Pr(Q, A, R)$ (Figure 2A), where we slightly abuse notation and use Q, A and R both as symbols for the neuron states and as variables counting the neurons in the corresponding states, i.e. non-negative integers. The time evolution of this probability distribution captures stochastic collective population dynamics, and is given by a master equation that describes the change in density for a given state $\{Q, A, R\}$ in terms of the probability of entering, minus the probability of leaving, said state:

$$\begin{aligned} \partial_t \Pr(Q,A,R) = {} & \Pr(Q, A{+}1, R{-}1)\rho_{\mathrm{a}}(A{+}1) && \text{(transition } A{\to}R\text{)} \\ & + \Pr(Q{-}1, A, R{+}1)\rho_{\mathrm{r}}(R{+}1) && \text{(transition } R{\to}Q\text{)} \\ & + \Pr(Q{+}1, A{-}1, R)\left[\rho_q + \rho_{\mathrm{e}}(A{-}1)\right](Q{+}1) && (Q{\to}A \text{ and } A{+}Q{\to}A{+}A) \\ & - \Pr(Q,A,R)\left[(\rho_{\mathrm{e}}A + \rho_{\mathrm{q}})Q + \rho_{\mathrm{a}}A + \rho_{\mathrm{r}}R\right] && \text{(outgoing transitions)} \end{aligned} \tag{2}$$

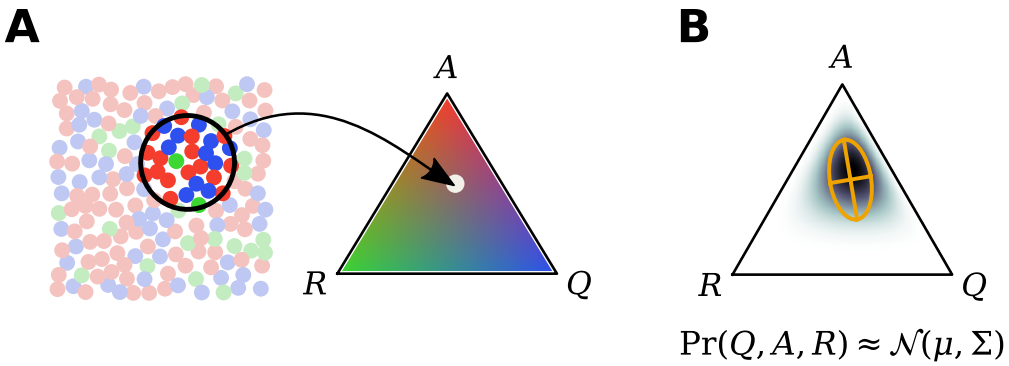

Figure 2: **Summarizing estimated neural state as population moments. A** The activity within a local spatial region (encircled, left) can be summarized by the fraction of cells in the quiescent (blue), active (red), and refractory (green) states (Q, A, R, right). **B** An estimate of the population state can be summarized as a probability distribution $\Pr(Q, A, R)$ over the possible proportions of neurons in each state. A Gaussian moment-closure approximates this distribution as Gaussian, with given mean and covariance (orange crosshairs).

Even in this simplified non-spatial scenario, no analytic solutions are known for the master equation. However, from Eq. 2 one can derive equations for the mean and covariance of the process. Due to the nonlinear excitatory interaction $Q{+}A \xrightarrow{\rho_e} A{+}A$, the evolution of the means is coupled to the covariance Σ_{AQ}, and the evolution of the covariance to the third moment. The moment equations are therefore not closed. We use Gaussian moment closure (Figure 2B) to approximate $\Pr(Q, A, R)$ with a multivariate normal distribution at each time-point, thereby replacing counts of neurons with continuous variables. The Gaussian moment closure approximation sets all cumulants beyond the variance to zero, yielding an expression for the third moment in terms of the first and second moments of Eq. 2 and giving closed ordinary differential equations for the means and covariances (Schnoerr et al., 2017; Rule and Sanguinetti, 2018). The evolution of the first moment (mean concentrations) is as follows:

$$\begin{aligned} \partial_t \langle Q \rangle &= r_{rq} - r_{qa} & r_{qa} &= \rho_q \langle Q \rangle + \rho_e \left(\langle A \rangle \langle Q \rangle + \Sigma_{AQ} \right) \\ \partial_t \langle A \rangle &= r_{qa} - r_{ar} & r_{ar} &= \rho_a \langle A \rangle \\ \partial_t \langle R \rangle &= r_{ar} - r_{rq} & r_{rq} &= \rho_r \langle R \rangle , \end{aligned} \tag{3}$$

where the rate variables $r_{(\cdot)(\cdot)}$ describe the population rates of each state transition, and $\langle \cdot \rangle$ denotes the expectation w.r.t. to $\Pr(Q, A, R)$. In the Gaussian moment-closure, the covariance of the estimated state distribution evolves as

$$\begin{aligned} \partial_t \Sigma &= \mathbf{J}\Sigma + \Sigma \mathbf{J}^T + \Sigma_{\text{noise}}, \\ \Sigma_{\text{noise}} &= \begin{bmatrix} r_{qa} + r_{rq} & -r_{qa} & -r_{rq} \\ -r_{qa} & r_{qa} + r_{ar} & -r_{ar} \\ -r_{rq} & -r_{ar} & r_{ar} + r_{qa} \end{bmatrix} \end{aligned} \tag{4}$$

where $\mathbf{J}$ is the Jacobian of the equations for the deterministic means in Eq. 3, and the Σ_{noise} fluctuations are Poisson and therefore proportional to the mean reactions rates (Eq. 3). Together, equations 3 and 4 provide approximate equations for the evolution of the first two moments of

the master equation (Eq. 2), expressed in terms of ordinary differential equations governing the mean and covariance of a multivariate Gaussian distribution. Here, we have illustrated equations for a 3-state system, but the approach is general and can be applied to any system composed of spontaneous and pairwise state transitions.

2.3 Generalization to spatial (neural field) system

So far we have considered a single local population. For a two-dimensional spatial system, the mean concentrations become density functions (fields) that depend on spatial coordinates $\mathbf{x} = (x_1, x_2)$, e.g. $\langle Q \rangle$ becomes $\langle \mathbf{Q}(\mathbf{x}) \rangle$. Similarly, the covariances become two-point correlation functions, e.g. $\Sigma_{QA}(\mathbf{x}, \mathbf{x}')$ denotes the covariance between the number of neurons in the quiescent state at location $\mathbf{x}$ and the number of neurons in the active state at location $\mathbf{x}'$. We introduce spatial interactions as an integral taken over nonlocal excitatory couplings, weighted by a Gaussian kernel that depends on the distance $||\Delta\mathbf{x}||$ with standard deviation σ_e:

$$\mathrm{k}(\Delta\mathbf{x}) \propto \exp(-||\Delta\mathbf{x}||^2/2\sigma_e^2). \tag{5}$$

With this coupling, the transition rate (compare to Eq. 3) from the quiescent to active state at position $\mathbf{x}$ is given by the integral:

$$\begin{aligned} r_{qa}(\mathbf{x}) &= \rho_q \langle Q(\mathbf{x}) \rangle + \rho_e \int k(\mathbf{x}-\mathbf{x}') \langle Q(\mathbf{x}) A(\mathbf{x}') \rangle \, \mathrm{d}\mathbf{x}' \\ \langle Q(\mathbf{x}) A(\mathbf{x}') \rangle &= \langle Q(\mathbf{x}) \rangle \langle A(\mathbf{x}') \rangle + \Sigma_{QA}(\mathbf{x}, \mathbf{x}') \end{aligned} \tag{6}$$

This kernel thus constitutes an effective coupling in the network, which could in principle reflect synaptic interactions, diffusing neurotransmitters, gap junction coupling, or combinations thereof. Equation 6 gives rise to integro-differential equations for the means and covariances of neural states (Methods: *Moment closure for the spatial system*). We verify that sampling from the spatially-extended system (Methods: *Langevin equations and sampling*) exhibits self-organized multi-scale wave phenomena (Figure 3).

2.4 Neural field models as latent-variable state-space models

The equations for the mean-fields and correlations can be integrated forward in time and used as a state-space model to explain population spiking activity (Figure 4; Methods: *Bayesian filtering*). In HDMEA recordings, we do not directly observe the intensity functions $\langle \mathbf{Q}(\mathbf{x}) \rangle$, $\langle \mathbf{A}(\mathbf{x}) \rangle$, and $\langle \mathbf{R}(\mathbf{x}) \rangle$. Instead, we observe the spikes that neurons emit, represented as a spatiotemporal point process $y(\mathbf{x}, t)$, where each spike has an associated time t and spatial location $\mathbf{x}$. We use a linear Poisson likelihood for which the point-process intensity

$$\lambda(\mathbf{x}, t) = \gamma(\mathbf{x}) \mathbf{A}(\mathbf{x}, t) + \beta(\mathbf{x}) \tag{7}$$

depends linearly on the number of active neurons $\mathbf{A}(\mathbf{x}, t)$ with gain γ and bias β. The combination of this Poisson observation model with the moment-closure state-space model yields a hidden Markov model for the latent neural field states (Figure 4). Given a spatiotemporal point-process of observed spikes, the latent neural field states and correlations can then be inferred using Bayesian filtering (Methods: *Bayesian filtering*).

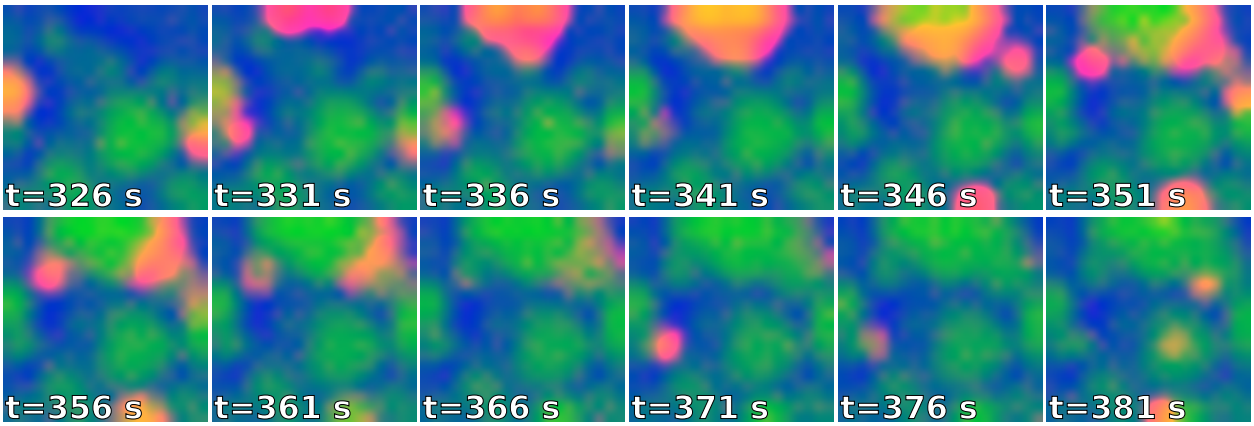

Figure 3: **Spatial 3-state neural-field model exhibits self-organized multi-scale wave phenomena.** Simulated example states at selected time-points on a $[0, 1]^2$ unit interval using 20×20 grid with effective population density of ρ=50 cells per simulation area, and rate parameters σ=7.5e-2, ρ_a=4e-1, ρ_r=3.2e-3, ρ_e=2.8e-2, and ρ_q=2.5e-1 (Methods: *Langevin equations and sampling*). As, for instance, in neonatal retinal waves, spontaneous excitation of quiescent cells (blue) lead to propagating waves of activity (red), which establish localized patches in which cells are refractory (green) to subsequent wave propagation. Over time, this leads to diverse patterns of waves at a range of spatial scales.

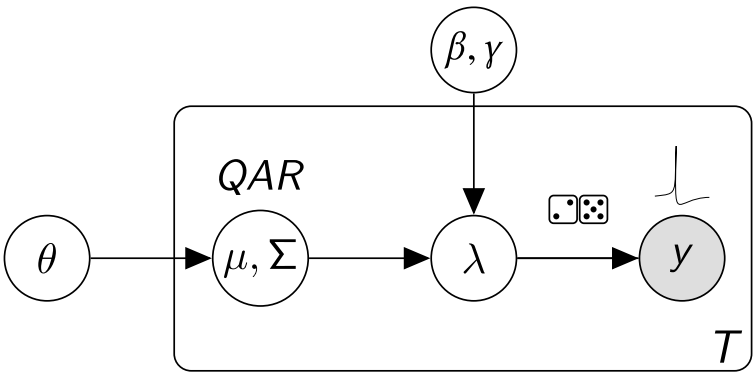

Figure 4: **Hidden Markov model for latent neural fields.** For all time-points T, state transition parameters θ=(ρ_q, ρ_a, ρ_r, ρ_e, σ) dictate the evolution of a multivariate Gaussian model μ, Σ of latent fields Q, A, R. The observation model (β, γ) is a linear map with adjustable gain and threshold, and reflects how field A couples to firing intensity λ. Point-process observations (spikes) y are Poisson with intensity λ.

The state-space formulation of the neural field model allows the latent states to be inferred from spiking point-process observations, as in a SSM framework. This in turn provides access to unobserved physiological states of neurons. To verify this SSM approach, we sampled ground-truth simulations using a Langevin approximation to the stochastic neural field equations (Methods: *Langevin equations and sampling*), and inferred latent neural states and confidence intervals via Bayesian filtering, using known parameters. In the example simulation (Figure 5), we used

a gain of γ=15 spikes/second per simulation area, corresponding to a relatively low spike rate, indicating that state inference can recover latent states in the presence of limited measurement information.

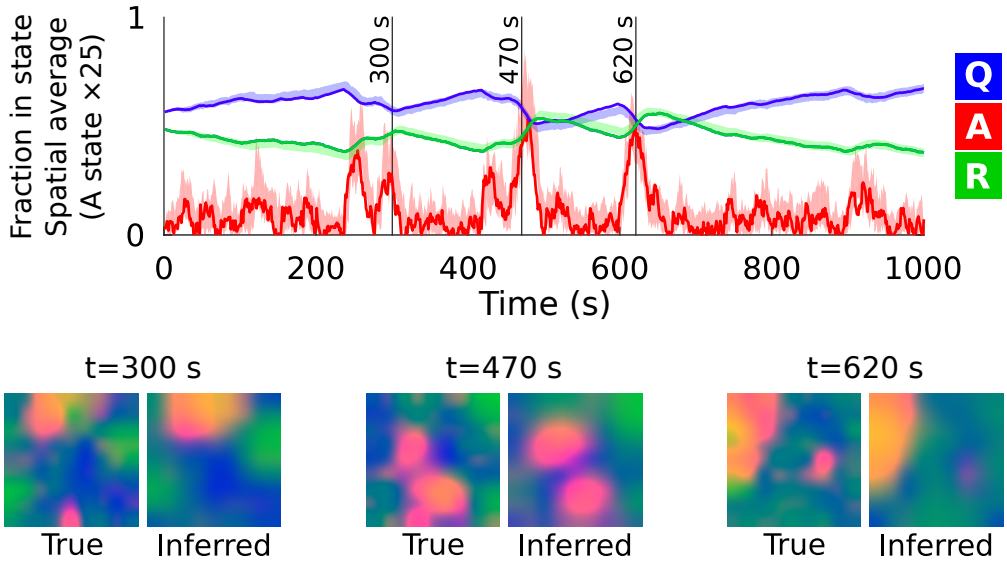

Figure 5: **State inference via filtering: ground-truth simulation.** Filtering recovers latent states in ground-truth simulated data. Spatially averaged state occupancy (**Q**, **A**, and **R**) (y-axis) is plotted over time (x-axis). Solid lines represent true values sampled from the model, and shaded regions represent the 95% confidence interval estimated by filtering. The active (**A**) state density has been scaled up by a factor of 25 for visualization. Colored plots (below) show the qualitative spatial organization of quiescent (blue), active (red), and refractory (green) neurons is recovered by filtering during example wave events. Model parameters are the same as Figure 3, with the exception of the spatial resolution, which has been reduced to a 9×9 grid. Conditionally-Poisson spikes were sampled with bias β=0 and gain γ=15 spikes/second per simulation area.

3 Application to retinal wave datasets

Having developed an interpretation of neural field equations as a latent-variable state-space model, we next applied this model to the analysis of spatiotemporal spiking data from spontaneous traveling wave activity occurring in the neonatal vertebrate retina (e.g. Figure 7; Sernagor et al. 2003; Hennig et al. 2009a; Blankenship et al. 2009; Meister et al. 1991; Zhou and Zhao 2000; Feller et al. 1996; Maccione et al. 2014).

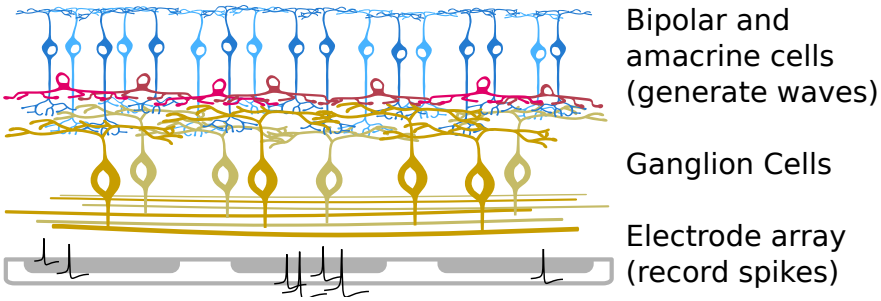

Figure 6: **Illustration of inner retina and recording setup.** Spontaneous retinal waves are generated in the inner retina via laterally interacting bipolar (blue) and amacrine (red) cells, depending on the developmental age. These waves activate Retinal Ganglion Cells (RGCs; yellow), the output cells of the retina. RGC electrical activity is recorded from the neonatal mouse retina via a 64×64 4096-electrode array with 42 μm spacing.

3.1 State inference in developmental retinal waves

During retinal development, the cell types that participate in wave generation change (Maccione et al., 2014; Sernagor et al., 2003; Zhou and Zhao, 2000), but the three-state model globally describes dynamics in the inner retina at all developmental stages (Figure 6). The Active (A) state describes a sustained bursting state, such as the depolarization characteristic of starburst amacrine cells (Figure 6) during cholinergic early-stage (Stage 2) waves between P0 and P9 (Feller et al., 1996; Zhou and Zhao, 2000), and late-stage (Stage 3) glutamate-dependent waves (Bansal et al., 2000; Zhou and Zhao, 2000). For example, Figure 7 illustrates spontaneous retinal wave activity recorded from a postnatal day 6 mouse pup (Stage 2). In addition, at least for cholinergic waves, the slow refractory state R is essential for restricting wave propagation into previously active areas (Zheng et al., 2006). We note that the multi-scale wave activity exhibited in the three-state neural field model (e.g. Figure 3) recapitulates the phenomenology of retinal wave activity explored in the discrete three-state model of Hennig et al. (2009b).

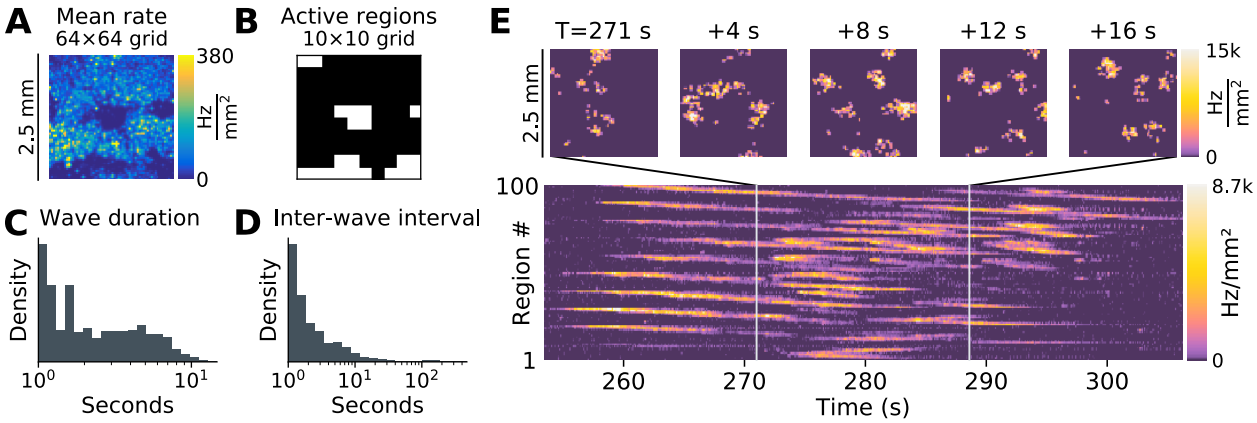

Figure 7: **Developmental retinal waves.** Example neonatal mouse retinal waves recorded on a 4096-electrode array on postnatal day 6. Recorded spikes were binned at 100 ms resolution, and assigned to 10×10 spatial regions for analysis. **A** Average firing rate from RGCs recorded across the retina (the central region devoid of recorded spikes is the optic disc). **B** Spatial regions with spiking activity detected for further analysis **C** Distribution of wave durations. To segment waves, spiking activity on each channel was segmented into "up" states (during wave activity) and "down" states (quiescent) using a two-state hidden Markov model with Poisson observations. **D** Average inter-wave interval. **E** Example wave event, traveling across multiple spatial regions and lasting for a duration of 16-20 seconds.

Using RGC spikes recorded with a 4,096 channel HDMEA (Figure 6), we demonstrate the practicality of latent-state inference using heuristically initialized rate parameters and illustrate an example of inference for a retinal wave dataset from postnatal day 11 (Stage 3) (Figure 8). For retinal wave inference, we employ a population-size normalized model (Methods: *System-size scaling*), so that the gain and bias may be set independently in normalized units, rather than depending on the local neuronal population size. Model parameters were initialized heuristically based on observed timescales at ρ_e=ρ_a=15, ρ_r=0.15, and σ=0.15. As in Lansdell et al. (2014), lateral interactions in our model reflect an effective coupling that combines both excitatory synaptic interactions and the putative effect of diffusing excitatory neurotransmitters, which has been shown to promote late-stage glutamatergic wave propagation (Blankenship et al., 2009). The spontaneous excitation rate ρ_q was set to zero, such that the spontaneous wave-initiation events are captured by the system as an extrinsic noise source. The Poisson noise was re-scaled to reflect an effective population size of 16 neurons/mm², significantly smaller than the true population density (Jeon et al., 1998). However, due to the recurrent architecture and correlated neuronal firing, the effective population size is expected to be smaller than the true population size. Equivalently, this amounts to assuming supra-Poisson scaling of fluctuations for the neural population responsible for retinal waves.

Bayesian filtering recovers the expected features of the retinal waves (Figure 8): the excito-excitatory transition $Q{+}A{\rightarrow}A{+}A$ and the onset of refractoriness $A{\rightarrow}R$ are rapid compared to the slow refractory dynamics, and therefore the A state is briefly occupied and mediates an effective $Q{\rightarrow}R$ transition during wave events. The second-order structure provided by the covariance is essential, as it allows us to model posterior variance (shaded regions in Figure 8), while also

capturing strong anti-correlations due to the conservation of reacting agents, and the effect of correlated fluctuations on the evolution of the means. Furthermore, spatial correlations allow localized RGC spiking events to be interpreted as evidence of regional (spatially-extended) latent neuronal activity.

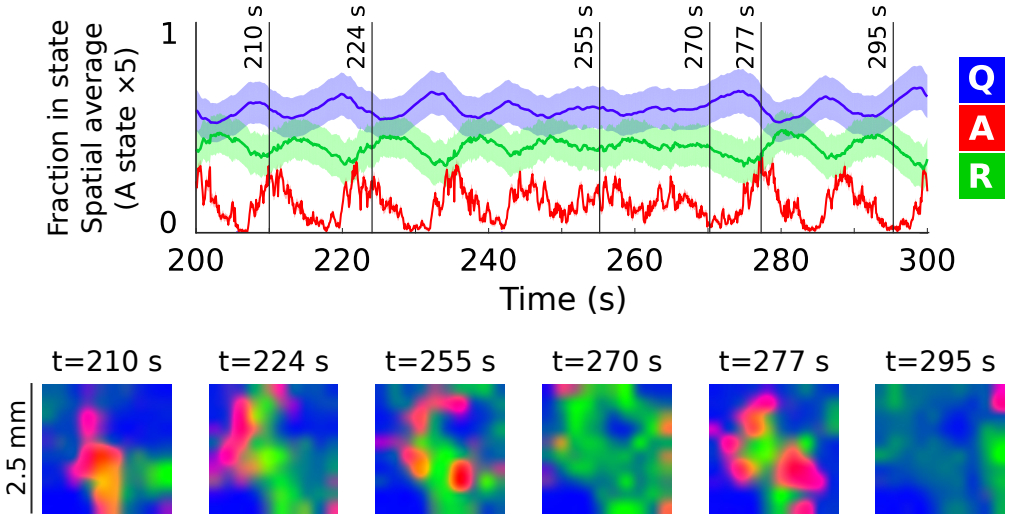

Figure 8: **State inference via filtering: retinal datasets.** *Filtering of spontaneous retinal waves (postnatal day 11).* Solid lines indicate inferred means, shaded regions the 95% confidence bound. The magnitude of the *A* state and the counts have been scaled up by a factor of 5 for visualization. Grey vertical lines indicate example time slices, which are shown in the colored plots below. Colored plots are the same as in Figure 5, with red, green, and blue reflecting (normalized) densities of active, refractory, and quiescent cells, respectively.

3.2 Open challenges in model identification

So far, we have demonstrated good recovery of states when the true rate parameters are known (Figure 5), and also shown that plausible latent-states can also be inferred from neural point-process datasets using heuristically initialized parameters (Figure 8). A natural question then is whether one can use the Bayesian state-space framework to estimate a posterior likelihood on the rate parameter values, and infer model parameters directly from data. At present, model inference remains very challenging for four reasons: under-constrained parameters, computational time complexity, numerical errors from successive approximations, and non-convexity in the joint posterior. It is worth reviewing these open challenges as they relate to open problems in machine learning and data assimilation.

First, the effective population size, the typical fraction of units in quiescent vs. refractory states, and the gain parameter mapping latent activations to spiking, are all essential to setting appropriate rates, and are not accessible from observation of RGC spiking alone. Without direct measurement or appropriate physiological priors on parameter values, one cannot recover a physiologically-realistic model. In effect, this means that a large number of equivalent systems can explain the observed RGC spiking activity, a phenomenon that has been termed "sloppiness" in biological systems (Transtrum et al., 2015; Panas et al., 2015). Indeed, Hennig et al. (2011) show that developmental waves are robust to pharmacological perturbations, suggesting that the

retina itself can use different configurations to achieve similar wave patterns. *Second*, although state inference is computationally feasible, parameter inference requires many thousands of state-inference evaluations. A Matlab implementation of state-inference running on a 2.9 GHz 8-core Xeon CPU can process ~85 samples/s for a 3-state system on a 10×10 spatial basis. For a thirty-minute recording of retinal wave activity, state inference is feasible, but repeated state inference for parameter inference is impractical. *Third*, model likelihood must be computed recursively, and is subject to well-known loss of numerical accuracy due to back-propagation through time (Pascanu et al., 2013; Bengio et al., 1994; Hochreiter et al., 2001). For early-stage retinal waves, the large separation of timescales between the refractory effects, and the fast-timescale excitation, also makes it difficult to estimate gradients for the slow-timescale parameters. Furthermore, the inferred likelihood approximated as the product of a large number of high-dimensional Laplace approximations (or similar Gaussian approximations, e.g. variational), which makes the inferred likelihood numerically unstable. *Fourth* and finally, the overall likelihood surface is not, in general, convex, and may contain multiple local optima. In addition, regions of parameters space can exhibit vanishing gradient for one or model parameters.

Overall, parameter inference via Bayesian filtering presents a formidable technical challenge that hinges upon several open problems for efficient model inference in high-dimensional spatiotemporal point process models undergoing latent, nonlinear dynamics. At present, it would seem that traditional parameter identification methods, based on mathematical expertise and matching observable physical quantities (e.g. wavefront speed, c.f. Lansdell et al. 2014), remain the best-available approach to model estimation. Nevertheless, the state-space formulation of neural field models enables Bayesian state inference from candidate neural field models, and opens the possibility of likelihood-based parameter inference in the future.

4 Discussion

In this work, we have demonstrated that classical neural-field models can be interpreted as state-space models. This is achieved by interpreting a second-order neural field model as defining equations on the first two moments of a latent-variable process that is coupled to spiking observations. In the state-space model interpretation, latent neural field states can be recovered from Bayesian filtering. This allows inferring the internal states of individual neurons in large networks based solely on recorded spiking activity, information that can experimentally only be obtained with whole cell recordings. We demonstrated successful state inference for simulated data, where the correct model and parameters were known. Next, we applied the model to large-scale recordings of developmental retinal waves. Here the correct latent state model is unknown, but a relatively simple three-state model with slow refractoriness is well motivated by experimental observations (Zheng et al., 2006). Consistent with previous work (Feller et al., 1997; Zheng et al., 2006; Godfrey and Swindale, 2007; Hennig et al., 2009a), the state inference revealed that activity-dependent refractoriness restricts the spatial spreading of waves.

In contrast to phenomenological latent state-space models, the latent states here are motivated by an (albeit simplified) description of single-neuron dynamics, and the state-space equations arise directly from considering the evolution of collective activity as a stochastic process.

In the example explored here, we use Gaussian moment-closure to arrive at a second-order approximation of the distribution of latent states and their evolution. In principle, other distri-

butional assumptions may also be used to close the moment expansion. Other mathematical approaches that yield second-order models could also be employed, for example the linear noise approximation (Van Kampen, 1992) or diagrammatic perturbation (Buice and Cowan, 2007b; Ocker et al., 2017). The approach applied here to a three-state system can generally be applied to systems composed of linear and quadratic state transitions. Importantly, systems with only linear and pairwise (quadratic) interactions can be viewed as a locally-quadratic approximation of a more general smooth nonlinear system (Ale et al., 2013), and Gaussian moment closure therefore provides a general approach to deriving approximate state-space models in neural population dynamics.

The state-space interpretation of neural field models opens up future work to leverage the algorithmic tools of SSM estimation for data assimilation with spiking point-process datasets. However, challenges remain regarding the retinal waves explored here, and future work is needed to address these challenges. Model likelihood estimation is especially challenging. Despite this, the connection between neural-field models and state-space models derived here will allow neural field modeling to incorporate future advances in estimating recursive, nonlinear, spatiotemporal models. We also emphasize that some of the numerical challenges inherent to high-dimensional spatially extended neural field models do not apply to simpler, low-dimensional neural mass models, and the moment-closure framework may therefore provide a practical avenue to parameter inference in such models.

In summary, this report connects neural field models, which are grounded in models of stochastic population dynamics, to latent state-space models for population spiking activity. This connection opens up new approaches to fitting neural field models to spiking data. We hope that this interpretation is a step toward the design of coarse-grained models of neural activity that have physically interpretable parameters, have physically measurable states, and retain an explicit connection between microscopic activity and emergent collective dynamics. Such models will be essential for building models of collective dynamics that can predict the effects of manipulations on single-cells on emergent population activity.

Acknowledgements: Funding provided by EPSRC EP/L027208/1 *Large scale spatio-temporal point processes: novel machine learning methodologies and application to neural multi-electrode arrays*. We thank Gerrit Hilgen for important discussions in establishing biologically-plausible parameter regimes for the three-state model. We thank Evelyne Sernagor for the retinal wave datasets, as well as ongoing advice and invaluable feedback on the manuscript.

5 Methods

5.1 Data acquisition and preparation

Example retinal wave datasets are taken from Maccione et al. (2014). For analysis, spikes were binned in time at a 100 ms resolution. Spiking activity over time at each region was segmented into wave-like and quiescent states using a two-state hidden Markov model with a Poisson observations. Regions without spiking observations were excluded. To address heterogeneity in the Retinal Ganglion Cell (RGC) outputs, the observation model was adapted to each spatial region based on firing rates. Background activity was used to establish per-region biases, defined as the

mean activity in a region in quiescent periods. The scaling between latent state and firing rate (gain) was adjusted locally based on the mean firing rate during wave events. The overall (global) gain for the observation model was then adjusted so that 99% of wave events on all channels correspond to an expected A-state fraction no greater than one.

5.2 Moment closure for the spatial system

To extend the moment equations (Eq. 3, 4) to the spatial system, we denote the intensity fields as $\mathbf{Q}$, $\mathbf{A}$, and $\mathbf{R}$, which are now vectors with spatial indices (in continuum: scalar functions of coordinates $\mathbf{x}$). In contrast to the non-spatial system, the number of active (bursting) neurons $\mathbf{A}$ is replaced by a weighted sum over the number of active neurons in a local neighborhood. This is modeled by a distance-dependent coupling kernel $K(\Delta\mathbf{x})$ (Eq. 6). Denote the convolution integral in equation 6 as a linear operator $\mathbf{K}$ such that

$$\mathbf{KA} = K(\Delta\mathbf{x}) * A(\mathbf{x}). \tag{8}$$

Using the notation of Eq. 8, the excitatory reaction rate mediated by the interaction between active and quiescent cells is given by the product $\rho_e(\mathbf{KA}) \circ \mathbf{Q}$, where $\circ$ denotes element-wise (in continuum: function) multiplication. For the time evolution of the first moment (mean intensity) of $\mathbf{Q}$ in the spatial system, one therefore considers the expectation $\left\langle \mathbf{KAQ}^\top \right\rangle$, as opposed to $\langle AQ \rangle$ in the non-spatial system. Since $\mathbf{K}$ is a linear operator, and the extension of the Gaussian state-space model over the spatial domain $\mathbf{x}$ is a Gaussian process, the second moment of the nonlocal interactions $\mathbf{KA}$ with $\mathbf{Q}$ can be obtained in the same way as one obtains the correlation for a linear projection of a multivariate Gaussian:

$$\begin{aligned} \langle \mathbf{KAQ}^\top \rangle &= \mathbf{K} \langle \mathbf{AQ}^\top \rangle \\ &= \mathbf{K} \left(\Sigma_{\mathbf{A},\mathbf{Q}} + \langle \mathbf{A} \rangle \langle \mathbf{Q} \rangle^\top \right) \end{aligned} \tag{9}$$

The resulting equations for the spatial means are similar to the nonspatial system (Eq. 3), with the exception of the rate at which quiescent cells enter the active state, which now includes spatial coupling terms:

$$\begin{aligned} r_{qa} &= \rho_q \langle \mathbf{Q} \rangle + \rho_e \operatorname{Diag} \left[\langle \mathbf{KAQ}^\top \rangle \right] \\ &= \rho_q \langle \mathbf{Q} \rangle + \rho_e \operatorname{Diag} \left[\mathbf{K} \left(\Sigma_{\mathbf{A},\mathbf{Q}} + \mathbf{K} \langle \mathbf{A} \rangle \langle \mathbf{Q} \rangle^\top \right) \right] \\ &= \rho_q \langle \mathbf{Q} \rangle + \rho_e \left[\operatorname{Diag} \left(\mathbf{K} \Sigma_{\mathbf{A},\mathbf{Q}} \right) + \mathbf{K} \langle \mathbf{A} \rangle \circ \langle \mathbf{Q} \rangle \right], \end{aligned} \tag{10}$$

The linear reactions remain local, and so the linear contribution to the Jacobian (Eq. 4) is similar to the non-spatial case. Nonlocal interaction terms, however, emerge in the nonlinear contribution to the Jacobian:

$$\mathbf{J}_{\text{nonlinear}} = \rho_e \begin{bmatrix} -\operatorname{Diag}(\mathbf{K} \langle \mathbf{A} \rangle) & -\operatorname{Diag}(\langle \mathbf{Q} \rangle \mathbf{K}) & 0 \\ \operatorname{Diag}(\mathbf{K} \langle \mathbf{A} \rangle) & \operatorname{Diag}(\langle \mathbf{Q} \rangle \mathbf{K}) & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{11}$$

5.3 Basis projection

The continuous neural field equations are simulated by projection onto a finite spatial basis B. Each basis element is an integral over a spatial volume. Means for each basis element are defined

as an integral over said volume, and correlations are defined as a double integral. For example, consider the number of quiescent neurons associated with the i^{th} basis function, Q_i. The mean $\langle Q_i \rangle$ and covariance Σ_{QA}^{ij} between the quiescent and active states are given by the projections:

$$\begin{aligned} \langle Q_i \rangle &= \int B_i(\mathbf{x}) \mathbf{Q}(\mathbf{x})\, d\mathbf{x} \\ \Sigma_{QA}^{ij} &= \iint B_i(\mathbf{x}) B_j(\mathbf{x}') \Sigma_{\mathbf{QA}}(\mathbf{x}, \mathbf{x}')\, d\mathbf{x}\, d\mathbf{x}', \end{aligned} \tag{12}$$

where $\mathbf{x}$ and $\mathbf{x}'$ range over spatial coordinates as in Eq. 5 and 6. When selecting a basis B, assumptions must be made about the minimum spatial scale to model. A natural choice is the radius of lateral (i.e. spatially nonlocal) interactions in the model σ_e (Eq. 5), as structure below this scale is attenuated by the averaging over many nearby neurons in the dendritic inputs.

5.4 Langevin equations and sampling

For ground-truth stimulations, we sample from a hybrid stochastic model derived from a Langevin approximation to the three-state neural field equation. In the Langevin approximation, the deterministic evolution of the state is given by the mean-field equations (Eq. 3 for local reactions, Eq. 10 for spatial excitation), and the (continuous) stochastic noise for Poisson state transitions is given by second-order terms (using local noise as in Eq. 4; see Schnoerr et al. 2017 for further details). Spontaneous wave initiation events are too rare to approximate as Gaussian, and instead are sampled as Poisson (shot) noise:

$$r_q(t) \sim \text{Poisson}(\rho_q \cdot dt) \cdot \delta(t), \tag{13}$$

where $\delta(t)$ is a Dirac delta (impulse). To avoid uniform spontaneous excitation, the excito-excitatory reaction rate is adjusted by a small finite threshold ϑ, i.e. $r_{qa} \leftarrow \max(0, r_{qa} - \vartheta)$ in Eq. 10. For our simulations (e.g. Figure 3), we let ϑ=8e-3. For the nonspatial system, the hybrid stochastic differential equation equation is:

$$\begin{bmatrix} dQ \\ dA \\ dR \end{bmatrix} = \left(\begin{bmatrix} -r_q(t) & 0 & \rho_r \\ r_q(t) & -\rho_a & 0 \\ 0 & \rho_a & -\rho_r \end{bmatrix} \begin{bmatrix} Q \\ A \\ R \end{bmatrix} + \rho_e \begin{bmatrix} -QA \\ QA \\ 0 \end{bmatrix} \right) dt + \Sigma_{\text{noise}}^{1/2} dW, \tag{14}$$

where Σ_{noise} is the fluctuation noise covariance as in Equation 4 (with ρ_q excluded, as it is addressed by the shot noise, Equation 13), and dW is the derivative of a multidimensional standard Wiener process, i.e. a spherical (white) Gaussian noise source. The deterministic component of the Langevin equation can be compared to Equation 3 for the means of the nonspatial system in the moment-closure system (without the covariance terms).

The stochastic differential equation for the spatial system is similar, amounting to a collection of local populations coupled through a spatial interaction kernel $\mathbf{K}$, and follows the same derivation used when extending the moment-closure to the spatial case (Equation 6 and Methods: *Moment closure for the spatial system*, Equations 9-11). Note that fluctuations are scaled by $\sqrt{dt\, \Delta\mathbf{x}}$ in the spatiotemporal implementation of the Euler-Maruyama method, where $\Delta\mathbf{x}$ is the volume of the spatial basis functions used to approximate the spatial system (See Methods: *System-size scaling for further detail).* When sampling from such systems using the Euler-Maruyama algorithm, the issue of negative intensities can be handled by the complex chemical Langevin equation (Schnoerr et al., 2014).

5.5 Bayesian filtering

Having established an approach to approximate moments and sample trajectories, we now discuss the Bayesian filtering methodology which allows us to incorporate observations in the estimation of the latent states. Suppose we have measurements $y_0, \ldots, y_N$ of the latent state x at time $t_0, \ldots, t_N$, given by a measurement process $\Pr(y_i|x_{t_i})$, which in our case is given by the point-process likelihood in Eq. 7. Bayesian filtering allows us to recursively estimate the *filtering distribution* $\Pr(x_{t_i}|y_i, \ldots, y_0)$ at time t_i, i.e. the posterior state probability at time t_i given the current and all previous observations. The procedure works by the following iterative scheme: i) suppose we know the filtering distribution $\Pr(x_{t_i}|y_i, \ldots, y_0)$ at time t_i. Solving the dynamics forward in time up to t_{i+1} gives the predictive distribution $\Pr(x_t|y_i, \ldots, y_0)$ for all times $t_i < t \leq t_{t+1}$. ii) at the time t_{i+1} the measurement y_{i+1} needs to be taken into account which can be done by means of the Bayes update:

$$\Pr(x_{i+1}|y_{i+1}, \ldots, y_0) = \frac{\Pr(y_{i+1}|x_{i+1}) \Pr(x_{i+1}|y_i, \ldots, y_0)}{\Pr(y_{i+1}|y_i, \ldots, y_0)}, \tag{15}$$

where we have used the Markov property and $\Pr(y_{i+1}|x_{i+1}, y_i, \ldots, y_0) = \Pr(y_{i+1}|x_{i+1})$ to obtain the right hand side. Eq. (15) gives the filtering $\Pr(x_{t_{i+1}}|y_{i+1}, \ldots, y_0)$ at time t_{i+1} which serves as the input of the next i step. Performing steps i) and ii) iteratively hence provides the filtering distribution for all times $t_0 \leq t \leq t_n$.

For our neural field model we have to compute both steps approximately: to obtain the predictive distribution in step i) we integrate forward the differential equations for mean and covariance derived from moment-closure (Eq. (3), (4) and Methods: *Moment closure for the spatial system*). In practice, we convert the continuous-time model to discrete time. If $F_{\partial t}$ denotes the local linearization of the mean dynamics in continuous time such that $\partial_t \mu(t) = F_{\partial t} \mu(t)$, then the approximated discrete-time forward operator is

$$F_{\Delta t} = \exp(F_{\partial t} \Delta t) \approx I + F_{\partial t} \Delta t. \tag{16}$$

We update the covariance using this discrete-time forward operator, combined with an Euler integration step for the Poisson fluctuations. A small constant diagonal regularization term Σ_{reg} can be added, if needed, to improve stability. The resulting equations read:

$$\begin{aligned} \mu_{t+\Delta|t} &= F_{\Delta t} \mu_t \\ \Sigma_{t+\Delta|t} &= F_{\Delta t} \Sigma_t F_{\Delta t}^T + \Sigma_t^{\text{noise}} \cdot \Delta t + \Sigma_{\text{reg}}. \end{aligned} \tag{17}$$

This form is similar to the update for a discrete-time Kalman filter (Kalman et al., 1960; Kalman and Bucy, 1961), the main difference being that the dynamics between observation times are non-linear and obey the nonlinear moment equations.

Consider next the measurement update of step ii) in Eq. (15). Since the Gaussian model for the latent states x is not conjugate with the Poisson distribution for observations y, we approximate the posterior $\Pr(x_{i+1}|y_{i+1}, \ldots, y_0)$ using the Laplace approximation (e.g. Paninski et al. 2010; Macke et al. 2011). The Laplace-approximated measurement update is computed using a Newton-Raphson algorithm. The measurement update is constrained to avoid negative values in the latent fields by adding a ε/x potential (compare to the log-barrier approach; Nazarpour et al. 2012), which ensures that the objective function gradient points away from this constraint

boundary, where x is the intensity of any of the three fields. The gradients and Hessian for the posterior measurement log-likelihood $\ln \mathcal{L}$ are

$$\begin{aligned} -\ln \mathcal{L} &= \tfrac{1}{2}(x-\mu)^T\Sigma^{-1}(x-\mu) + v(\gamma x+\beta) - y\ln(\gamma x+\beta) \\ -\frac{\partial \ln \mathcal{L}}{\partial x} &= \Sigma^{-1}(x-\mu) + v\gamma - y\left(\frac{\gamma}{\gamma x+\beta}\right) \\ -\frac{\partial^2 \ln \mathcal{L}}{\partial x^2} &= \Sigma^{-1} + y\left(\frac{\gamma}{\gamma x+\beta}\right)^2, \end{aligned} \tag{18}$$

where x is the latent state with prior mean μ and covariance Σ, and couples to point-process observations y linearly with gain γ and bias β as in Eq. 7. The parameter $v{=}\Delta x^2{\cdot}\Delta t$ is the spatiotemporal volume of the basis function of spatial region over which counts are observed.

5.6 System-size scaling

For clarity, the derivations in this paper are presented for a population of neurons with a known, finite size, such that the fields $\mathbf{Q}(\mathbf{x})$, $\mathbf{A}(\mathbf{x})$, and $\mathbf{R}(\mathbf{x})$ have units of *neurons*. In practice, the population size Ω of neurons is not known, and it becomes expedient to work in normalized intensities, where $\mathbf{Q}(\mathbf{x})$, $\mathbf{A}(\mathbf{x})$, and $\mathbf{R}(\mathbf{x})$ represent the *fraction* of neurons in a given state between 0 and 1, and $\mathbf{Q}(\mathbf{x}) + \mathbf{A}(\mathbf{x}) + \mathbf{R}(\mathbf{x}){=}1$. In this normalized model for population size Ω, quadratic interaction parameters (like ρ_e) as well as the gain are multiplied by Ω, to reflect the rescaled population. In contrast, noise variance should be *divided* by Ω to account for the fact that the coefficient of variation decreases as population size increases. The infinitesimal neural-field limit for the second-order model, however, is ill-defined. This is because, while the mean-field equations scale with the population size $O(\Omega)$, the standard deviation of Poisson fluctuations scales with the square root of the population size $O(\sqrt{\Omega})$. The ratio of fluctuations to the mean (coefficient of variation) therefore scales as $O(1/\sqrt{\Omega})$, which diverges as $\Omega{\to}0$.

This divergence is not an issue in practice as all numerical simulations are implemented on a set of basis functions with finite nonzero volumes, associated with finite nonzero population sizes. Even in the limit where fluctuations would begin to diverge, one can treat the neural field equations as if defined over a continuous set of overlapping basis functions with nonzero volume. In this approach, adjacent spatial areas also experience correlated fluctuations. Consider Poisson fluctuations as entering with some rate-density $\sigma^2(\mathbf{x})$. The observed noise variances and covariances, projected onto basis functions $B_i(\mathbf{x})$ and $B_j(\mathbf{x})$, are:

$$\Sigma^{\text{noise}}_{i,j} = \int B_i(\mathbf{x})B_j(\mathbf{x})\sigma^2(\mathbf{x})d\mathbf{x} \tag{19}$$

In effect, one defines a population density $\rho(\mathbf{x})$. The effective population size for a given basis function is then

$$\Omega_i = \int B_i(\mathbf{x})\rho(\mathbf{x})d\mathbf{x} \tag{20}$$

If the population density is uniform, and if basis functions have a constant volume v, we can write this more simply as $\Omega = v\rho$. Therefore, when normalizing the model for system size, the contributions of basis function volume cancel and the noise should be scaled simply as $1/\rho$.

References

Aghagolzadeh, M. and Truccolo, W. (2016a). Inference and decoding of motor cortex low-dimensional dynamics via latent state-space models. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 24(2):272–282.

Aghagolzadeh, M. and Truccolo, W. (2016b). Inference and Decoding of Motor Cortex Low-Dimensional Dynamics via Latent State-Space Models. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 24(2):272–282.

Ale, A., Kirk, P., and Stumpf, M. P. (2013). A general moment expansion method for stochastic kinetic models. *The Journal of Chemical Physics*, 138(17):174101.

Amari, S.-i. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27(2):77–87.

Bansal, A., Singer, J. H., Hwang, B. J., Xu, W., Beaudet, A., and Feller, M. B. (2000). Mice lacking specific nicotinic acetylcholine receptor subunits exhibit dramatically altered spontaneous activity patterns and reveal a limited role for retinal waves in forming on and off circuits in the inner retina. *Journal of Neuroscience*, 20(20):7672–7681.

Bengio, Y., Simard, P., and Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166.

Blankenship, A. G., Ford, K. J., Johnson, J., Seal, R. P., Edwards, R. H., Copenhagen, D. R., and Feller, M. B. (2009). Synaptic and extrasynaptic factors governing glutamatergic retinal waves. *Neuron*, 62(2):230–241.

Bojak, I., Oostendorp, T. F., Reid, A. T., and Kötter, R. (2010). Connecting Mean Field Models of Neural Activity to EEG and fMRI Data. *Brain Topography*, 23(2):139–149.

Bressloff, P. C. (2009). Stochastic neural field theory and the system-size expansion. *SIAM Journal on Applied Mathematics*, 70(5):1488–1521.

Bressloff, P. C. (2012). Spatiotemporal dynamics of continuum neural fields. *Journal of Physics A: Mathematical and Theoretical*, 45(3):033001.

Bressloff, P. C., Cowan, J. D., Golubitsky, M., Thomas, P. J., and Wiener, M. C. (2001). Geometric visual hallucinations, euclidean symmetry and the functional architecture of striate cortex. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 356(1407):299–330.

Buice, M. A. and Cowan, J. D. (2007a). Field-theoretic approach to fluctuation effects in neural networks. *Physical Review E*, 75(5):051919.

Buice, M. A. and Cowan, J. D. (2007b). Field-theoretic approach to fluctuation effects in neural networks. *Physical Review E*, 75(5):051919.

Buice, M. A. and Cowan, J. D. (2009). Statistical mechanics of the neocortex. *Progress in Biophysics and Molecular Biology*, 99(2-3):53–86.

Buice, M. A., Cowan, J. D., and Chow, C. C. (2010). Systematic fluctuation expansion for neural network activity equations. *Neural Computation*, 22(2):377–426.

Cowan, J. (2014). A Personal Account of the Development of the Field Theory of Large-Scale Brain Activity from 1945 Onward. In *Neural Fields*, pages 47–96. Springer Berlin Heidelberg, Berlin, Heidelberg.

Cseke, B., Zammit-Mangion, A., Heskes, T., and Sanguinetti, G. (2016). Sparse approximate inference for spatio-temporal point process models. *Journal of the American Statistical Association*, 111(516):1746–1763.

Durstewitz, D., Seamans, J. K., and Sejnowski, T. J. (2000). Neurocomputational models of working memory. *Nature Neuroscience*, 3(11s):1184.

Ermentrout, G. B. and Cowan, J. D. (1979). A mathematical theory of visual hallucination patterns. *Biological Cybernetics*, 34(3):137–150.

Feller, M. B., Butts, D. A., Aaron, H. L., Rokhsar, D. S., and Shatz, C. J. (1997). Dynamic processes shape spatiotemporal properties of retinal waves. *Neuron*, 19(2):293–306.

Feller, M. B., Wellis, D. P., Stellwagen, D., Werblin, F. S., and Shatz, C. J. (1996). Requirement for cholinergic synaptic transmission in the propagation of spontaneous retinal waves. *Science*, 272(5265):1182–1187.

Gao, Y., Archer, E. W., Paninski, L., and Cunningham, J. P. (2016). Linear dynamical neural population models through nonlinear embeddings. In *Advances in Neural Information Processing Systems*, pages 163–171.

Godfrey, K. B. and Swindale, N. V. (2007). Retinal wave behavior through activity-dependent refractory periods. *PLoS Computational Biology*, 3(11):e245.

González-Ramírez, L., Ahmed, O., Cash, S., Wayne, C., and Kramer, M. (2015). A biologically constrained, mathematical model of cortical wave propagation preceding seizure termination. *PLoS Computational Biology*, 11(2):e1004065–e1004065.

Hennig, M. H., Adams, C., Willshaw, D., and Sernagor, E. (2009a). Early-stage waves in the retinal network emerge close to a critical state transition between local and global functional connectivity. *Journal of Neuroscience*, 29(4):1077–1086.

Hennig, M. H., Adams, C., Willshaw, D., and Sernagor, E. (2009b). Early-Stage Waves in the Retinal Network Emerge Close to a Critical State Transition between Local and Global Functional Connectivity. *Journal of Neuroscience*, 29(4).

Hennig, M. H., Grady, J., van Coppenhagen, J., and Sernagor, E. (2011). Age-dependent homeostatic plasticity of gabaergic signaling in developing retinal networks. *Journal of Neuroscience*, 31(34):12159–12164.

Hochreiter, S., Bengio, Y., Frasconi, P., Schmidhuber, J., et al. (2001). Gradient flow in recurrent nets: the difficulty of learning long-term dependencies.

Jeon, C.-J., Strettoi, E., and Masland, R. H. (1998). The major cell populations of the mouse retina. *Journal of Neuroscience*, 18(21):8936–8946.

Kalman, R. E. and Bucy, R. S. (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83(1):95–108.

Kalman, R. E. et al. (1960). Contributions to the theory of optimal control. *Boletín de la Sociedad Matemática Mexicana*, 5(2):102–119.

Lansdell, B., Ford, K., Kutz, J. N., Rokhsar, D., and Plenz, D. (2014). A Reaction-Diffusion Model of Cholinergic Retinal Waves. *PLOS Computational Biology*, 10(12):e1003953.

Linderman, S. W., Tucker, A., and Johnson, M. J. (2016). Bayesian latent state space models of neural activity. *Computational and Systems Neuroscience (Cosyne) Abstracts*.

Maccione, A., Hennig, M. H., Gandolfo, M., Muthmann, O., Coppenhagen, J., Eglen, S. J., Berdondini, L., and Sernagor, E. (2014). Following the ontogeny of retinal waves: pan-retinal recordings of population dynamics in the neonatal mouse. *The Journal of physiology*, 592(7):1545–1563.

Macke, J. H., Buesing, L., Cunningham, J. P., Byron, M. Y., Shenoy, K. V., and Sahani, M. (2011). Empirical models of spiking in neural populations. In *Advances in Neural Information Processing Systems*, pages 1350–1358.

Martinet, L.-E., Fiddyment, G., Madsen, J., Eskandar, E., Truccolo, W., Eden, U. T., Cash, S., and Kramer, M. A. (2017). Human seizures couple across spatial scales through travelling wave dynamics. *Nature Communications*, 8:14896.

Meister, M., Wong, R., Baylor, D. A., and Shatz, C. J. (1991). Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina. *Science*, 252(5008):939–943.

Moran, R., Pinotsis, D. A., and Friston, K. (2013). Neural masses and fields in dynamic causal modeling. *Frontiers in Computational Neuroscience*, 7:57.

Nazarpour, K., Ethier, C., Paninski, L., Rebesco, J. M., Miall, R. C., and Miller, L. E. (2012). EMG Prediction From Motor Cortical Recordings via a Nonnegative Point-Process Filter. *IEEE Transactions on Biomedical Engineering*, 59(7):1829–1838.

Ocker, G. K., Josić, K., Shea-Brown, E., and Buice, M. A. (2017). Linking structure and activity in nonlinear spiking networks. *PLoS Computational Biology*, 13(6):e1005583.

Ohira, T. and Cowan, J. D. (1993). Master-equation approach to stochastic neurodynamics. *Physical Review E*, 48(3):2259.

Panas, D., Amin, H., Maccione, A., Muthmann, O., van Rossum, M., Berdondini, L., and Hennig, M. H. (2015). Sloppiness in spontaneously active neuronal networks. *Journal of Neuroscience*, 35(22):8480–8492.

Paninski, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15(4):243–262.

Paninski, L., Ahmadian, Y., Ferreira, D. G., Koyama, S., Rad, K. R., Vidne, M., Vogelstein, J., and Wu, W. (2010). A new look at state-space models for neural data. *Journal of Computational Neuroscience*, 29(1-2):107–126.

Pascanu, R., Mikolov, T., and Bengio, Y. (2013). On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pages 1310–1318.

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E., and Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995.

Pinotsis, D., Moran, R., and Friston, K. (2012). Dynamic causal modeling with neural fields. *Neuroimage*, 59(2):1261–1274.

Proix, T., Jirsa, V. K., Bartolomei, F., Guye, M., and Truccolo, W. (2018). Predicting the spatiotemporal diversity of seizure propagation and termination in human focal epilepsy. *Nature Communications*, 9(1):1088.

Riedler, M. G. and Buckwar, E. (2013). Laws of large numbers and langevin approximations for stochastic neural field equations. *The Journal of Mathematical Neuroscience*, 3(1):1.

Rule, M. and Sanguinetti, G. (2018). Autoregressive point processes as latent state-space models: A moment-closure approach to fluctuations and autocorrelations. *Neural Computation*, 30(10):2757–2780. PMID: 30148704.

Rule, M., Stoffregen, M., and Ermentrout, B. (2011). A model for the origin and properties of flicker-induced geometric phosphenes. *PLoS Comput. Biol*, 7(9):e1002.

Schnoerr, D., Grima, R., and Sanguinetti, G. (2016). Cox process representation and inference for stochastic reaction-diffusion processes. *Nature Communications*, 7:11729.

Schnoerr, D., Sanguinetti, G., and Grima, R. (2014). The complex chemical langevin equation. *The Journal of Chemical Physics*, 141(2):07B606_1.

Schnoerr, D., Sanguinetti, G., and Grima, R. (2017). Approximation and inference methods for stochastic biochemical kinetics—a tutorial review. *Journal of Physics A: Mathematical and Theoretical*, 50(9):093001.

Sernagor, E., Young, C., and Eglen, S. J. (2003). Developmental modulation of retinal wave dynamics: shedding light on the gaba saga. *Journal of Neuroscience*, 23(20):7621–7629.

Smith, A. C. and Brown, E. N. (2003). Estimating a state-space model from point process observations. *Neural Computation*, 15(5):965–991.

Sussillo, D., Jozefowicz, R., Abbott, L., and Pandarinath, C. (2016). LFADS-latent factor analysis via dynamical systems. *arXiv preprint arXiv:1608.06315*.

Transtrum, M. K., Machta, B. B., Brown, K. S., Daniels, B. C., Myers, C. R., and Sethna, J. P. (2015). Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *The Journal of Chemical Physics*, 143(1):07B201_1.

Truccolo, W. (2016). From point process observations to collective neural dynamics: Nonlinear hawkes process glms, low-dimensional dynamics and coarse graining. *Journal of Physiology-Paris*, 110(4):336–347.

Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., and Brown, E. N. (2005). A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects. *Journal of Neurophysiology*, 93(2).

Van Kampen, N. G. (1992). *Stochastic processes in physics and chemistry*, volume 1. Elsevier.

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

Wilson, H. R. and Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13(2):55–80.

Wilson, H. R., Cowan, J. D., Baker, T., Cowan, J., and van Drongelen, W. (1972). Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons. *Biophysical Journal*, 12(1):1–24.

Zammit-Mangion, A., Dewar, M., Kadirkamanathan, V., and Sanguinetti, G. (2012). Point process modelling of the Afghan War Diary. *Proceedings of the National Academy of Sciences of the United States of America*, 109(31):12414–9.

Zammit-Mangion, A., Kadirkamanathan, V., and Sanguinetti, G. (2011). Efficient system identification of dynamic spatio-temporal systems from point process observations.

Zhang, H. and Xiao, P. (2018). Seizure dynamics of coupled oscillators with epileptor field model. *International Journal of Bifurcation and Chaos*, 28(03):1850041.

Zhao, Y. and Park, I. M. (2016). Variational latent gaussian process for recovering single-trial dynamics from population spike trains. *arXiv preprint arXiv:1604.03053*.

Zhao, Y. and Park, I. M. (2017). Recursive variational bayesian dual estimation for nonlinear dynamics and non-gaussian observations. *arXiv preprint arXiv:1707.09049*.

Zheng, J., Lee, S., and Zhou, Z. J. (2006). A transient network of intrinsically bursting starburst cells underlies the generation of retinal waves. *Nature Neuroscience*, 9(3):363.

Zhou, Z. J. and Zhao, D. (2000). Coordinated transitions in neurotransmitter systems for the initiation and propagation of spontaneous retinal waves. *Journal of Neuroscience*, 20(17):6570–6577.