

This is a pre-copyedited, author-produced version of an article accepted for publication in *Bioinformatics* following peer review. The version of record Sean Robinson, Michael J Courtney; Spatial quantification of the synaptic activity phenotype across large populations of neurons with Markov random fields, *Bioinformatics*, , bty322, is available online at <https://doi.org/10.1093/bioinformatics/bty322>

Bioinformatics

doi.10.1093/bioinformatics/xxxxx

Advance Access Publication Date: Day Month Year

Manuscript Category

OXFORD

Subject Section

Spatial quantification of the synaptic activity phenotype across large populations of neurons with Markov random fields

Sean Robinson^{1,2*}, and Michael J. Courtney^{3,4,5*}

¹Department of Mathematics and Statistics, University of Turku, Turku, Finland ²Université Grenoble Alpes, CEA, INSERM, Biology of Cancer and Infection UMR S 1036, F-38000 Grenoble, France ³Neuronal Signalling Lab, Turku Centre for Biotechnology, University of Turku, Turku, Finland ⁴Screening Unit, University of Turku, Turku, Finland ⁵Turku Brain and Mind Center, Turku, Finland

*To whom correspondence should be addressed.

Associate Editor: XXXXXXX

Received on XXXXX; revised on XXXXX; accepted on XXXXX

Abstract

Motivation: The collective and coordinated synaptic activity of large neuronal populations is relevant to neuronal development as well as a range of neurological diseases. Quantification of synaptically-mediated neuronal signalling permits further downstream analysis as well as potential application in target validation and *in vitro* screening assays. Our aim is to develop a phenotypic quantification for neuronal activity imaging data of large populations of neurons, in particular relating to the spatial component of the activity.

Results: We extend the use of Markov random field (MRF) models to achieve this aim. In particular, we consider Bayesian posterior densities of model parameters in Gaussian MRFs to directly model changes in calcium fluorescence intensity rather than using spike trains. The basis of our model is defining neuron 'neighbours' by the relative spatial positions of the neuronal somata as obtained from the image data whereas previously this has been limited to defining an artificial square grid across the field of view and spike binning. This both explicitly incorporates a spatial component within the model and allows for feasible inference. We demonstrate that our spatial phenotypic quantification is applicable for both *in vitro* and *in vivo* data consisting of thousands of neurons over hundreds of time points. We show how our approach provides insight beyond that attained by conventional spike counting and discuss how it could be used to facilitate screening assays for modifiers of disease-associated defects of communication between cells.

Availability: We supply the MATLAB code and data to obtain all of the results in the paper.

Contact: sean.j.robinson@utu.fi and michael.courtney@utu.fi

Supplementary information: Supplementary data are available at *Bioinformatics* online.

1 Introduction

Synaptic neurotransmission and neuronal activity are essential for the development and refinement of the mature nervous system [Blanquie *et al.*, 2017, Sigler *et al.*, 2017]. Perturbations in synaptic neurotransmission and plasticity are known contributors to diseases such as epilepsy, neuropathic pain and depression [Doucet *et al.*, 2012, Jacobs *et al.*, 2000, Zhuo, 2008]. Spontaneous neuronal activity has been found to play a particularly important role in the establishment of brain circuits during development [Kutsarova *et al.*, 2017, Pratt *et al.*, 2016, Leighton and Lohmann, 2016] and may also contribute to disorders such as phantom

limb pain [Blumberg and Dooley, 2017]. The use of *ex vivo* culture systems as models for neuronal development and disease provides numerous benefits such as the potential to assist in the discovery of underlying regulatory mechanisms. Neurons isolated from immature cerebral cortex are able to develop functional synaptic connectivity in culture. They exhibit spontaneous wave-like propagations of calcium spiking (Supplementary Movie 1) that have pharmacological sensitivity indicating the involvement of action potentials and neurotransmitter receptors [Dravid and Murray, 2004]. These properties suggest that such cultures may have the capacity to model aspects of development and disease related to synaptic function and could therefore be utilised in screens to identify molecules altering synaptic behaviour.

It is well understood that the contribution of neural activity to brain function lies within the coordinated activity of large populations of neurons rather than the independent activity of individual neurons [Tkačik *et al.*, 2014]. We therefore do not propose to analyse the behaviour of individual neurons but instead seek a representation of neuronal communication throughout larger populations. We aim to quantify this synaptically-driven activity, in particular relating to the spatial structure of the activity, as the phenotype of interest within a screen or experimental setup. This requires an aggregate analysis of up to thousands of neurons over hundreds of time points.

A general first step in the analysis of large, neuronal time-course data is dimensionality reduction applied either to the time series [Cunningham and Byron, 2014, Breakspear, 2017] or to low-level features such as cell morphology [Sharma *et al.*, 2012]. This approach has directly led to applications in decoding (classification), for example classifying images/videos of natural scenes using the corresponding induced activity of neurons in the retina [Ahmadian *et al.*, 2011, Onken *et al.*, 2016]. Other general analytical techniques include pairwise or population cross-correlation studies of the time series [Smith and Häusser, 2010, Okun *et al.*, 2015], or the use of hidden Markov models to infer underlying states of activity [Mazzucato *et al.*, 2016]. Non-spatial analysis of spike trains (binary time series of neuron spikes) makes use of ‘spike counting’ statistics such as the mean firing rate and interspike interval, or averaging over the data in some other way, for example by using time histograms [Hill *et al.*, 2015].

Specifically spatial analysis of neuronal data has typically addressed a small number of cells (~ 10) and data indices that need not meaningfully index space [?]. This is generally the case where Ising models are used for the analysis of spike train data [?]. Such analyses are computationally infeasible for larger numbers of neurons since the interactions within the model are generally defined between sets of all possible combinations of neurons [?]. For example, pairwise interactions are considered for every possible pair of neurons. In other words, all neurons are considered equally adjacent in space.

One possible way forward is to limit the interactions considered to just pairwise and additionally consider interactions based on the relative spatial location of the cells [?]. Although neurons are well known for their capacity to transmit a signal over large distances, there are also many local connections between neurons, which are particularly relevant during neuronal development [?]. Furthermore, pathological changes in local circuit properties are implicated in multiple neurodevelopmental and psychiatric disorders [?]. Computational limitations will necessarily result in compromises to either the number of neurons considered or the inclusion of long range connections. The local wave-like spread of spontaneous activity within a single field (Supplementary Movie 1), propagating at least $50\times$ slower than an action potential and involving trans-synaptic signalling (see below) prompted us to consider the applicability of an approach prioritising the size of the neuron population over long range connectivity.

Here we use a Gaussian Markov random field (MRF) model [?] with pairwise interactions and where neuronal connectivity is modelled based on the relative spatial locations of the neuron somata, resulting in a sparse adjacency matrix (Supplementary Figure 1). Our methodology is based on a novel use of Gaussian MRFs to directly model the process of interest as opposed to their more standard use as one component within hierarchical models for disease mapping [?]. We analyse the neuronal activity as extracted from image data (fluorescent dye-based measurements of calcium levels as opposed to binary spike trains) and quantify the spatial component of the activity through our model.

A particular exception to non-spatial analysis of neuronal activity data is the use of Markov random field (MRF) models [Makarenko *et al.*, 1997, François *et al.*, 2000, Abdallahi *et al.*, 2003]. However, up until

this point such analysis has been limited to spike train data which requires a spike binning scheme in time. Ising models are a specific type of MRF which have also been used in spike train analysis [Schneidman *et al.*, 2006, Shlens *et al.*, 2006, Nirenberg and Victor, 2007, Roudi *et al.*, 2009] but where pairwise interactions have been considered for every possible pair of neurons, or in other words, all neurons are considered equally adjacent in space. An explicitly spatial component can be achieved using an MRF model by defining a square grid across the field of view and where each region only interacts with its four immediate neighbours [Makarenko *et al.*, 1997, François *et al.*, 2000, Abdallahi *et al.*, 2003]. This sparse interaction setup means that much more data can be feasibly analysed than in the previous use of Ising models (only ~ 10 cells), however the definition of such a grid requires a spike binning scheme in space as well.

Here we use a Gaussian MRF model [Rue and Held, 2005] with pairwise interactions and where neuronal connectivity is modelled based on the relative spatial locations of the neuron somata. Hence we do not consider a square grid but rather a triangulation actually corresponding to individual neurons yet which also results in a sparse adjacency matrix (Supplementary Figure 1). We directly analyse the neuronal activity as extracted from image data, that is, the fluorescent dye-based measurements of calcium levels as opposed to spike trains. Furthermore, rather than point estimation of model parameters as has been previously achieved, we will consider posterior densities of the relevant model parameters in a fully Bayesian framework. Hence we quantify the spatial component of the neuronal activity through our model, which is based on a novel use of Gaussian MRFs to directly model the process of interest.

Recent advances in high-resolution imaging of large populations of neurons across multiple brain areas [Sofroniew *et al.*, 2016] as well as potential large-field *in vitro* screening applications (~ 5000 individual cells; ~ 6000 time points) has provided the opportunity to specifically leverage the spatial component of synaptic activity data. We consider data from simulation experiments, data acquired from cultures of *ex vivo* neurons specifically for this study and freely available data from *in vivo* studies published in the literature [Li *et al.*, 2015b, Chen *et al.*, 2016, Sofroniew *et al.*, 2016]. We show that our approach can achieve insights beyond conventional spike counting, as well as conforming to and complementing previously attained results. We also discuss how our approach could be used to facilitate screening assays. [through the analysis of the synaptic activity of large populations of neurons.](#)

2 Materials and Methods

2.1 Neuronal culture preparation

Rat cortical neuron cultures were prepared as previously described [Li *et al.*, 2013], in NeurobasalA medium supplemented with 1.5 to 2% B27 (Life Technologies-Thermo). Half of the medium was replaced with fresh medium every three days. Cells were plated in Cellstar 96 well microclear plates (Greiner, 0.19mm bottom) for large-field imaging or Cellstar 96 well plates (Greiner, 1mm bottom) for medium-field imaging. At 9 (large-field experiment) or 16 (medium-field experiment) days *in vitro*, the medium on the cortical cultures was replaced with wash-free calcium reporter dye loading medium (with final concentrations of $2\ \mu\text{M}$ Fluo4AM, 0.01% pluronic acid, 0.1% BSA and Hoechst 33342 DNA dye in SGF-based solution (NaCl 137.5 mM, KCl 5.3 mM, CaCl_2 2 mM, MgCl_2 1 mM, HEPES 10mM, Glucose 30 mM, Glycine 1 mM; [Bading *et al.*, 1993]) supplemented with 0.5 mM sodium pyruvate, 2.5mM additional NaHCO_3 and 10% minimal essential medium (MEM Gibco-Thermo cat #11700077 for imaging under ambient conditions) and incubated for 1 hour to allow uptake of the dye.

2.2 Image data acquisition

For large-field imaging, images were taken through a $4\times$ objective (NA 0.16, Olympus) using a customized IX70 microscope as previously described [Li *et al.*, 2015a]. The field was illuminated with a 460nm LED source (XLED1, LDGI). Emission light was collected through 510nm long-pass filter. The Hamamatsu Flash 4.0 camera was run using HCImage software (Hamamatsu) with binning 2×2 to generate 1024×1024 pixel images at 16 bit depth with an integration time of 0.300 seconds. The field (size $3.2\text{ mm} \times 3.2\text{ mm}$) contained up to 10^4 cells and was imaged every ~ 0.375 seconds for a total of 418 frames (total time ~ 156.75 seconds). The data corresponding to this setup is hereafter referred to as the large-field (LF) data.

To acquire images of cells before, during and after addition of pharmacological inhibitors, an automated imager equipped with a cell-imaging chamber (set at 5% CO_2 , 37°C) containing a pipetting gantry was used (BD Pathway 855 High Content Analyzer). Cell culture medium was first manually removed and replaced with MEM ($64\ \mu\text{l}/\text{well}$) supplemented with $5\ \mu\text{M}$ final picrotoxin, an inhibitor of the GABA-A receptor, to promote a reproducible level of spontaneous activity in the cultures. Each of 30 cultures in wells was sequentially loaded (at 45 second intervals) with calcium dye by on-stage addition of $6\ \mu\text{l}$ of $10\times$ wash-free calcium reporter dye loading medium in MEM, using the single-channel pipettor of the automated imager.

Cells were treated during the third imaging cycle with either $0.1\ \mu\text{M}$ tetrodotoxin (Ascent Scientific) to fully block action potentials, or MK-801 (Tocris) $0.2\ \mu\text{M}$ to partially reduce NMDA receptor-gated calcium responses, in 10 replicate wells for each condition. Additions were made from $10\times$ stocks in MEM using the on stage-pipettor. Control samples were left untouched after the loading dye addition. This system uses a less sensitive camera (Hamamatsu Orca-ER CCD, 1344×1024 pixels, used at binning of 2×2), necessitating the use of higher magnification and higher numerical aperture objectives, leading to reduced number of cells per field (~ 300). Images were acquired through a long-working distance $10\times$ objective (NA 0.3, Olympus) for 45 frames per acquisition window (~ 35 seconds within a single acquisition window, allowing ~ 10 seconds for the objective to move to the next well and refocus, resulting in ~ 20 minutes between acquisition cycles). There were 2 windows of measurements pre-treatment, 1 during treatment and 4 windows post-treatment. The data corresponding to this setup is hereafter referred to as the medium-field (MF) data.

2.3 Voronoi tessellations and triangulations

A tessellation is defined as a partition of 2D space into a set of regions called tiles. When the set of tiles is in one-to-one correspondence with a set of seed points in space such that any point lying within a tile is closest to that tile's corresponding seed point, this is known as a Voronoi tessellation. We can then construct a corresponding graph by considering the seed points as vertices and that there is an edge between any two vertices that have adjacent corresponding tiles within the field of view. Hence we end up with a triangulation, although note that this is not a Delaunay triangulation since the adjacency of the Voronoi tiles is constrained to exist only within a specified field of view.

An example Voronoi tessellation obtained from neuronal culture image data is given in Figure 1. We have the standard deviation projection of the time series image data (Figure 1(a)), which is then used for the segmentation mask for the neuronal cells (Figure 1(b)). The centres of the segmented regions are the seed points used to define the Voronoi tessellation using Euclidean distance (Figure 1(c)). Finally, the triangulation is found resulting in a graph where each vertex corresponds to a neuronal cell and each edge corresponds to neuronal cell 'neighbours' as seen in the Voronoi tessellation (Figure 1(d)). Hence neurons are only

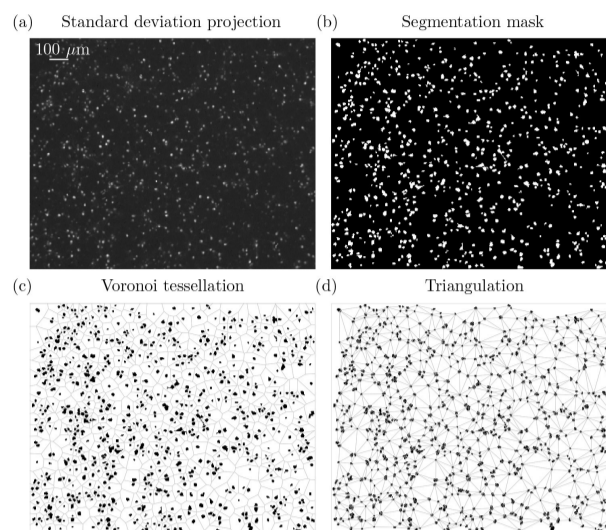


Fig. 1. Obtaining the Voronoi tessellation and graph. (a) The standard deviation projection (in time) for an example field of view of cultured neurons. (b) The segmentation mask of the neuronal cells. (c) The Voronoi tessellation where the seed points are the centres of each of the segmented cells. (d) The corresponding triangulation graph. That is, neurons represented as vertices have neighbours represented through edges corresponding to adjacent Voronoi tiles, not all other neurons (see Supplementary Figure 1).

'neighbours' corresponding to adjacent Voronoi tiles, not with all other neurons (see Supplementary Figure 1).

Supplementary Figure 2(a) shows an example of the same fluorescence image data, F_t presented as a 'film strip'. The first difference images, $\Delta F_t = F_t - F_{t-1}$ are also given as a 'film strip' (Supplementary Figure 2(b)), which are then overlaid with the corresponding segmentation mask and Voronoi tessellation (Supplementary Figure 2(c)). Then the extracted neuron activity time series data, the median fluorescence intensity value of each cell at each time point, is visualised on the Voronoi tessellation by colouring each tile based on the value of the time series (Supplementary Figure 2(d)). The same data is presented in Supplementary Movie 2. Note that the sizes of the Voronoi tiles do not have any meaning in our analysis but the tessellation is a useful visualisation of the time series data that maintains the spatial adjacency information.

2.4 Spatial model for neuronal activity

Consider that we have a collection of neurons with an associated graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ and that X_i for $i \in \mathcal{V}$ are the random variables of the measured neuron activity. For convenience in the following we will consider that the index i corresponds to individual neuronal cells while the considerations for time are addressed below. We write $X = \{X_i \mid i \in \mathcal{V}\}$ (all the measurements) and $X_{\mathcal{V}\setminus i} = \{X_j \mid j \in \mathcal{V}\setminus i\}$ (all the measurements except the i th). Let the edge set $\mathcal{E} = \{e_{ij} \in \{0, 1\} \mid e_{ij} = e_{ji}, \text{ for all } i, j \in \mathcal{V}\}$ be the set of all edges between every pair of neurons where neurons i and j are neighbours if and only if $e_{ij} = 1$. Let the degree of vertex i be $\partial_i = \sum_{j \in \mathcal{V}} e_{ij}$ with $e_{ii} = 0$ for all $i \in \mathcal{V}$.

Conditional autoregressive (CAR) models are Gaussian MRFs defined by the series of full conditional distributions

$$X_i | X_{\mathcal{V}\setminus i} \sim N\left(\phi \sum_{j \in \mathcal{V}\setminus i} \frac{e_{ij}}{\partial_i} X_j, \frac{\sigma^2}{\partial_i}\right)$$

for all $i \in \mathcal{V}$. That is, each variable X_i has a conditional Gaussian distribution with mean, the ϕ -scaled average of its neighbouring values and variance σ^2 scaled by the number of its neighbours. The joint distribution

of X exists and is unique [Rue and Held, 2005], written

$$X \sim N(0, \sigma^2(D(I - \phi W))^{-1})$$

where D is a diagonal matrix with entries $D_{ii} = \partial_i$ for all $i \in \mathcal{V}$, I is the identity matrix and W is known as the adjacency matrix with entries $W_{ij} = e_{ij}$ for all $i, j \in \mathcal{V}$. For the covariance matrix to be positive definite, $\phi \in (-1, 1)$ and $\sigma^2 > 0$ [Gelfand and Vounatsou, 2003]. If $\phi = 0$ then X has a diagonal covariance matrix and hence each of the univariate marginal distributions are independent. A value of $\phi > 0$ results in spatial autocorrelation between neighbouring variables (easiest to see in the full conditional distributions), while $\phi < 0$ results in spatial anti-autocorrelation. Hence ϕ is known as the spatial autocorrelation parameter.

Likelihood based methods [Cressie *et al.*, 2005] as well as approximate algorithmic methods [Banerjee *et al.*, 2014] exist for inference of the spatial autocorrelation parameter ϕ and the precision $\tau = \frac{1}{\sigma^2}$. Although CAR models are often utilised as one component within larger hierarchical models [Banerjee *et al.*, 2014], unlike disease mapping with covariates, there is no need here and we can find an explicit expression for the posterior density of the spatial autocorrelation parameter ϕ within a Bayesian framework [Bell and Broemeling, 2000, De Oliveira, 2012, Ren and Sun, 2013]. We set priors $\pi(\phi) \propto \text{constant}$ and $\pi(\tau) \propto \tau^{-1}$, and we assume *a priori* that ϕ and τ are independent, that is $\pi(\phi, \tau) = \pi(\phi)\pi(\tau) = \tau^{-1}$.

Now the likelihood is

$$\pi(X = x|\phi, \tau) \propto |D(I - \phi W)|^{\frac{1}{2}} \tau^{\frac{n}{2}} \exp\left\{-\frac{\tau}{2} X^T (D(I - \phi W)) X\right\}.$$

Hence the joint posterior distribution is

$$\begin{aligned} \pi(\phi, \tau|X = x) &\propto \pi(X = x|\phi, \tau)\pi(\phi, \tau) \\ &= |D(I - \phi W)|^{\frac{1}{2}} \tau^{\frac{n}{2} - 1} \exp\left\{-\frac{\tau}{2} X^T (D(I - \phi W)) X\right\}. \end{aligned}$$

For fixed ϕ it can be seen that τ has a Gamma distribution,

$$\tau|\{\phi, X = x\} \sim \Gamma\left(\frac{n}{2}, \frac{1}{2} X^T (D(I - \phi W)) X\right).$$

In order to find the posterior density of ϕ we integrate $\pi(\phi, \tau|X = x)$ with respect to τ by using the properties of the Gamma distribution to obtain

$$\pi(\phi|X = x) \propto |D(I - \phi W)|^{\frac{1}{2}} \left(X^T (D(I - \phi W)) X\right)^{-\frac{n}{2}}.$$

Hence we are able to sample from $\pi(\phi|X = x)$. That is, for a given set of measurements of neuron activity $X = x$, we are able to sample from the posterior distribution of ϕ under a CAR model in order to assess the amount of spatial autocorrelation present in the data as measured by ϕ .

In the above presentation the index of the random variables X does not explicitly correspond to either space or time. The spatial adjacency matrix can correspond to any graph, for example the one given in Figure 1(d). In order to model through time, identical graphs can be considered for each time point and extra edges added between those graphs to represent temporal adjacency between each neuron in time. Then the autocorrelation parameter can be decomposed into a spatial and temporal component, $\phi W = \phi_{\text{space}} W_{\text{space}} + \phi_{\text{time}} W_{\text{time}}$ where W_{space} and W_{time} are the adjacency matrices corresponding to spatial and temporal adjacency respectively. Although the adjacency matrix W can take any form in principle, in practice W is generally a sparse matrix, which is what allows for the model inference to be feasibly achieved. Sparseness results automatically from having the adjacency matrix correspond to adjacent tiles in a Voronoi tessellation in our model.

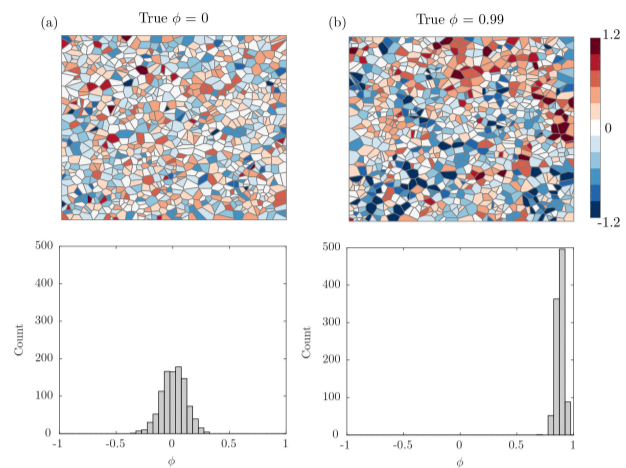


Fig. 2. Simulations from a MRF based on the Voronoi tessellation of 1000 uniformly distributed data points in space. The Voronoi tiles are coloured in an analogous way as in Supplementary Figure 2(d). (a) True autocorrelation parameter $\phi = 0$ and histogram of a sample from the posterior density given below. (b) True autocorrelation parameter $\phi = 0.99$ and histogram of a sample from the posterior density given below.

3 Results and Discussion

3.1 The MRF model obtains accurate estimates of the autocorrelation parameter ϕ in simulated experiments

We first consider a simulation from an MRF model with known parameters. One thousand points were uniformly distributed in the field of view and the corresponding Voronoi tessellation and graph was obtained. Figure 2 shows realisations of the MRF model for both high and low values of the spatial autocorrelation parameter ϕ (note we are just considering a single time point so there is no need to consider ϕ_{time} and we set $\tau = \sigma^2 = 1$ for all simulations). Below the simulated data in Figure 2 is the histogram of a sample from the corresponding posterior density of ϕ . We can see that when the true value of $\phi = 0$ (no spatial autocorrelation between adjacent cells), the posterior density is centred around 0 and when the true value of $\phi = 0.99$ (high spatial autocorrelation between adjacent cells), the posterior density is centred around 0.9.

Supplementary Figure 3 shows boxplots of samples from the posterior density of ϕ for a range of true values of ϕ for each of 10, 100, 1000 and 10^4 simulated points in the field of view. We can see that for the most part the posterior samples are centred around the true corresponding values of ϕ , with greater accuracy and smaller variance as the number of data points increases as expected. This plot suggests a lower limit of around 1000 data points needed to obtain reasonable posterior distributions. When the true value of $\phi = 0.99$, the centres of the posterior samples for 1000 and 10^4 simulated points appear to be biased, sitting around 0.9. However, in this case we are right at the edge of the domain of ϕ and so it is not completely surprising that the centres of the posterior distributions are further away from that edge.

The above simulations correspond to a single time point whereas in practice we have time series data. As discussed above, incorporating a time component into the model can be achieved by setting $\phi W = \phi_{\text{space}} W_{\text{space}} + \phi_{\text{time}} W_{\text{time}}$. Supplementary Figure 4 shows boxplots of samples from the posterior densities of ϕ_{space} for a range of true values of both ϕ_{space} and ϕ_{time} . However in each case the posterior densities of ϕ_{space} were found when holding $\phi_{\text{time}} = 0$ fixed. That is, ϕ_{time} appears to be ancillary for the accurate inference of ϕ_{space} , at least in terms of the centre of the posterior sample. Hence in all cases below we set $\phi_{\text{time}} = 0$ and write ϕ instead of ϕ_{space} . Since we are most interested in the spatial autocorrelation of populations of neuronal cells, this makes the analysis

much clearer and also means there is much less computational cost to finding the posterior density.

In the case where we have less than 1000 cells in the field of view (for example with the MF data) we may expect low accuracy and high variance in our samples of the posterior density based on Supplementary Figure 3. However, Supplementary Figure 4 also demonstrates that in cases where we have time series data, the equivalent number of data points used to calculate the posterior density is the number of cells in the field of view multiplied by the number of time points. Hence even though there are only ~ 200 cells per field of view for the MF data, we have 45 frames in each acquisition window and hence the equivalent number of data points is $200 \times 45 = 9000$. Hence we are more generally able to expect posterior samples with both low variance and high accuracy.

3.2 The MRF model can distinguish between purely spatial differences in wave-like neuronal activity

We now consider a simulation experiment based on the LF data. This data itself was generated as a test of the system and so only a single field of view is considered with untreated cells exhibiting spontaneous network activity and wave-like signal propagation (Supplementary Movie 1). The propagation speed of the wave (< 8 mm/s) is well below the reported action potential propagation speeds in cortical and hippocampal neurons (~ 240 - 470 mm/s [Gulledge and Stuart, 2003, Kress *et al.*, 2008, Schmidt-Hieber *et al.*, 2008]). This, together with the elimination of spontaneous activity by tetrodotoxin (see below; [Draavid and Murray, 2004] strongly argue that the observed wave-like spread results from synaptic activation across multiple local synapses in series, rather than the arrival of a single long range input synapsing on multiple cells. In the latter case the maximal delay in activation across the imaging field would be 10 ms, not 300 ms as observed here.

There are a total of 5366 segmented cells in the field of view and the data consists of 418 time frames (417 first difference frames). Supplementary Figure 5 shows the degree distribution of the vertices in the corresponding graph. We consider 2 different simulation schemes using the LF data by permuting the time series either in space or in time. In the ‘permute space’ case, the observed time series are randomly permuted between the cells. In the ‘permute time’ case, the time series are randomly shifted in time for each cell (and wrapped around as necessary). The original data and the two simulated cases can be seen in Supplementary Movie 3 while the corresponding time series is given in Supplementary Figure 6.

Boxplots of the samples from the posterior densities of ϕ are given for different time windows in Figure 3. We can see the spatial component of the data is clearly apparent in the posterior samples of ϕ . That is, the wave-like propagation of activity that is visible in the original data and is no longer visible after spatial permutation (Supplementary Movie 3). Note that analyses such as spike counting would give the same output for both the original data and the spatially permuted case as the spatial component of the data is simply not taken into account. The spatial autocorrelation is zero after temporal permutation since the spiking has been completely desynchronised. If the time series were spike trains, statistics like the mean firing rate would not change under either spatial or temporal permutation.

3.3 The MRF model can quantify suppression of cellular communication induced by pharmacological compounds

Our aim is to develop a phenotypic quantification of neuronal communication that could be used for cell-based screens of reagents with potential pharmacological activity. We therefore evaluated the use of the spatial autocorrelation parameter to quantify the effects of specific compounds on cellular communication. In the MF data, a control condition

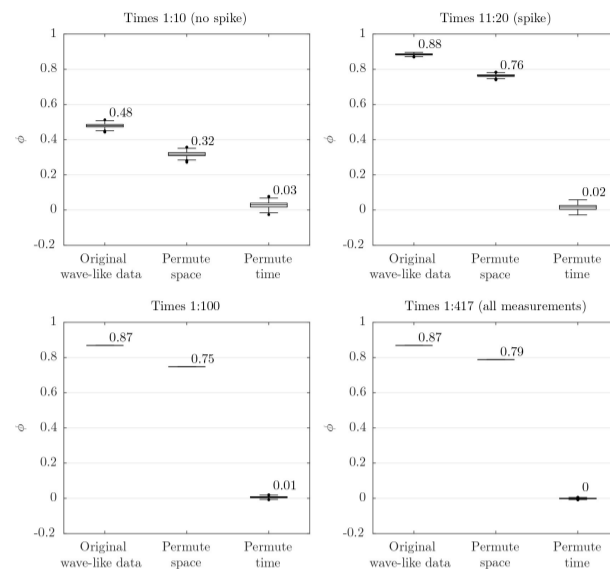


Fig. 3. Boxplots of samples from the posterior density of ϕ for the simulation scheme based on the LF data for different time windows with the median value also provided. When considering 10 time points the equivalent number of data points is $\sim 5 \times 10^4$ while when considering all time points the equivalent number of data points $\sim 2 \times 10^6$ and hence the sample variance is so low (compare to Supplementary Figure 3). The corresponding neuron activity data is presented in Supplementary Movie 3.

was compared with 2 pharmacological treatments, tetrodotoxin (TTX) at $0.1 \mu\text{M}$ to abolish action potentials, and MK-801 at $0.2 \mu\text{M}$ to partially limit calcium flux through NMDA receptor channels, each of which was replicated in 10 wells. The distribution of vertex degrees (Supplementary Figure 7) is very similar for each well with an average degree ~ 5 and low spread (comparison with Supplementary Figure 5 demonstrates the invariance of the size of the cellular neighbourhood when defined by a Voronoi tessellation for different field sizes, suggesting the scalability of our approach). The number of cells per well are also quite similar with a total median of just under 200. There are 45 time frames per acquisition window so for each well and each window we sample from the posterior of ϕ where there are the equivalent of ~ 9000 data points.

For each well we consider the 2 pre-treatment and 4 post-treatment acquisition windows. Figure 4 shows boxplots of the spatial autocorrelation parameter where for each treatment and for each acquisition window there is a boxplot for each of the 10 replicates. Example time series for 3 wells are also given in Supplementary Figure 8. Recall that cells were pre-treated with a low concentration of picrotoxin, an inhibitor of the GABA-A receptor, to initially promote spontaneous activity of the networks. Therefore, in all pre-treatment imaging time-windows, neurons are spiking in a highly synchronised way ($\phi \approx 0.99$).

Addition of TTX leads to a dramatic suppression of calcium spikes as the treatment effectively eliminates the action potentials on which intraneuronal communication depends (Figure 4). Thus the cells are not expected to be exposed to synaptically released transmitter and network activity is terminated. In the untreated condition used as a control, we can see that the cells appear to gradually become less synchronised (Figure 4) and exhibit increasing firing rates (Supplementary Figure 8). This likely reflects adaptation of the network to the loss of inhibitory input resulting from continued exposure of the network to picrotoxin, by synaptic scaling or other homeostatic mechanism [Fernandes and Carvalho, 2016].

The NMDA receptor open-channel blocker MK-801 at $2 \mu\text{M}$ prevents NMDA-evoked calcium responses [Li *et al.*, 2013, Courtney *et al.*, 1990] and, in pilot experiments, eliminates all activity (data not shown) not unlike

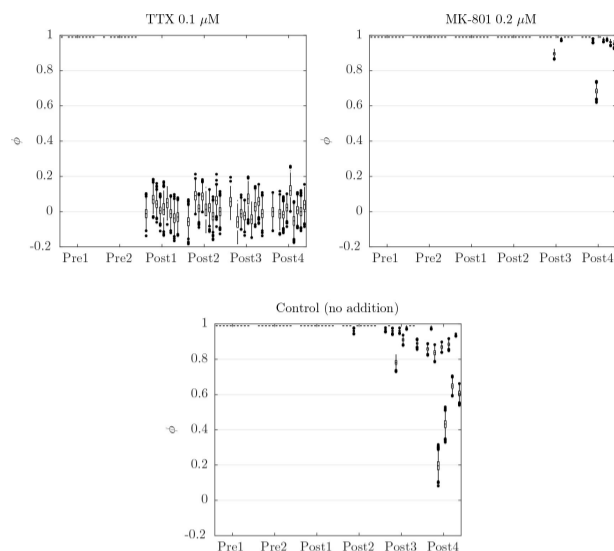


Fig. 4. Boxplots of samples from the posterior density of ϕ for the MF data. There were 6 acquisition windows (2 pre-treatment, 4 post-treatment) with 10 replicates for each condition.

treatment with TTX $0.1 \mu\text{M}$. For this reason, we used $0.2 \mu\text{M}$ MK-801 to only partially reduce the NMDA-evoked calcium responses. This condition only leads to a slight and very delayed drop in autocorrelation, compared with the TTX addition. This is only seen during the final acquisition window (Figure 4 and Supplementary Figure 8). This behaviour suggests that the rate of adaptation of the network activity to picrotoxin, measured as autocorrelation of calcium responses, itself depends on the calcium influx through the NMDA receptors that leads to the autocorrelated calcium activity. This is consistent with the adaptation of the network to the source of the change of activity and with the known role of NMDA receptor activity to synaptic scaling events occurring in the hours time range [Fernandes and Carvalho, 2016, Pawlak *et al.*, 2005].

It should be emphasised that the control sample did not receive a blank addition but was untouched during the measurement period. That is, even though the apparent adaptation to picrotoxin (added 1 hour before the start of the first time-window) becomes visible, **no treatment was applied. addition was actually made.** Furthermore it is interesting to note that the behaviour of the control MF data converges over time towards that of the LF data (see the time series in Supplementary Figures 6 and 8 and the spatial autocorrelation in Figures 3 and 4), which expressed baseline spontaneous activity and received no picrotoxin supplementation. This supports our proposal above that the low level picrotoxin-treated control samples gradually adapt and start to recover a baseline state. It is clear from the raw traces (Supplementary Figure 8) that not only do the calcium spikes gradually disappear but a proportion of the cells also acquire a heterogenous calcium state distinct from their initial inter-spike levels ('Pre1' panel) and inhibited levels in either TTX or MK treated cells.

We investigated potential phototoxicity in the MF data since illumination was greater and over a longer time period in this experiment compared with the LF setup. Toxicity was evaluated by lysis of cells during the experiment which was measured as a notably reduced fluorescence signal, that is if the cell exhibited a maximum intensity in the last imaging window less than the minimum intensity in the first imaging window acquired ~ 2 hours earlier. Toxicity might have resulted from the loading procedure, the de-esterification of the acetoxymethyl-derivitised dye, the inadequate humidification in the imaging chamber, phototoxicity, continuous exposure to $5 \mu\text{M}$ picrotoxin used to promote synaptic activity, or a combination of these factors. In addition to these non-optimal

conditions common to all samples, some samples were in addition exposed to inhibitors part way through the imaging. In spite of this, toxicity as determined above was less than 1% during the 2 hour period for most wells (18/20) in the control and MK-801 treated conditions. Hence overall toxicity was minimal, suggesting phototoxicity was also not a major problem during the course of this experiment. Interestingly, wells treated with sufficient TTX to prevent calcium spikes showed greater average cell loss (mean $\sim 4\%$; maximum of $\sim 8\%$), consistent with calcium spiking rather than illumination being the main source of toxicity under our conditions.

The data here was acquired as a proof of concept to evaluate the application of the spatial autocorrelation parameter in distinguishing varying neuronal activity induced by standard pharmacological manipulations. In a screening scenario, thousands of data points per sample would be necessary but this can be achieved from image data acquired over multiple time points. Such a system has the potential to be applied to cell-based models of disease and to identify compounds correcting disease-associated defects of cell-to-cell communication.

3.4 The MRF model can be used for the analysis of *in vivo* data and generates conclusions consistent with previous analyses

We aimed to additionally evaluate the application of the MRF approach to *in vivo* calcium imaging time course data. We considered image data from two different *in vivo* experiments using mice expressing fluorescent calcium sensor proteins in their neurons. In the first instance we consider the 2-photon random access mesoscope (2pRAM) data, which consists of high-resolution images of the activity of populations of neurons simultaneously acquired from multiple brain areas [Sofroniew *et al.*, 2016]. The authors of the original study selected 4 different regions within a larger brain-wide image, with ~ 100 neurons per region and 6000 time frames (their Figure 8). As an example of downstream analysis, the Pearson correlation of the 'average' activity (computed by taking the maximum values across the time series) in region 1 (corresponding to the somatosensory cortex) against individual cells in the other regions was analysed [Sofroniew *et al.*, 2016]. As a result, it was shown that the activity of individual neurons in one part of the brain could be compared to the

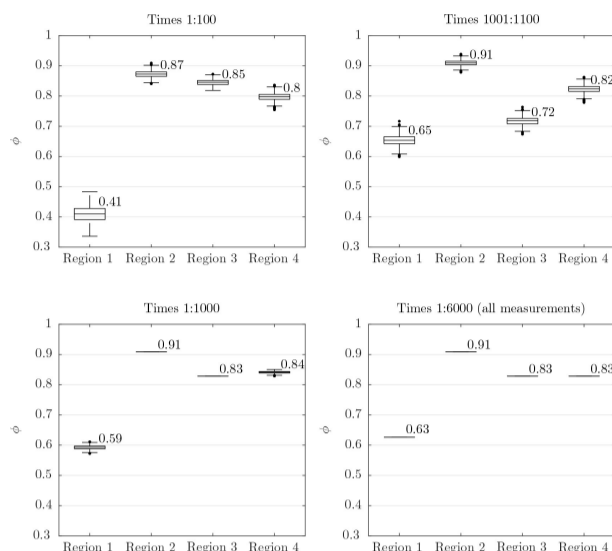


Fig. 5. Boxplots of samples from the posterior density of ϕ for the 2pRAM data [Sofroniew *et al.*, 2016] (the 4 regions can be seen in their Figure 8) for different time windows with the median value also provided. When considering all time points the equivalent number of data points is $\sim 6 \times 10^5$.

activity in other regions that had been simultaneously imaged [Sofroniew *et al.*, 2016].

We consider the same 4 regions of the 2pRAM data and investigate the spatial autocorrelation parameter. In Figure 5 and we can see that there is different spatial autocorrelation observed in each of the 4 regions within different time windows. Interestingly, regions 1 (somatosensory cortex) and 3 (retrosplenial cortex) exhibit quite different spatial autocorrelation across the different time windows while regions 2 and 4, which both appear to be in the primary somatomotor cortex, have more consistent spatial autocorrelation. In all cases, as more and more time frames are considered the autocorrelation parameter appears to converge and stabilise. Once again, there is little variance in the samples since we have the equivalent of up to $\sim 6 \times 10^5$ data points (although this is a similar order as the LF data, note that there are fewer cells and many more time points here).

These observed differences between regions were not identified by the use of the Pearson correlation as originally described [Sofroniew *et al.*, 2016]. Supplementary Figure 9 gives boxplots of the Pearson correlation calculated between every pair of neurons in each region for the same time windows as Figure 5. The Pearson correlation is centred close to zero in all cases and as with the original output [Sofroniew *et al.*, 2016], it is difficult to interpret. Our spatial autocorrelation parameter gives a model-based quantification of the neuronal activity whereas with a more ad hoc approach there are abundant possibilities with no clearly most appropriate choice. For example, the Pearson correlation could be calculated between every pair of neurons, or between each neuron and an average [Okun *et al.*, 2015] or maximum [Sofroniew *et al.*, 2016] time series, not to mention considering cross-correlation and lag.

Secondly, we considered image data taken from the left anterior motor cortex (ALM) in mice performing a directed licking task for a study aimed at localising the planning of motor activity [Li *et al.*, 2015b]. Data are available from four different mice licking either to the left or right in response to a tactile cue [Chen *et al.*, 2016]. Each imaging session corresponds to a single mouse and defined depth within the left ALM. There were 59 sessions in total, of which 53 were used (where the time series didn't contain either Inf or NaN). There were 20 sessions available for mouse A19 and 11 sessions available for each other mouse. Multiple

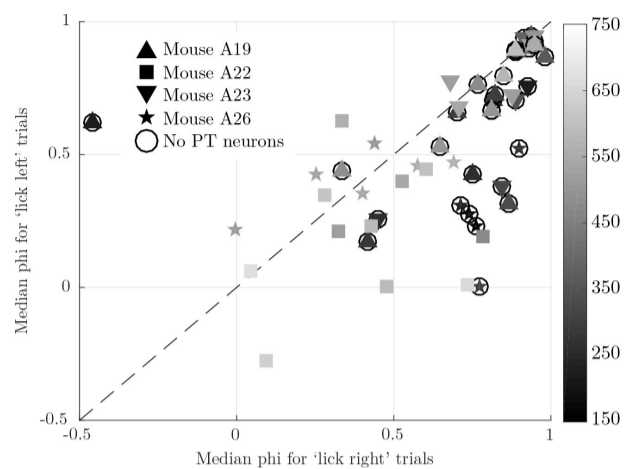


Fig. 6. Median values of samples from the posterior density of ϕ for each session (mice A19, A22, A23 and A26 at different imaging depths) in the ALM data [Li *et al.*, 2015b, Chen *et al.*, 2016]. The depth of the imaging is given in the colourbar (μm). There were no labelled pyramidal tract (PT) neurons for sessions with depths $< 450 \mu\text{m}$ (above layer 5) as well as for some sessions at greater depths.

replicates of both licking tasks were measured in a single session. The median number of cells in a field of view is 88 and with 91 time points for each session, the equivalent number of data points is ~ 8000 .

We calculate the spatial autocorrelation of the neuronal activity observed within the imaging field for each session for the left and right licking tasks separately and plot the median values in Figure 6. For the most part, we found that the neuron activity (which was only recorded in the left ALM) appears to be more spatially autocorrelated for the 'lick right' cases compared to the 'lick left'. This is consistent with the original report on this data that the imaged neurons (from the left ALM) responded more during right lick trials [Li *et al.*, 2015b]. Notably, this results was not originally found from spike counts across the recorded population of neurons but from additional optogenetic experimentation [Li *et al.*, 2015b], suggesting that the explicit spatial component of our approach may provide a significant benefit for the analysis of activity-behaviour correlation *in vivo*.

The data in Figure 6 are also coloured by depth of measurement highlighting that the largest differences in ϕ mostly correspond to imaging sessions at lower depths. Imaging sessions without pyramidal tract (PT) neurons are also identified and appear to similarly correspond to this difference. Note that there are no PT neurons for sessions with depth $< 450 \mu\text{m}$ (that is, above layer 5) but there are some deeper sessions also lacking identified PT neurons. Also note that in our analysis all neurons were considered regardless of PT identification. Once again, we also consider the Pearson correlation between every pair of neurons in each session and can see that the median Pearson correlation is generally low (Supplementary Figure 10). Although it is arguable that there is higher Pearson correlation in the 'lick right' trials, this result is much clearer and more interpretable with our model-based spatial autocorrelation parameter in Figure 6.

3.5 Model fitting discussion

Kinetic large-field imaging of neural tissue and *in vivo* samples generally doesn't allow for the complete capture of all neurites, which can include structures smaller than the resolving power of the microscopy methods typically employed. Using a Voronoi tessellation as the basis for neuronal connectivity in our model clearly does not represent the complete physical and chemical interactions between neurons. Nevertheless it allows us

to feasibly describe collective behaviour such as the motivating wave-like synaptic activity (Supplementary Movie 1), without exhaustively considering minute details, some of which may not even have been captured in the original image data. The resultant model adjacency matrix is sparse which is what allows for feasible inference. Additionally, using a Voronoi tessellation results in graphs with very similar vertex degree distributions regardless of the total number of neurons (Supplementary Figures 5 and 7) indicated our method is inherently scalable for different field sizes.

Within our setup the maximum clique size is 3 (triangulation), however cliques of size 2 already create spatial structure and those of larger sizes are rarely used, as they are considered to introduce complexity with little benefit [Banerjee *et al.*, 2014]. Hence we only considered pairs of neurons (the cliques of size 2), which have been more generally preferred [Tkačik *et al.*, 2014, Nirenberg and Victor, 2007]. The analysis could easily be extended to three-dimensional (3D) image data as there is a 3D analogue to a Voronoi tessellation. The input to our approach is also not necessarily image data but could in principle be spatially reconstructed neuronal activity from electrophysiological data acquired with large-scale multielectrode arrays for example [Hilgen *et al.*, 2017]. **Other opportunities for data integration include how the spatial autocorrelation parameter relates to concurrent measurements, for example fMRI BOLD signals.**

We considered pairwise interactions in our model based on the relative spatial location of the neuronal somata. Although neurons are well known for their capacity to transmit a signal over large distances, there are also many local connections between neurons, which are particularly relevant during neuronal development [Luhmann *et al.*, 2016]. Furthermore, pathological changes in local circuit properties are implicated in multiple neurodevelopmental and psychiatric disorders [Tatti *et al.*, 2017]. The local wave-like spread of spontaneous activity within a single field (Supplementary Movie 1), propagating at least $50\times$ slower than an action potential and involving trans-synaptic signalling (Section 3.2) prompted us to consider the applicability of an approach prioritising the size of the neuron population over long range connectivity.

In all of our applications we considered that the spatial autocorrelation parameter ϕ was constant in space. This corresponds to an assumption that this component of the neuron activity is the same across the entire field of view. This assumption seems plausible for culture since in principle there should be no distinct regions within the field of view nor orientation to the culture. Generally this assumption will not be realistic for *in vivo* or even *ex vivo* tissue imaging data. However, when the field of view is small enough, as in the case of the ALM data (~ 88 cells), or the 4 regions of the 2pRAM data (~ 100 cells), this still may be a reasonable assumption. For larger field image data our method could be applied by windowing across the field of view to obtain a series of local values of ϕ , which was essentially the case with the 4 regions from the larger 2pRAM image data.

An additional assumption in the model is that the spatial autocorrelation parameter was constant in time, that is, it itself does not change in time. The assumption of constant ϕ in time fitted nicely with the experimental setup of the MF data as any change in ϕ is considered between acquisition windows rather than within any one window. This is also not particularly problematic for the ALM data since interest is in the comparison across sessions with the same length time frames. More generally, the assumption of constant ϕ in time could also be considered within the framework of windowing statistics for time series analysis, as was presented for the LF and 2pRAM data.

Another aspect relating to the model parameters is that we did not consider ϕ_{time} nor the precision parameter τ . The main issue here is computational since setting $\phi_{\text{time}} = 0$ results in a (sparse) block diagonal adjacency matrix W and so adding 100s of time points makes little difference to the computation time. If we consider $\phi_{\text{time}} \neq 0$ then the adjacency matrix will be block tridiagonal and calculating the determinant

becomes computationally infeasible with a large number of time points. However we showed that we can still obtain accurate estimates for ϕ_{space} by simply holding $\phi_{\text{time}} = 0$ fixed. It was also the case that not considering the precision parameter τ simplifies the analysis and the interpretation of ϕ since the posterior density of τ is conditional on ϕ anyway. Investigation of both ϕ_{time} and τ could be considered in further work.

3.6 Implementation in MATLAB and computational expense

The MRF model was implemented in MATLAB on a mid 2012 MacBook Pro with a 2.6 GHz Intel Core i7 processor (quad-core) and 16 GB of RAM. The LF data was the most computationally expensive with 5366 identified cells within a single field of view. In this case it took ~ 15 seconds to calculate the Voronoi tessellation and ~ 60 seconds to calculate the posterior density of ϕ (using all 417 time frames). Calculation of the posterior density for the 2pRAM data took only ~ 10 seconds using all 6000 time frames since there are only ~ 100 cells in the field of view for each region. We supply the MATLAB code to obtain the results presented in the paper.

4 Conclusion

We have shown the MRF model can quantify the spatial autocorrelation in neuronal activity data and can be applied to data obtained from both *in vitro* or *in vivo* experiments. We first analysed simulation experiments to demonstrate the model before showing its use in identifying purely spatial differences in neuron activity. We considered a proof of concept screening setup where we quantified changes in neuronal activity *in vitro* in response to distinct defined conditions. Finally, we showed that the MRF model is applicable to data from *in vivo* experiments and generates conclusions consistent with previous analysis. Moreover, we demonstrated how our approach may provide inference from observational calcium data that was not attainable from the analysis of spike counts, and nevertheless consistent with independent, more invasive, optogenetic experimentation.

The aim of the phenotypic quantification matches the particular features of the data we consider. That is, the analysis of the aggregate activity of large populations of upwards of thousands of neurons, not the specific activity of individual neurons. We make particular use of the spatial aspects of the data where the neuron ‘neighbours’ in the model are based on the relative spatial positions of the cells and a corresponding Voronoi tessellation. Then the resulting sparse adjacency matrix and Markov modelling assumption is what allows for feasible inference. Considering all possible pairs of neurons as is done elsewhere, in particular with Ising models, limits the analysis to only a handful of cells due to computational cost. We considered data with up to 5366 neurons (LF data) and 6000 time points (2pRAM data). We supply the data and MATLAB code to obtain all of the results in the paper.

Acknowledgements

Neurons were generously provided by Lili Li (BTK, University of Turku). The facilities of the Turku Bioimaging Screening Unit were used for this work. **We wish to thank the editor and reviewer for valuable comments and in particular suggestions of related literature.**

Funding

This work has been supported by the Department of Mathematics and Statistics at the University of Turku and The Finnish Cultural Foundation, Varsinais-Suomi Regional Fund (SR), and the University of Turku and National Institutes of Health National Cancer Institute Grant No. R01CA200417 (MJC).

References

- Abdollahi, L. M. O. M. *et al.* (2003). [Parameter estimation in a model for multidimensional recording of neuronal data: A Gibbsian approximation approach.](#) *Biol Cybern*, **89**(3), 170–178.
- Ahmadian, Y. *et al.* (2011). Efficient Markov chain Monte Carlo methods for decoding neural spike trains. *Neural Comput*, **23**(1), 46–96.
- Bading, H. *et al.* (1993). Regulation of gene expression in hippocampal neurons by distinct calcium signaling pathways. *Science*, **260**(5105), 181–186.
- Banerjee, S. *et al.* (2014). *Hierarchical Modeling and Analysis for Spatial Data*. CRC Press, second edition.
- Bell, B. S. and Broemeling, L. D. (2000). A Bayesian analysis for spatial processes with application to disease mapping. *Stat Med*, **19**(7), 957–974.
- Blanquie, O. *et al.* (2017). Homeostatic interplay between electrical activity and neuronal apoptosis in the developing neocortex. *Neuroscience*, **358**, 190–200.
- Blumberg, M. S. and Dooley, J. C. (2017). Phantom limbs, neuroprosthetics, and the developmental origins of embodiment. *Trends Neurosci*, **40**(10), 603–612.
- Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nat Neurosci*, **20**(3), 340–352.
- Chen, T.-W. *et al.* (2016). Calcium imaging responses from anterior lateral motor cortex (ALM) neurons of adult mice performing a tactile decision behavior. [CRCNS doi:10.6080/K04M92GX](#).
- Courtney, M. *et al.* (1990). The interactions between plasma membrane depolarization and glutamate receptor activation in the regulation of cytoplasmic free calcium in cultured cerebellar granule cells. *J Neurosci*, **10**(12), 3873–3879.
- Cressie, N. *et al.* (2005). Likelihood-based estimation for Gaussian MRFs. *Stat Methodol*, **2**(1), 1–16.
- Cunningham, J. P. and Byron, M. Y. (2014). Dimensionality reduction for large-scale neural recordings. *Nat Neurosci*, **17**(11), 1500–1509.
- De Oliveira, V. (2012). Bayesian analysis of conditional autoregressive models. *Ann Inst Stat Math*, **64**(1), 107–133.
- Doucet, M. V. *et al.* (2012). The PSD-95/nNOS complex: New drugs for depression? *Pharmacol Ther*, **133**(2), 218–229.
- Dravid, S. M. and Murray, T. F. (2004). Spontaneous synchronized calcium oscillations in neocortical neurons in the presence of physiological [Mg²⁺]: Involvement of AMPA/kainate and metabotropic glutamate receptors. *Brain Res*, **1006**(1), 8–17.
- Fernandes, D. and Carvalho, A. L. (2016). Mechanisms of homeostatic plasticity in the excitatory synapse. *J Neurochem*, **139**(6), 973–996.
- François, O. *et al.* (2000). [Statistical procedures for spatiotemporal neuronal data with applications to optical recording of the auditory cortex](#) *Neural Comput*, **12**(8), 1821–1838.
- Gelfand, A. E. and Vounatsou, P. (2003). Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics*, **4**(1), 11–15.
- Gulledge, A. T. and Stuart, G. J. (2003). Excitatory actions of GABA in the cortex. *Neuron*, **37**(2), 299–309.
- Harrison, M. T. *et al.* (2014). [Spatiotemporal conditional inference and hypothesis tests for neural ensemble spiking precision.](#) *Neural Comput*, **27**(1), 104–50
- Hilgen, G. *et al.* (2017). Unsupervised spike sorting for large-scale, high-density multielectrode arrays. *Cell Rep*, **18**(10), 2521–2532.
- Hill, M. R. *et al.* (2015). Quantification and classification of neuronal responses in kernel-smoothed peristimulus time histograms. *J Neurophysiol*, **113**(4), 1260–1274.
- Jacobs, K. *et al.* (2000). Postlesional epilepsy: The ultimate brain plasticity. *Epilepsia*, **41**(s6).
- Kress, G. J. *et al.* (2008). High threshold, proximal initiation, and slow conduction velocity of action potentials in dentate granule neuron mossy fibers. *J Neurophysiol*, **100**(1), 281–291.
- Kutsarova, E. *et al.* (2017). Rules for shaping neural connections in the developing brain. *Front Neural Circuits*, **10**, 111.
- Leighton, A. H. and Lohmann, C. (2016). The wiring of developing sensory circuits—from patterned spontaneous activity to synaptic plasticity mechanisms. *Front Neural Circuits*, **10**, 71.
- Li, L.-L. *et al.* (2013). The nNOS-p38MAPK pathway is mediated by NOS1AP during neuronal death. *J Neurosci*, **33**(19), 8185–8201.
- Li, L.-L. *et al.* (2015a). Unexpected heterodivalent recruitment of NOS1AP to nNOS reveals multiple sites for pharmacological intervention in neuronal disease models. *J Neurosci*, **35**(19), 7349–7364.
- Li, N. *et al.* (2015b). A motor cortex circuit for motor planning and movement. *Nature*, **519**(7541), 51–56.
- Luhmann, H. J. *et al.* (2016). Spontaneous neuronal activity in developing neocortical networks: From single cells to large-scale interactions. *Front Neural Circuits*, **10**, 40.
- Makarenko, V. I. *et al.* (1997). [A new approach to the analysis of multidimensional neuronal activity: Markov random fields](#) *Neural Netw*, **10**(5), 785–789.
- Mazzucato, L. *et al.* (2016). Stimuli reduce the dimensionality of cortical activity. *Front Syst Neurosci*, **10**, 11.
- Nirenberg, S. H. and Victor, J. D. (2007). Analyzing the activity of large populations of neurons: How tractable is the problem? *Curr Opin Neurobiol*, **17**(4), 397–400.
- Okun, M. *et al.* (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature*, **521**(7553), 511–515.
- Onken, A. *et al.* (2016). Using matrix and tensor factorizations for the single-trial analysis of population spike trains. *PLOS Comput Biol*, **12**(11), e1005189.
- Pawlak, V. *et al.* (2005). Impaired synaptic scaling in mouse hippocampal neurons expressing NMDA receptors with reduced calcium permeability. *J Physiol*, **562**(3), 771–783.
- Pratt, K. G. *et al.* (2016). An evolutionarily conserved mechanism for activity-dependent visual circuit development. *Front Neural Circuits*, **10**, 79.
- Ren, C. and Sun, D. (2013). Objective Bayesian analysis for CAR models. *Ann Inst Stat Math*, **65**(3), 457–472.
- Roudi, Y. *et al.* (2009). Pairwise maximum entropy models for studying large biological systems: when they can work and when they can't. *PLOS Comput Biol*, **5**(5), e1000380.
- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: Theory and applications*. CRC Press.
- Schmidt-Hieber, C. *et al.* (2008). Action potential initiation and propagation in hippocampal mossy fibre axons. *J Physiol*, **586**(7), 1849–1857.
- Schneidman, E. *et al.* (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, **440**(7087), 1007–1012.
- Sharma, P. *et al.* (2012). High-throughput screening in primary neurons. *Methods Enzymol*, **506**, 331–360.
- Shlens, J. *et al.* (2006). The structure of multi-neuron firing patterns in primate retina. *J Neurosci*, **26**(32), 8254–8266.
- Sigler, A. *et al.* (2017). Formation and maintenance of functional spines in the absence of presynaptic glutamate release. *Neuron*, **94**(2), 304–311.
- Smith, S. L. and Häusser, M. (2010). Parallel processing of visual space by neighboring neurons in mouse visual cortex. *Nat Neurosci*, **13**(9), 1144–1149.
- Sofroniew, N. J. *et al.* (2016). A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *Elife*, **5**, e14472.
- Tatti, R. *et al.* (2017). Neurophysiology and regulation of the balance between excitation and inhibition in neocortical circuits. *Biol Psychiatry*, **81**(10), 821–831.
- Tkačik, G. *et al.* (2014). Searching for collective behavior in a large network of sensory neurons. *PLOS Comput Biol*, **10**(1), e1003408.
- Zhuo, M. (2008). Cortical excitation and chronic pain. *Trends Neurosci*, **31**(4), 199–207.