



Modeling statistical dependencies in multi-region spike train data

Stephen L Keeley, David M Zoltowski, Mikio C Aoi and Jonathan W Pillow

Neural computations underlying cognition and behavior rely on the coordination of neural activity across multiple brain areas. Understanding how brain areas interact to process information or generate behavior is thus a central question in neuroscience. Here we provide an overview of statistical approaches for characterizing statistical dependencies in multi-region spike train recordings. We focus on two classes of models in particular: regression-based models and shared latent variable models. Regression-based models describe interactions in terms of a directed transformation of information from one region to another. Shared latent variable models, on the other hand, seek to describe interactions in terms of sources that capture common fluctuations in spiking activity across regions. We discuss the advantages and limitations of each of these approaches and future directions for the field. We intend this review to be an introduction to the statistical methods in multi-region models for computational neuroscientists and experimentalists alike.

Address

Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

Corresponding author: Pillow, Jonathan W (pillow@princeton.edu)

Current Opinion in Neurobiology 2020, **65**:194–202

This review comes from a themed issue on **Whole-brain interactions between neural circuits**

Edited by **Larry Abbott** and **Karel Svoboda**

For a complete overview see the [Issue](#) and the [Editorial](#)

Available online 14th December 2020

<https://doi.org/10.1016/j.conb.2020.11.005>

0959-4388/© 2020 Elsevier Ltd. All rights reserved.

Introduction

The idea that distinct areas of the brain support different cognitive and behavioral functions dates back to the early 19th century [1]. During this time, there was debate as to whether various cortical regions had specific functionality or rather participated in all cognitive and psychological functions [2,3]. These competing ideas remain a topic of study today [4], and despite significant progress in understanding brain function over the past two centuries, the extent to which different brain areas support distinct

functions and convey different information to other brain areas remains largely unknown.

In the past decade, advances in neural recording technologies have provided exciting new opportunities to study how distinct brain regions interact. In particular, advances in wide field calcium imaging [5–7] and high-throughput neural probes [8] allow for simultaneous recording of neural activity from hundreds to thousands of neurons in multiple brain regions (Figure 1a). Recent work has made substantial progress in characterizing neural population codes within single brain regions [9–15]. The challenge remains to extend these models to study how populations interact. This necessitates both adding to existing statistical tools and developing new tools to extract accurate and interpretable descriptions of interactions from multi-region neural recordings.

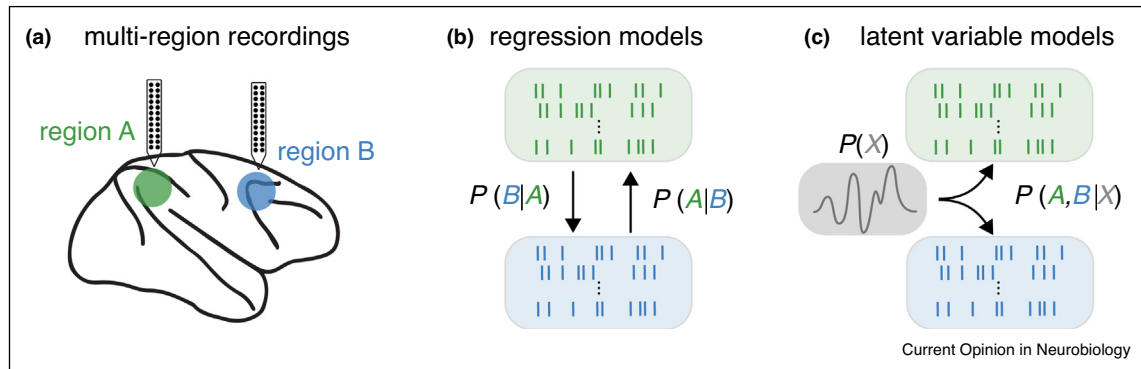
A simple approach to analyzing multi-region datasets is to compute descriptive statistics that summarize the statistical relationships between regions. For example, an extensive literature has focused on using the cross-correlogram or cross-covariance between spike trains to make inferences about connectivity [16–19]. In recent work, covariance-based measures have been used to identify functional subnetworks across visual areas in rodent visual cortex [5,20,15]. However, moment-based methods are typically limited by the lack of an explicit statistical model for neural responses, and may lack statistical power for identifying weak interactions between areas. In this review, we therefore focus on model-based approaches, which provide explicit descriptions of the dependencies within and between brain areas.

Here we provide an overview of two modern approaches for characterizing dependencies in multi-region electrophysiological recordings. We begin by discussing regression-based approaches, which seek to build predictive models for neural activity in one region using neural activity in other regions (Figure 1b). We then discuss latent variable models, which characterize the structure of multi-region data in terms of shared low-dimensional time series (Figure 1c). We conclude by proposing future directions and ways of extending existing models to help further uncover the ways distinct brain regions interact.

Regression-based approaches

Regression models provide one framework for characterizing statistical dependencies between brain areas using

Figure 1



Approaches to characterizing neural activity from simultaneously recorded brain regions. **(a)** Schematic for simultaneous electrophysiological recording from two separate brain regions (adapted from [22]). **(b)** Regression models are directional statistical descriptions that seek to characterize the firing of one brain region using the activity of the other region. **(c)** Latent variable models characterize shared structure underlying response fluctuations in multiple brain regions.

data from multi-area recordings. The basic idea is to fit a model that predicts activity in one brain area using the activity from one or more other brain areas (Figure 1b). Models of this form contain a set of linear weights that specify how spikes from a group of input neurons modulate the firing rate of an output neuron, with Gaussian [21••] or Poisson observation noise and time-lag between regions [22•,23•]. The inferred regression weights in this class of models provide a measure of statistical dependency, often referred to as *functional connectivity*, between areas [22•,24]. When the regression model in question is a time-lagged linear-Gaussian ‘least-squares’ regression, the resulting dependency measure is known as *Granger causality* [25,26]. When the model consists of a linear-mapping following by a point-wise non-linearity and Poisson observations, the method is known as the Poisson generalized linear model (GLM). Recent work has extended this framework to allow for non-linear dependencies between neurons using deep neural networks [27]).

For this review, we focus on the Poisson GLM, which provides a simple yet powerful regression modeling framework for spike train data. The single-neuron Poisson GLM describes how the spike rate of a single neuron depends on external covariates such as visual stimuli, spatial position, or theta oscillation phase [28–32]. So called ‘auto-regressive’ Poisson GLMs include weights on spike history from one or more neurons, which allow them to characterize dependencies between neurons in multi-region recordings [22•,23•,24,33–39].

Formally, a Poisson GLM describes each neuron by an instantaneous firing rate, also known as the *conditional intensity*, denoted λ_t , which defines the probability of observing a spike in a small time bin at time t . This firing

rate is given by the projection of an input vector \mathbf{x}_t onto the linear weights \mathbf{w} , transformed by a rectifying nonlinear function f :

$$\lambda_t = f(\mathbf{w}^\top \mathbf{x}_t + b), \quad (1)$$

where b is a constant offset or bias term. Common choices of nonlinearity f include exponential and the soft-rectified or ‘softplus’ function, $f(x) = \log(1 + e^x)$, both of which ensure that firing rates are non-negative. The spike count y_t in a time bin of size Δ has a conditionally Poisson distribution:

$$y_t | \mathbf{x}_t \sim \text{Poisson}(\Delta \lambda_t), \quad (2)$$


although other count distributions (e.g. Bernoulli, binomial, negative binomial) have also been considered [40–43].

To model statistical dependencies between neurons, the spike history of other neurons can be included as a regressor. In this case the firing rate of neuron j can be re-written as:

$$\lambda_t^{(j)} = f(\mathbf{w}_j^\top \mathbf{x}_t + \mathbf{h}_j^\top \mathbf{y}_{\text{hist}_t}^{(j)} + \sum_{i \neq j} \mathbf{c}_{ij}^\top \mathbf{y}_{\text{hist}_t}^{(i)} + b_j) \quad (3)$$

where \mathbf{w}_j and b_j are stimulus weights and offset for neuron j , \mathbf{h}_j denotes the ‘self-coupling’ spike-history filter for neuron j , \mathbf{c}_{ij} denotes a ‘coupling’ filter from neuron i to neuron j , and $\mathbf{y}_{\text{hist}_t}^{(i)}$ denotes a vector representing the spike history of neuron i at time t . This is typically defined to be a vector of binned spike counts from time bins

Figure 2



The coupled Poisson GLM is a regression model for capturing dependencies in multi-region spike train recordings. To predict spike trains in region B using spike trains from region A , the model contains 'coupling' filters C_{ij} that capture statistical dependencies between previous spikes from neuron i and the spike rate of neuron j at time t . It may also contain filters capturing dependencies on external stimuli, and on the spike history of other neurons in region B (not shown). The full Poisson GLM for a single brain region contains an independent set of filters for every spike train, meaning that it is really a collection of independent models, one for every neuron in downstream population B . One can similarly construct a Poisson GLM to predict spike trains in region A from the spikes in region B .

$(t - \tau, \dots, t - 1)$, for some number of time lags τ before the current time bin.

When defined this way, it can be seen that the coupled Poisson GLM is actually a collection of single-neuron models, one for each neuron in the population (schematized for a single neuron in Figure 2). Each neuron is characterized by the set of filters that determine its inputs. Fitting the coupled Poisson GLM to multi-neuron data therefore involves fitting a separate GLM to each neuron in the population, which can be carried out in parallel.

One way to characterize statistical dependencies with the Poisson GLM is to analyze the coupling filters between neurons in different brain regions. The strength of dependencies between regions can be assessed by computing the predictive performance of models fit with and without coupling filters between regions [22^{*}]. A simple metric is the difference in cross-validated log-likelihood, given by $\log P(Y_{test}|\theta_{coupled}) - \log P(Y_{test}|\theta_{uncoupled})$, where Y_{test}

represents held-out test data not used for fitting, and $\theta_{coupled}$ and $\theta_{uncoupled}$ represent fitted parameters of the coupled and uncoupled models, respectively. This quantity has units of bits when taking log base 2, and can be divided by the number of seconds or number of samples to get units of bits/s or bits/sample [31,42]. Another way to use GLMs to characterize dependencies between brain regions is to directly analyze the shape and amplitude of fitted coupling filters. The peak amplitude of each coupling filter provides a summary of coupling strength, while the sign of the peak indicates whether the effective dependencies are excitatory or suppressive [38^{*}].

One general advantage of regression-based approaches is the ability to characterize directed dependencies between brain regions. Because we regress the activity in each brain area against the time-lagged activity from other areas, we can identify asymmetric relationships between areas. For example, we might identify that spikes in area A predict increased spiking area B , whereas spikes in area B predict decreased firing in area A . However, it is important to emphasize that the statistical dependencies identified by GLMs are statistical and *not* causal, and are susceptible to omitted variable bias [44]. Thus, for example, the dependencies between two brain regions identified by GLM analysis may in fact reflect shared input with different time lags from a third, unobserved brain region [45,46].

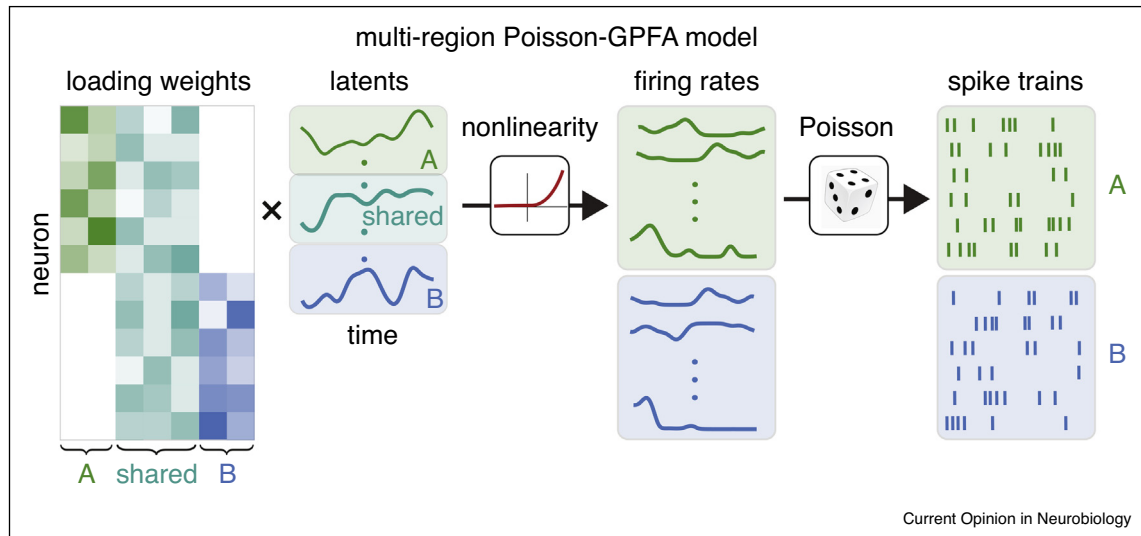
Another advantage of GLMs is the ability to parallelize over neurons when fitting, since the full model is simply a collection of single-neuron models. This confers computational advantages over latent variable models (see next section), which must be fit jointly to all neurons at once. However, the total number of coupling filters in a GLM grows quadratically with the number of neurons, which can lead to overfitting and problems of interpretation in large neural populations. Sparsity-inducing priors, which drive some coupling filters to zero, can help mitigate both problems; these priors reduce the number of model parameters while preserving the most significant connections between neurons [47,35,38^{*}].

Latent variable modeling approaches

An alternative to regression-based approaches is to model dependencies between brain regions using shared latent variables. Here, the standard approach is to model the response fluctuations in a large population of neurons as arising from a small number of hidden or 'latent' sources. These models are typically not directed, meaning that they seek to describe dependencies between brain regions in terms of shared modulation by a small number of hidden or unobserved sources, as opposed to seeking to describe dependencies in terms of separate feedforward and feedback components.

In recent years, latent variable models have become a popular approach for analyzing structure in single-region

Figure 3



The multi-region Poisson GPFA (Gaussian-Process Factor Analysis) model contains a block-structured loadings matrix, with latent variables local to each region, and ‘shared’ latents that capture shared variability across regions. Latent firing rates are passed through a rectifying nonlinearity f to obtain firing rates, which drive spiking via a Poisson process.

datasets. Typically, the time-course of the latent variable is either assumed to follow a linear dynamical system (LDS) [48,45,49–53], or is simply assumed to be smooth in time using a Gaussian process (GP) prior [54–58,43,59,60]. Both LDS and GP latent variable models have been augmented and modified for different purposes, and here we will highlight how each form can be extended to understand multi-region data.

The key distinction between LDS and GP latent variable models lies in the assumed prior over the latent variable. In LDS models, the latent vector \mathbf{x}_t evolves according to linear dynamics with Gaussian noise:

$$\mathbf{x}_t = A\mathbf{x}_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, Q) \quad (4)$$

where A describes the discrete-time linear dynamics. Although the assumption of linear dynamics is almost certainly too restrictive for real neural circuits, the noise component ε_t , often referred to as ‘innovations’ noise, can be viewed as accounting for both the true noise as well model mismatch. One notable extension of this model, the switching Linear Dynamical System (sLDS), allows for discrete switching between different linear dynamics matrices, which can approximate nonlinear dynamics in the same way that a series of line segments can approximate a curve [61–65].

An alternate approach to LDS models is to model the latent variable with a Gaussian process (GP), which

makes no assumption about the dynamics beyond smoothness in time. A GP defines a joint Gaussian distribution over the latent x_t (an element of the latent vector \mathbf{x}_t) at all time points $t \in [0, T]$. The covariance of this latent process is specified by a kernel function:

$$\text{cov}(x_t, x_{t'}) = k(t, t') \quad (5)$$

which determines the *a priori* correlation between the latent variable at any pair of times t and t' . A common choice of kernel function is the ‘Gaussian’ or ‘radial basis function’ (RBF) covariance:

$$k(t, t') = e^{-\frac{(t-t')^2}{2\ell^2}} \quad (6)$$

which is governed by a single hyperparameter ℓ known as the ‘length scale’ that determines the degree of smoothness. In a classic model known as Gaussian process factor analysis (GPFA), each element of the latent vector \mathbf{x}_t is assumed to have an independent zero-mean GP prior, each with its own length scale [54].

In both LDS and GP latent variable models, the probability distribution over neural data is defined in terms of a noisy projection of the low-dimensional latent variable to the space of high-dimensional neural activity. The simplest case is to assume a linear mapping, as found in a class of so-called ‘factor models’. If we assume Gaussian additive noise with covariance R , we obtain [54]:

$$\mathbf{y}_t = W\mathbf{x}_t + \varepsilon_t, \quad \varepsilon \sim N(0, R) \quad (7)$$

where \mathbf{y}_t is the set of n neurons' spike counts at time t , \mathbf{x}_t is the p -element latent vector at time t , and W is a $(n \times p)$ matrix of *loading weights*. Each column of W describes how a single latent affects the population spiking activity, and each row describes how a single neuron's activity depends on the different latents contained in \mathbf{x}_t .

A more appropriate model for spike count data is to assume Poisson noise after transformation by a fixed nonlinearity. In this case, the factor model can be written:

$$\mathbf{y}_t | \mathbf{x}_t \sim \text{Pois}(\Delta f(W\mathbf{x}_t)), \quad (8)$$

whereas before Δ denotes time bin width and $f(\cdot)$ is a fixed non-negative output function [64,57,43]. Gaussian models are tractable and simple to fit because the log-likelihood can be evaluated in closed form. Poisson models, by contrast, may be more accurate for spike count data, particularly when spike rates are low, but require approximate inference methods to fit [66,67,57,64,43]. Recent work has introduced latent variable models with flexible nonlinear mappings between latent variables and observations [66–68,60,69], but for the purposes of this review we focus on models with linear mappings.

To extend the two latent variable modeling approaches described above to multi-region data, we will first consider recording from two brain regions, denoted A and B . We can perform inference on an LDS or GP factor model by concatenating data from two regions, and then examine the inferred loadings matrix, W . In the standard inference case, the model will uncover latent structure agnostic to the multi-region nature of the data. The latent factors will reflect the dimensions of shared variability, across time, for the entire population. A way to isolate per-region latent structure is to constrain the loadings matrix, W , to have block structure, with a subset of latents mapping only to region A , and another subset mapping only to region B [81]. The model, shown for the Poisson observation case, takes the following form:

$$\begin{bmatrix} \mathbf{y}_A \\ \mathbf{y}_B \end{bmatrix} = \text{Pois} \left(f \left(\begin{bmatrix} W_A & 0 \\ 0 & W_B \end{bmatrix} \begin{bmatrix} \mathbf{x}_A \\ \mathbf{x}_B \end{bmatrix} \right) \right), \quad (9)$$

where \mathbf{x}_A and \mathbf{x}_B reflect the latents corresponding to regions A and B and $W_{A,B}$ the respective loadings for each region, and for simplicity we have dropped indexing on time.

If the latent variables \mathbf{x}_A and \mathbf{x}_B have independent noise and dynamics, this model is equivalent to fitting two

entirely separate models, one for region A and another for region B , with no shared structure between them. We wish to contrast the above model with one that includes both shared and independent latent variables:

$$\begin{bmatrix} \mathbf{y}_A \\ \mathbf{y}_B \end{bmatrix} = \text{Pois} \left(f \left(\begin{bmatrix} W_A & W_{SA} & 0 \\ 0 & W_{SB} & W_B \end{bmatrix} \begin{bmatrix} \mathbf{x}_A \\ \mathbf{x}_S \\ \mathbf{x}_B \end{bmatrix} \right) \right). \quad (10)$$

Here \mathbf{x}_S refers to latent variables that are shared between regions, while \mathbf{x}_A and \mathbf{x}_B denote region-specific latent variables. This block structure in the loadings matrix W allows us to extract shared and independent sources of variability using either GP [70•,71•] or LDS [65] priors over the latent variables. We show a schematic for this model in Figure 3. Here, a loadings matrix with independent components and shared components map independent and shared latents to a set of firing rates. These firing rates describe Poisson spiking from populations of neurons in each brain region.

These shared latent variable models for multi-region data are structurally related to canonical correlation analysis (CCA) [72], which has been used in neuroscience settings for multi-region data before [4]. In particular, the LVMs outlined above relate closely to statistical variants of CCA known in the machine learning community as probabilistic CCA [73,81] and Bayesian CCA (BCCA) [74]. Though probabilistic CCA and BCCA do not typically include neuroscience-specific model features such as point-wise nonlinearities, Poisson observations, and GP priors, they are factor analytic models which identify shared low-dimensional subspaces that map to distinct upstream observations [73,81]. In the case of BCCA, these statistical models, like the LVMs described above, assert region-specific and shared factors through a block-structured loadings matrix [74,75]. The ongoing work in the machine learning community on probabilistic CCA and Bayesian CCA may have useful implications to the neuroscience community. In particular, recent developments of incorporating sparsity priors over the block loadings matrix may be well suited to neuroscience contexts [75].

Implementing LVMs with block loadings structure is an exciting future for multi-region analysis. Some preliminary results suggest that models of this sort can incorporate time-delays to identify latents that represent feed-forward and feedback communication between brain regions [71•]. Other work has isolated trial-varying latent structure that is shared across regions during repeated presentations of a visual stimulus [70••]. However, a technical challenge here remains in properly identifying the dimensionality of the shared and independent subspaces, as an exhaustive search over all possibilities is computationally cumbersome. Greedy methods, which add dimensions one-by-one if it improves cross-validation performance, are a sensible initial approach [70••]. An

exciting possibility would be to impose an explicit regularization over the loadings matrix, similar to the approaches used in [75], such as an automatic relevance determination prior, or nuclear norm regularization. These more sophisticated statistical methods would isolate within and across region subspaces without the need directly specify the nature of the block structured loadings structure.

Hybrid approaches: low-dimensional regression models

So far we have considered regression models, which seek to explain activity in one brain area as a function of activity in another, and latent variable models, which seek to explain activity in multiple brain areas using a small number of components or factors. It is of course natural to seek to combine the strengths of these two approaches: models that can explain activity in one brain area in terms of a small number of components arising from another brain area. Dimensionality-reduction methods have had great success in extracting interpretable features from high-dimensional neural recordings [9,82]. Combining dimensionality reduction with regression thus provides a powerful approach for extracting more interpretable statistical interactions between brain regions [21,81]. Although we have emphasized models with Poisson noise up to now, this section will focus primarily on models with Gaussian noise.

A simple approach to low-dimensional regression for multi-region data is to work in two stages: first perform dimensionality reduction on the upstream region, and then use regression on this low dimensional representation to predict downstream activity. When using principal components analysis (PCA) for dimensionality reduction followed by ordinary least squares regression, this technique is known as *principal components regression* (PCR). Kaufman *et al.* [76] used PCR to analyze the relationship between neural activity in pre-motor (PMd) and primary motor cortex (M1) during movement preparation, arguing that PMd activity falls primarily in the null space of the linear mapping to M1. Semedo *et al.* [21**] used a similar approach, substituting factor analysis for PCA, to analyze dependencies in neural activity between V1 and V2, concluding that principal axes of V1 variability were *not* the dimensions that best predicted activity in downstream area V2.

An alternate approach to low-dimensional regression modeling is reduced rank regression (RRR), which unifies dimensionality reduction and regression into a single model. The RRR model explicitly imposes a rank constraint on the *parameters* of a regression model that govern the mapping from one brain region to another. For linear regression with Gaussian noise, the RRR objective is

$$\min_{W \in \mathbb{R}^{m \times n}} \|Y_B - Y_A W\|_2^2, \quad \text{such that } \text{rank}(W) \leq r \quad (11)$$

where Y_A denotes a $T \times m$ matrix of responses from m area-A neurons over T time bins, Y_B denotes a $T \times n$ matrix of responses from n area-B downstream neurons, and W is a $m \times n$ matrix of regression weights of rank at most r .

Remarkably, there is a closed-form solution to this optimization problem. This solution can be obtained using the top r right singular vectors of $Y_A \hat{W}_{OLS}$, where \hat{W}_{OLS} is the matrix of ordinary least squares regression weights, $\hat{W}_{OLS} = (Y_A^T Y_A)^{-1} Y_A^T Y_B$. If we assemble these singular vectors into a matrix V , the RRR solution is simply $\hat{W}_{RRR} = \hat{W}_{OLS} V V^T$. The reduced-rank regression model can of course be extended to the setting of Poisson GLMs, although in this case fitting requires numerical optimization of the log-likelihood, as there is no longer an eigenvector solution.

The rank constraint in the reduced-rank regression model can be seen as imposing a bottleneck for information transmission between brain areas. For example, if we use a rank-2 RRR model, the interactions between brain areas are confined to a 2-dimensional subspace, regardless of the total number of latent dimensions in each area. Recent work from Semedo *et al.* [21**] coined the term ‘communication subspace’ to refer to this phenomenon, and showed that the dependencies between neural activity in visual areas V1 and V2 are indeed low-dimensional relative to the dimensionality of activity in each brain region.

The rank constraint in RRR also provides a form of regularization by reducing the number of parameters and allowing for parameter-sharing across neurons. The RRR model may therefore achieve better generalization performance than a (full-rank) GLM with all-to-all coupling due to reduced overfitting. More importantly, RRR allows for direct estimation of the communication subspace from one population to another, without requiring separate steps for dimensionality-reduction and regression. In addition to being more accurate, the RRR model is also more interpretable due to the fact that it provides a concise summary of the statistical dependencies between areas.

Challenges and opportunities

Recent advances in technology have ushered in a new era for multi-region optical and electrophysiological recording. These techniques hold great potential for deciphering the complex flow of information within neural circuits. However, the challenges posed by analyzing such datasets highlight the need for innovative new statistical tools and approaches. As we move forward with the development of such tools, we propose that future work be guided by the principles of interpretable structure, scalability, and statistical efficiency.

Thinking carefully about interpretable statistical structure is important due to the large number of possible modelling options for multi-region data. When developing a model for multi-region data, scientists should consider what the model fits might tell us in the context of a particular hypothesis or set of hypotheses of interest. Successful examples include multi-region models built to understand signal and noise representations [70^{**},21^{**}], feedforward and feedback signals [71^{*}], or functional coupling properties across and within regions [23^{*}]. In each of these cases, a particular neuroscience question or network property motivates the development of the model, and the resultant fits provide insight into the network property of interest. Developing or using a multi-region model without first thinking about how to interpret it may be a useful mathematical exercise, but a good fit to neural data will not on its own provide any insight into multi-region neural activity.

The scalability of statistical methods for analyzing multi-region data is vital for future work. High-throughput imaging and electrophysiology preparations will soon be routinely recording from thousands of neurons simultaneously. The estimation procedures used to fit the models described here must be able to contend not only with an expanding volume of data, but with the exploding degrees of freedom that accompany it. Fitting regression models in the large-scale setting, including parameters characterizing the coupling between cells, results in geometric growth in the number of parameters with the number of cells. One promising solution is to rely on approximate modeling methods [38^{*}]. Hybrid approaches like those described in the previous section also serve to solve the scalability problem by reducing the degrees of freedom of the model.

We are especially excited about possibilities for combining latent variable models and regression models, and the potential for such combined models to reveal new insights into the flow of information between and within brain regions. Model-based targeted dimensionality reduction [77,78] represents one recent method that combines a regression model with explicit low-dimensional constraints on population activity. Bayesian latent structure discovery [79] is another approach that merges regression and latent variable models. Both of these approaches model population-level activity by sharing degrees of freedom among cells but leave enough flexibility to identify functionally significant, low dimensional latent structure. These approaches have yet to be used in the multi-region setting but we expect their applicability to the multi-region setting to be straight forward.

Conclusion

The development of statistical analysis tools for multi-region data is in its infancy. We believe that a conceptual understanding of information processing in the brain

depends on appropriate characterization of population-level activity [80]. As multi-region datasets become increasingly common, there will be greater need for analytical tools and appropriate statistical summaries for each region as well as inter-region interactions. Here, we have described two general methodological approaches for understanding multi-region spike train data, which represents an important avenue for current and future research..

Conflict of interest statement

Nothing declared.

Acknowledgements

SLK was supported by NIH grant F32MH115445-03. JWP was supported by grants from the Simons Collaboration on the Global Brain (SCGB AWD543027), the NIH BRAIN Initiative (NS104899 and R01EB026946), and a U19 NIH-NINDS BRAIN Initiative Award (5U19NS104648).

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
 - of outstanding interest
1. Gall FJ: *On the Functions of the Brain and of Each of Its Parts: With Observations on the Possibility of Determining the Instincts, Propensities, and Talents, or the Moral and Intellectual Dispositions of Men and Animals, by the Configuration of the Brain and Head*, vol 1. Marsh, Capen & Lyon; 1835.
 2. Gross CG: *A Hole in the Head: More Tales in the History of Neuroscience*. MIT Press; 2012.
 3. Clarke E, O'Malley CD: *The Human Brain and Spinal Cord: A Historical Study Illustrated By Writings From Antiquity to the Twentieth Century*. Norman Publishing; 1996.
 4. Steinmetz NA, Zátka-Haas P, Carandini M, Harris KD: **Distributed coding of choice, action and engagement across the mouse brain**. *Nature* 2019, **576**:266-273.
 5. Stirman JN, Smith IT, Kudenov MW, Smith SL: **Wide field-of-view, multi-region, two-photon imaging of neuronal activity in the mammalian brain**. *Nat Biotechnol* 2016, **34**:857.
 6. Stosiek C, Garaschuk O, Holthoff K, Konnerth A: **In vivo two-photon calcium imaging of neuronal networks**. *Proc Natl Acad Sci U S A* 2003, **100**:7319-7324.
 7. Tian L, Andrew Hires S, Mao T, Huber D, Eugenia Chiappe M, Chalasani SH, Petreanu L, Akerboom J, McKinney SA, Schreiner ER *et al.*: **Imaging neural activity in worms, flies and mice with improved gcamp calcium indicators**. *Nat Methods* 2009, **6**:875.
 8. Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, Lee AK, Anastassiou CA, Andrei A, Aydn Ç *et al.*: **Fully integrated silicon probes for high-density recording of neural activity**. *Nature* 2017, **551**:232-236.
 9. Cunningham JP, Yu Byron M: **Dimensionality reduction for large-scale neural recordings**. *Nat Neurosci* 2014, **17**:1500-1509.
 10. Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A: **Information-limiting correlations**. *Nat Neurosci* 2014, **17**:1410-1417.
 11. Tkačik G, Marre O, Amodei D, Schneidman E, Bialek W, Berry M II: **Searching for collective behavior in a large network of sensory neurons**. *PLoS Comput Biol* 2014, **10**:e1003408.
 12. Stringer C, Pachitariu M, Steinmetz N, Carandini M, Harris KD: **High-dimensional geometry of population responses in visual cortex**. *Nature* 2019:1.

13. Stringer C, Michaelos M, Pachitariu M: **High Precision Coding in Visual Cortex.** 2019:679324. BioRxiv.
14. Rummyantsev OI, Lecoq JA, Hernandez O, Zhang Y, Savall J, Chrapkiewicz R, Li J, Zeng H, Ganguli S, Schnitzer MJ: **Fundamental bounds on the fidelity of sensory cortical coding.** *Nature* 2020, **580**:100-105.
15. de Vries SEJ, Lecoq JA, Buice MA, Groblewski PA, Ocker GK, Oliver M, Feng D, Cain N, Ledochowitsch P, Millman D et al.: **A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex.** *Nat Neurosci* 2020, **23**:138-151.
16. Moore GP, Segundo JP, Perkel DH, Levitan H: **Statistical signs of synaptic interaction in neurons.** *Biophys J* 1970, **10**:876.
17. Allum JHJ, Hepp-Reymond M-C, Gysin R: **Cross-correlation analysis of interneuronal connectivity in the motor cortex of the monkey.** *Brain Res* 1982, **231**:325-334.
18. Toyama K, Kimura M, Tanaka K: **Cross-correlation analysis of interneuronal connectivity in cat visual cortex.** *J Neurophysiol* 1981, **46**:191-201.
19. De Blasi S, Ciba M, Bahmer A, Thielemann C: **Total spiking probability edges: a cross-correlation based method for effective connectivity estimation of cortical spiking neurons.** *J Neurosci Methods* 2019, **312**:169-181.
20. Yu Y, Stirman JN, Dorsett CR, Smith SL: **Mesoscale correlation structure with single cell resolution during visual coding.** *bioRxiv* 2019.
21. Semedo JD, Zandvakili A, Machens CK, Yu Byron M, Kohn A: **Cortical areas interact through a communication subspace.** *Neuron* 2019, **102**:249-259
 Use regression based approaches in combination with dimensionality reduction techniques to demonstrate that V1 and V2 activity interact via a 'communication subspace'. In particular, the authors demonstrate that low-dimensional mappings of V1 activity best predict V2 activity, and this dimensionality does not align with the principle dimensions of V1.
22. Yates JL, Park IL, Katz LN, Pillow JW, Huk AC: **Functional dissection of signal and noise in mt and lip during decision-making.** *Nat Neurosci* 2017, **20**:1285
 Fit multi-region Poisson GLMs with coupling filters to investigate functional connections between MT and LIP neurons. The authors show that trial-by-trial activity in LIP does not depend on MT activity.
23. Hart E, Huk AC: **Recurrent circuit dynamics underlie persistent activity in the macaque frontoparietal network.** *eLife* 2020, **9**:e52460
 Use a generalized linear model (GLM) to identify functional coupling between activity in LIP and FEF (parietal-frontal) areas. They used smooth temporal basis functions in their GLM reduce dimensionality of the coupling filters. Model fits show strong long-timescale recurrent excitation between LIP and FEF.
24. Okatan M, Wilson M, Brown E: **Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity.** *Neural Comput* 2005, **17**:1927-1961.
25. Quinn C, Coleman T, Kiyavash N, Hatsopoulos N: **Estimating the directed information to infer causal relationships in ensemble neural spike train recordings.** *J Comput Neurosci* 2011, **30**:17-44 10.1007/s10827-010-0247-2.
26. Seth AK, Barrett AB, Barnett L: **Granger causality analysis in neuroscience and neuroimaging.** *J Neurosci* 2015, **35**:3293-3297.
27. Moore BJ, Berger T, Song D: **Validation of a convolutional neural network model for spike transformation using a generalized linear model.** 2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC) 2020:3236-3239.
28. Brillinger DR: **Maximum likelihood analysis of spike trains of interacting nerve cells.** *Biol Cybern* 1988, **59**:189-200.
29. Paninski L: **Maximum likelihood estimation of cascade point-process neural encoding models.** *Network: Comput Neural Syst* 2004, **15**:243-262.
30. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN: **A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects.** *J Neurophysiol* 2005, **93**:1074-1089.
31. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP: **Spatio-temporal correlations and visual signalling in a complete neuronal population.** *Nature* 2008, **454**:995-999.
32. Hardcastle K, Maheswaranathan N, Ganguli S, Giocomo LM: **A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex.** *Neuron* 2017, **94**:375-387.
33. Brown EN, Kass RE, Mitra PP: **Multiple neural spike train data analysis: state-of-the-art and future challenges.** *Nat Neurosci* 2004, **7**:456-461.
34. Ventura V, Cai C, Kass RE: **Trial-to-trial variability and its effect on time-varying dependency between two neurons.** *J Neurophysiol* 2005, **94**:2928-2939.
35. Stevenson IH, Rebesch JM, Hatsopoulos NG, Haga Z, Miller LE, Kording KP: **Bayesian inference of functional connectivity and network structure from spikes.** *IEEE Trans. Neural Syst. Rehabil. Eng.* 2008, **17**:203-213.
36. Rikhye RV, Gilra A, Halassa MM: **Thalamic regulation of switching between cortical representations enables cognitive flexibility.** *Nat Neurosci* 2018, **21**:1753-1763.
37. Perich MG, Gallego JA, Miller LE: **A neural population mechanism for rapid learning.** *Neuron* 2018, **100**:964-976.
38. Zoltowski D, Pillow JW: **Scaling the Poisson glm to massive neural datasets through polynomial approximations.** *Advances in Neural Information Processing Systems* 2018:3517-3527
 Develop scalable methods for fitting Poisson GLMs. They apply the methods to fit a fully coupled GLM with sparse priors on connection weights to simultaneous recordings of ~800 neurons.
39. Dowling M, Zhao Y, Park IM: *Non-Parametric Generalized Linear Model.* 2020.
40. Pillow JW, Scott J: **Fully Bayesian inference for neural models with negative-binomial spiking.** *Advances in Neural Information Processing Systems* 2012:1898-1906.
41. Goris RLT, Movshon JA, Simoncelli EP: **Partitioning neuronal variability.** *Nat Neurosci* 2014, **17**:858-865.
42. Williamson RS, Sahani M, Pillow JW: **The equivalence of information-theoretic and likelihood-based methods for neural dimensionality reduction.** *PLoS Comput Biol* 2015, **11**:e1004141 04.
43. Keeley S, Zoltowski D, Yu Y, Smith S, Pillow J: **Efficient Non-Conjugate Gaussian process factor models for spike count data using polynomial approximations.** *In International Conference on Machine Learning.* PMLR; 2020:5177-5186.
44. Stevenson IH: **Omitted variable bias in glms of neural spiking activity.** *Neural Comput* 2018, **30**:3227-3258.
45. Macke JH, Buesing L, Cunningham JP, Yu Byron M, Shenoy KV, Sahani M: **Empirical models of spiking in neural populations.** *Advances in Neural Information Processing Systems* 2011:1350-1358.
46. Vidne M, Ahmadian Y, Shlens J, Pillow JW, Kulkarni J, Litke AM, Chichilnisky EJ, Simoncelli E, Paninski L: **Modeling the impact of common noise inputs on the network activity of retinal ganglion cells.** *J Comput Neurosci* 2012, **33**:97-121.
47. Gerwin S, Macke JH, Bethge M: **Bayesian inference for generalized linear models for spiking neurons.** *Front Comput Neurosci* 2010.
48. Paninski L, Ahmadian Y, Ferreira DG, Koyama S, Rad KR, Vidne M, Vogelstein J, Wu W: **A new look at state-space models for neural data.** *J Comput Neurosci* 2010, **29**:107-126.
49. Buesing L, Macke JH, Sahani M: **Spectral learning of linear dynamics from generalised-linear observations with application to neural population data.** *Advances in Neural Information Processing Systems* 2012:1682-1690.

50. Archer EW, Koster U, Pillow JW, Macke JH: **Low-dimensional models of neural population activity in sensory cortical circuits**. *Advances in Neural Information Processing Systems* 2014:343-351.
51. Macke JH, Buesing L, Sahani M: **Estimating state and parameters in state space models of spike trains**. *Advanced State Space Methods for Neural and Clinical Data* 2015:137.
52. Gao Y, Busing L, Shenoy KV, Cunningham John P: **High-dimensional neural spike train analysis with generalized count linear dynamical systems**. *Advances in Neural Information Processing Systems* 2015:2044-2052.
53. Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, Shenoy KV: **Single-trial dynamics of motor cortex and their applications to brain-machine interfaces**. *Nat Commun* 2015, **6**.
54. Yu Byron M, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M: **Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity**. *Advances in Neural Information Processing Systems* 2009:1881-1888.
55. Pfau D, Pneumatikakis EA, Paninski L: **Robust learning of low-dimensional dynamics from large neural ensembles**. *Advances in Neural Information Processing Systems* 2013:2391-2399.
56. Nam H: *Poisson Extension of Gaussian Process Factor Analysis for Modeling Spiking Neural Populations (Master's thesis)*. Tubingen: Department of Neural Computation and Behaviour, Max Planck Institute for Biological Cybernetics; 2015, 8.
57. Zhao Y, Park IM: **Variational latent Gaussian process for recovering single-trial dynamics from population spike trains**. *Neural Comput* 2017, **29**:1293-1316.
58. Duncker L, Sahani M: **Temporal alignment and latent gaussian process factor inference in population spike trains**. *Advances in Neural Information Processing Systems* 2018:10445-10455.
59. Zhao Y, Yates JL, Levi AJ, Huk AC, Park IM: **Stimulus-choice (mis) alignment in primate mt cortex**. *bioRxiv* 2019.
60. Wu A, Roy NA, Keeley S, Pillow JW: **Gaussian process based nonlinear latent structure discovery in multivariate spike train data**. *Advances in Neural Information Processing Systems* 2017:3496-3505.
61. Linderman S, Johnson M, Miller A, Adams R, Blei D, Paninski L: **Bayesian learning and inference in recurrent switching linear dynamical systems**. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)* 2017:914-922.
62. Nassar J, Linderman S, Bugallo M, Park IM: **Tree-structured recurrent switching linear dynamical systems for multi-scale modeling**. *International Conference on Learning Representations (ICLR)* 2019.
63. Linderman SW, Nichols ALA, Blei DM, Zimmer M, Paninski L: **Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in *C. elegans***. *bioRxiv* 2019:621540.
64. Zoltowski DM, Pillow JW, Linderman SW: **A general recurrent state space framework for modeling neural dynamics during decision-making**. *International Conference on Machine Learning* 2020. (pp. 11680-11691).
65. Glaser J, Whiteway M, Cunningham JP, Paninski L, Linderman S: **Recurrent switching Dynamical Systems Models for Multiple Interacting Neural Populations**. *Advances in Neural Information Processing Systems*, **33** (2020).
66. Archer E, Park IM, Buesing L, Cunningham J, Paninski L: *Black Box Variational Inference for State Space Models*. 2015arXiv:1511.07367.
67. Gao Y, Archer EW, Paninski L, Cunningham JP: **Linear dynamical neural population models through nonlinear embeddings**. *Advances in Neural Information Processing Systems* 2016:163-171.
68. Ainsworth SK, Foti NJ, Lee AKC, Fox EB: **oi-vae: output interpretable vaes for nonlinear group factor analysis**. *International Conference on Machine Learning* 2018:119-128.
69. Pandarinath C, O'Shea DJ, Collins J, Jozefowicz R, Stavisky SD, Kao JC, Trautmann EM, Kaufman MT, Ryu SI, Hochberg LR, Henderson JM, Shenoy KV, Abbott LF, Sussillo D: **Inferring single-trial neural population dynamics using sequential auto-encoders**. *Nat Methods* 2018, **15**:805-815.
70. Keeley S, Aoi M, Yu Y, Smith S, Pillow JW: **Identifying signal and noise structure in neural population activity with Gaussian process factor models**. *Advances in Neural Information Processing Systems*, **33** (2020).
 Find latent structure in multi-region recordings that is shared across trials (signal) and trial specific (noise). Using cross validation methods, the authors find that there is trail-by-trial latent structure that is shared across brain region. These results somewhat disagree with the findings of [22], though the data in [70**] are from rodent visual cortex.
71. Gokcen E, Semedo J, Zandvakili A, Machens C, Kohn A, Yu B: **Dissecting feedforward and feedback interactions between populations of neurons**. *Cosyne Abstracts; Denver, CO: 2020*
 Use a Gaussian Process latent variable model with temporal delays to understand inter-region communication in visual cortex. They identify distinct feedforward and feedback latent structures across visual cortical areas by associating these structures with particular delay values learned in the latent variable model. They call this model Delayed Latents Across Groups (DLAG).
72. Hotelling H: **Simplified calculation of principal components**. *Psychometrika* 1936, **1**:27-35.
73. Bach FR, Jordan MI: *A Probabilistic Interpretation of Canonical Correlation Analysis. Technical Report*. Berkeley: University of California; 2005.
74. Klami A, Virtanen S, Kaski S: **Bayesian canonical correlation analysis**. *J Mach Learn Res* 2013, **14**:965-1003.
75. Zhao S, Gao C, Mukherjee S, Engelhardt BE: **Bayesian group factor analysis with structured sparsity**. *J Mach Learn Res* 2016, **17**:6868-6914.
76. Kaufman MT, Churchland MM, Ryu SI, Shenoy KV: **Cortical activity in the null space: permitting preparation without movement**. *Nat Neurosci* 2014, **17**:440-448.
77. Aoi M, Pillow JW: **Model-based targeted dimensionality reduction for neuronal population data**. *Advances in Neural Information Processing Systems* 2018:6690-6699.
78. Aoi MC, Mante V, Pillow JW: **Prefrontal cortex exhibits multidimensional dynamic encoding during decision-making**. *Nat Neurosci* 2020. October.
79. Linderman S, Adams RP, Pillow JW: **Bayesian latent structure discovery from multi-neuron recordings**. *Advances in Neural Information Processing Systems* 2016:2002-2010.
80. Saxena S, Cunningham JP: **Towards the neural population doctrine**. *Curr Opin Neurobiol* 2019, **55**:103-111.
81. Semedo J, Amin Z, Adam K, Christian KM, Byron MY: **Extracting latent structure from multiple interacting neural populations**. *Adv Neural Inf Process Syst* 2014, **27**:2942-2950.
82. Cowley B, Semedo J, Zandvakili A, Smith M, Kohn A, Yu B: **Distance Covariance Analysis Proceedings of the 20th International Conference on Artificial Intelligence and Statistics**. *PMLR* 2017, **54**:242-251.