UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

Scalable Bayesian Functional Connectivity Inference for Multi-Electrode Array Recordings

Permalink

https://escholarship.org/uc/item/6m2244p7

Authors

Zhao, Yun Jiang, Richard Xu, Zhenni <u>et al.</u>

Publication Date

2020-07-04

Peer reviewed

Scalable Bayesian Functional Connectivity Inference for Multi-Electrode Array Recordings

Yun Zhao, Richard Jiang, Zhenni Xu, Elmer Guzman, Paul K. Hansma, Linda Petzold

University of California, Santa Barbara

Santa Barbara, CA

yunzhao@cs.ucsb.edu

ABSTRACT

Multi-electrode arrays (MEAs) can record extracellular action potentials (also known as 'spikes') from hundreds or thousands of neurons simultaneously. Inference of a functional network from a spike train is a fundamental and formidable computational task in neuroscience. With the advancement of MEA technology, it has become increasingly crucial to develop statistical tools for analyzing multiple neuronal activity as a network. In this paper, we propose a scalable Bayesian framework for inference of functional networks from MEA data. Our framework makes use of the hierarchical structure of networks of neurons. We split the large scale recordings into smaller local networks for network inference, which not only eases the computational burden from Bayesian sampling but also provides useful insights on regional connections in organoids and brains. We speed up the expensive Bayesian sampling process by using parallel computing. Experiments on both synthetic datasets and large-scale real-world MEA recordings show the effectiveness and efficiency of the scalable Bayesian framework. Inference of networks from controlled experiments exposing neural cultures to cadmium presents distinguishable results and further confirms the utility of our framework.

CCS CONCEPTS

• Applied computing → Life and medical sciences; Computational biology; Biological networks;

KEYWORDS

Functional connectivity; Multi-electrode array; Bayesian inference;

ACM Reference Format:

Yun Zhao, Richard Jiang, Zhenni Xu, Elmer Guzman, Paul K. Hansma, Linda Petzold. 2020. Scalable Bayesian Functional Connectivity Inference for Multi-Electrode Array Recordings. In San Diego '20: 19th International Workshop on Data Mining in Bioinformatics, August 24, 2020, San Diego, CA. ACM, New York, NY, USA, 10 pages.

1 INTRODUCTION

Neuroscience deals with how networks of neurons are organized and how they function [2]. Understanding connectivity between

San Diego '20, August 24, 2020, San Diego, CA

© 2020 Association for Computing Machinery.

neurons and within the brain is a fundamental problem in neurobiology [11]. Functional connectivity, defined as the statistical dependencies between different brain regions with similar patterns, is widely used in various neural tasks [1, 13]. For instance, functional connectivity magnetic resonance imaging (MRI) is crucial for diagnosing and comprehending autism spectrum disorders [20]. MEAs [24] can record extracellular action potentials from hundreds or thousands of neurons and provide insights on neuronal connectivity [17]. For hours or weeks, action potentials can be non-invasively monitored, when neurons are grown on planar MEAs [14]. Further, there is a trend towards increasing the density of the arrays [26] to better understand the neuron connectivities.

MEA recordings provide researchers opportunities to understand neuron activities in many regions such as the brain, retina, and heart [18]. However, the analysis of this data is challenging, in part because of its high dimensionality. Summary statistics could be used to measure the connection weights between electrodes. These include, for example, Pearson correlation [12], cross correlogram (CCG), the maximal information coefficient (MIC) [25] as well as biophysically-inspired metrics [4]. However, a data generative model is required to understand the underlying structure and to make full use of the domain expert knowledge [16]. Furthermore, these summary statistic methods provide different functional connectivity results for the same recording since they are all deterministic metrics, which present fixed connection weights between every two electrodes instead of a probabilistic estimation.

Bayesian inference can address the requirements for the inference, since it provides distributions for parameters using probabilistic models and observation data. In contrast to deterministic optimization procedures that give point estimates of the unknown functional connectivity, computing a Bayesian posterior yields probability distributions for the neuronal network functional connectivity. Bayesian inference has been combined with the generalized linear model (GLM), with graph-based priors to infer the neuron connectivity pattern for analysis [15]. However, there is a lack of scalable Bayesian techniques for inference of network structure, which is particularly acute for inference from high-density recordings.

A considerable challenge for Bayesian techniques is the rapid growth of computation time in accordance with the increasing scale of the network. In this paper, we propose a scalable framework of Bayesian inference, inspired by the hierarchical structure of networks of neurons. Experiments on both synthetic datasets and large scale real-world MEA recordings show that our framework provides accurate and insightful results. Furthermore, we apply the proposed framework to a controlled cadmium dataset, and the results confirm its utility.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

San Diego '20, August 24, 2020, San Diego, CA

Yun Zhao, Richard Jiang, Zhenni Xu, Elmer Guzman, Paul K. Hansma, Linda Petzold



Figure 1: The workflow for Bayesian functional connectivity inference. We consider only negative deflections that exceeded 6 times the standard deviation of the median noise level as spikes.

The key contributions of this paper include:

1) We propose a scalable functional connectivity inference framework shown in Fig. 1 for MEA recording data. We speed up the expensive Bayesian sampling process through the use of parallel computing.

2) We infer the network by splitting the large scale recordings into smaller local networks. The splitting strategy decreases the average sampling time quadratically in accordance with the number of smaller local networks. We also provide a strategy for inferring the regional connectivity between local networks. This not only eases the computational burden from sampling but also provides useful insights on regional connections in organoids and brains.

3) Experiments on both synthetic dataset and large-scale realworld MEA recordings show the effectiveness and efficiency of the Bayesian framework. Inference of network structure of Cadmiumexposed neuron cultures further demonstrates the usefulness of our framework.

The remainder of this paper is organized as follows. Section 2 describes the MEA data collection. We delineate the probabilistic models in Section 3 and demonstrate the Bayesian inference details in Section 4. Section 5 describes the hierarchical setup. Results for both synthetic and real data are provided in Sections 6 and 7, respectively. Related work is described in Section 8. Section 9 is the Discussion.

2 DATA COLLECTION

2.1 Cell Culture

We prepared hippocampal neurons from postnatal day 0 (P0) mice with C57BL/6 genetic background using a previously described protocol [28]. Cleaned and sterilized MEAs (120MEA100/30iR-ITO arrays; Multi Channel Systems) were incubated with poly-L-lysine (0.1 mg/ml) for at least one hour at 37 °C, rinsed 3 times with sterile deionized water and allowed to air dry before cell plating. Glial cultures were maintained in separate T-75 flasks. 100, 000 - 125, 000 dissociated glial cells were used for the first plating of MEAs to obtain a confluent glial culture over the surface of the electrodes. Once glia were confluent, the hippocampi dissected from the brain followed by manual dissociation were plated at 250, 000 cells in the MEA chamber. Cultures were grown in a tissue culture incubator (37 °C, 5% *CO*₂) in a medium made with minimum essential medium + EarleâĂŹs salts (Thermo Scientific, catalog # 11090081) with 2mM Glutamax (Thermo Scientific), 5% heat-inactivated fetal bovine serum (Thermo Scientific), and 1 ml/l Mito+ serum extender (Corning) and supplemented with glucose to an added concentration of 21mM. To minimize the effects of evaporation, maintain cell culture sterility, and decrease degassing of the medium during recordings, the MEA chamber was covered by a gas permeable membrane that permits exchange of CO_2 when the plate is in the CO_2 incubator.

2.2 MEA recordings

Extracellular voltage recordings of neuronal cultures were performed using an MEA 2100-System (Multichannel Systems, Reutlingen, Germany). Each MEA contained 120 electrodes with a 100 Âţm inter-electrode distance. All data were acquired at a 20 kHz sampling rate. All recordings were performed in culture media. The head stage temperature was set to 30ÂřC with an external temperature controller, and the MEAs were equilibrated for 5 min on the head stage before data acquisition or after any pharmacological or temperature manipulation. Recording duration was 3 minutes. Only cultures at 14 days in vitro (DIV) or older were used for pharmacological experiments.

2.3 Data Processing

Raw data was converted to HDF5 file format and processed offline. Spike detection was done with Matlab tools Waveclus [7]. Note that we did not apply spike sorting, since it may introduce considerable noise due to unsupervised clustering methods when trying to obtain the neuron (or unit) information [8], and there are many different spike sorting algorithms [7, 10], which give different outputs. Extracellular voltage recordings were bandpass filtered using cutoff frequencies of 200Hz and 4000Hz. Only negative deflections in the voltage records were labelled as spikes when the amplitude exceeded 6 times the standard deviation of the median noise level. Spike times and amplitudes were recorded and used for downstream analysis.

3 PROBABILISTIC MODEL

In this section, we briefly review the probabilistic model of neuronal spike trains introduced in [15] along with our choice of parameterization. Table 1 summarizes some common notations that we will use in this paper. At a high level, the model describes how the underlying connectivity network affects the activation propensity Scalable Bayesian Functional Connectivity Inference for Multi-Electrode Array Recordings

San Diego '20, August 24, 2020, San Diego, CA

Notations	Description
$X_{t,n}$	The observed spike at time bin t for electrode n
A	Adjacency matrix
W	Weight matrix
b_n	The baseline activation of electrode n
N	Number of electrodes
T	Autoregressive window of influence
$\psi_{t,n}$	The activation of electrode n at time bin t
No	Number of overlapped electrodes when split
W_o	Weight matrix of the overlapped region
N_s	Sample number of Bayesian inference
ho	The prior for connection probability
μ_{w_n}	Mean for the n th row of W
μ_b	Mean of the bias vector
S_{w_n}	Covariance for the n th row of W
Sb	Covariance of the bias vector

Table 1: Notations

of each electrode over time, producing the observed spike firing pattern measured over the entire MEA. Specifically, the model is composed of three parts: a network model specifying the underlying connectivity of the electrodes, an activation propensity model detailing how a network along with past spike history affects the probability of a spike at a time bin, and a spiking observation model mapping the activation propensity to the observed binary spike trains. Note that an electrode can fire no more than once in one time bin because of the refractory period in neurons. A probabilistic graphical model of this is shown in Fig 2.



Figure 2: Probabilistic graphical model for MEA. The model describes how the underlying connectivity of a network of electrodes (A and W) can lead to the observed spike trains $X_{t,n}$.

3.1 Network Model

The network model aims to capture the key properties of the underlying functional network of the electrode population. Specifically, it seeks to represent that those connections are potentially directional and different in strength. To accommodate this, a weighted directed graph is used where the edge weights represent the strength of the connection between two electrodes. This is incorporated as two matrix-valued latent variables $A \in \{0, 1\}^{N \times N}$ and $W \in \mathbb{R}^{N \times N}$ corresponding to a binary adjacency matrix and a real-valued weight matrix respectively.

3.2 Activation Propensity Model

A neuron can either fire spontaneously or as a response to communications (spikes) it receives from incoming, connected neurons. Given a particular realization of the electrode network, *A* and *W*, the instantaneous activation of electrode *n* at time bin *t*, $\psi_{t,n}$ is modeled as a linear, autoregressive function of the lagged spikes from neighboring electrodes:

$$\psi_{t,n} = b_n + \sum_{m=1}^N \sum_{\Delta t=1}^T A_{m \to n} W_{m \to n} e^{-\Delta t/\tau} X_{t - \Delta t, m}, \qquad (1)$$

Here, b_n represents the baseline activation rate for electrode nin the absence of influence from any other electrode. $A_{m \to n} \in \{0, 1\}$ is a binary variable indicating whether or not there exist directed connections from electrode m to electrode n. The weight $W_{m \to n}$ is the connection strength from electrode m to electrode n. The activation rate ψ is linearly adjusted by the lagged spikes from neighboring electrodes. The strength of the lagged spike is weighted by the strength of the connection to the neighbor and an exponentially decreasing function of time, inspired by the synapse connectivity measurement in [4], with time constant of $\tau = 15ms$. This prioritizes recent spikes from strongly connected neighbors. We consider both positive and negative $W_{m \to n}$, which captures that neuronal connections may be excitatory or inhibitory in nature. It is possible for a spike to decrease the propensity of firing when a weight is negative.

3.3 Bernoulli Observation Model

Our spike train data consists of binary observations of whether electrode *n* fired at time bin *t*, $X_{t,n}$. This is modeled as a Bernoulli random variable with probability dependent on the activation propensity, $\sigma(\psi_{t,n}) = e^{\psi_{t,n}}(1 + e^{\psi_{t,n}})^{-1}$, where σ is the logistic function that maps the propensity to a probability.

4 BAYESIAN INFERENCE

Based on the previous section, the full model, including the priors, is as follows:

Average Time for One Bayesian Sample (s)

$$A_{i,j} \sim \text{Bernoulli}(\rho)$$

{ μ_{w_n}, μ_b }, $S_{w_n}, S_b \sim \text{Normal-Inverse-Wishart}(0, 1, I, 3)$
 $w_n | \mu_{w_n}, S_{w_n} \sim \text{Normal}(\mu_{w_n}, S_{w_n})$
 $b | \mu_b, S_b \sim \text{Normal}(\mu_b, S_b)$
 $\psi_{t,n} = b_n + \sum_{m=1}^N \sum_{\Delta t=1}^T A_{m \to n} W_{m \to n} e^{-\Delta t/\tau} s_{t-\Delta t, m}$
 $X_{t,n} \sim \text{Bernoulli}(\sigma(\psi_{t,n})),$
(2)

where w_n denotes the *n*th row of the matrix *W*. The hyper-parameter ρ affects the prior over the connectivity matrix *A*.

We apply an efficient Gibbs sampler, which exhibits scalable parallelism, derived from [15]. Sampling the posterior over the discrete adjacency matrix *A* is the most challenging step. Due to conjugacy, we can integrate over *W* and sample *A* from its collapsed conditional distribution. We update *A* and *W* by collapsing out *W* to directly sample each of $A\hat{a}\tilde{A}\hat{Z}s$ elements. We iterate over each $A_{m\to n}$ and sample it from its conditional distribution. After that, we sample *W* from its conditional Gaussian likelihood.

We employ a novel splitting strategy to make the inference method amenable to high density arrays. We present and verify the methodology in the following sections.

5 SPLIT

We have 120 electrodes in the MEA device in this paper. However, there are 26,400 electrodes in the MaxWell complementary metal oxide semiconductor (CMOS) MEA device. As we can see in Fig. 3, the average time for one sample increases drastically with the size of the array. That is because the dimensions of parameters (A and W) to be estimated exhibit quadratic growth according to the number of electrodes. The time reported in Fig. 3 is the sampling time, computed in parallel with 24 CPU cores. To deal with the time complexity challenge, we propose a hierarchical inference procedure. As shown in Fig. 4, we split the whole large array into fixed number (for example, 4) of regions. In the first level of the algorithm, we perform Bayesian inference individually on each of the smaller regions. In the second level, we treat each whole region as one group by taking the mean of all the spike trains in the region. We apply the same probabilistic model to infer the regional connectivity, which is the connection strength between each pair of regions. This heuristic splitting strategy is inspired by the biological phenomenon of regional connections in brains. The splitting strategy decreases the average sampling time quadratically in accordance with the number of sub-regions. To get a better understanding of the connections on the border of any two regions, we further propose an overlapping split mechanism, as shown in Fig. 5.

6 RESULTS ON SYNTHETIC DATA

Due to the lack of ground truth functional connectivity in an MEA, we first use synthetic data to check the effectiveness of the probabilistic model and our splitting strategies. We use the ground truth model to generate synthetic data and compare the functional connectivity inferred from synthetic data with the ground truth. For all



Figure 3: Average time for one Bayesian sampling exhibits quadratic growth according to the number of electrodes. The time reported is the sampling time computed in parallel with 24 central processing units (CPU) cores.

	А	В	С	D	Ε	F	G	Н	J	К	L	Μ	
1				•	•	•	•	•	•				
2			•	•	•	•	•	•	•	•			
3		•	•	•	•	•	•	•	•	•	•		
4	•	•	•	•	•	•	•	•	•	•	•	•	
5	•	٠	•	٠	٠	٠	٠	٠	٠	٠	•	٠	
6	•	٠	•	•	٠	•	٠	٠	•	٠	•	•	
7	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	
8	٠	•	•	•	•	•	٠	٠	•	٠	•	•	
9	٠	٠	•	•	٠	•	٠	٠	•	٠	٠	٠	
10		٠	•	•	٠	•	٠	٠	•	٠	٠		
11			•	•	•	•	٠	٠	•	٠			
12				•	٠	•	٠	٠	•				

Figure 4: Non-overlapping split. We split the large array into four regions. For the first level, we infer individually on the smaller regions. For the second level, we treat each whole region as one hidden super node and infer the regional connection strength between each pair of sub-regions.

the experiments in this section and the next section, we use spike trains with shape of 180, 000 * 120. We apply parallel sampling with 24 CPUs. The Gibbs sampler was run for 1000 iterations. The first 500 samples were discarded to account for burn-in. We verify that the chain has reached a steady state by observing that the parameter traces and the log-likelihood have converged.

In Fig. 6, the inferred posterior mean is almost the same as the ground truth connectivity matrices. Cosine similarity measures



Figure 5: Overlapping split. To better capture the connections on the border of the regions, an overlapping split is used.

the similarity between two matrices of an inner product space, which is widely used in high dimensional spaces [19]. The inferred functional network A and W both achieved a high cosine similarity of 0.99 compared with ground truth in this case. Similarly in Fig. 7, we obtain cosine similarity of 0.95 and 0.99 for A and W when the number of electrodes increases to 10.

We also use synthetic datasets to verify the effectiveness of our splitting strategies. Here, we assume all the electrodes lie in a line. For the non-overlapping split, we use N electrodes and split them into two sub-regions, "front" and "back", with equal size and apply the Bayesian inference on each region. We compared the separate split results with the Bayesian inference results without split and reported the cosine similarity in Tab. 2. It is shown that the cosine similarities are consistently high with all different parameter settings, which validates the effectiveness of our splitting strategy. Similarly, for overlapping split, we test the results from split with N_o overlapped electrodes. In Tab. 3, N_o indicates the number of overlapped electrodes. W_o is the weight matrix for the overlapped region. High cosine similarities in Tab. 3 confirm the effectiveness of overlapping split strategy. From both Tab. 2 and Tab. 3, the network prior does affect the inference of *A* and *W* but in an indirect way, which is small in degree compared to the effect of the data.

Fig. 8 shows the effectiveness of regional connection inference after split. Each element in the inferred regional connectivity matrix *W* summarizes the elements of the corresponding regions in ground truth *W* altogether. Regional inferences after split precisely indicate the strength and the nature of the connections between regions.

7 RESULTS ON REAL DATA

In this section, we apply our framework onto two sets of real MEA recordings. Throughout this section, we apply non-overlapping split with 4 regions of equal size as in Fig. 4. As we mentioned in Section 2.2, we use 3-minute recordings, which present as spike



Figure 6: Comparison of network inferred from synthetic data with 4 electrodes and ground truth. Zero indicates no connection. Minus one indicates an inhibitory connection, and one indicates an excitatory connection. Cosine similarity between the two *As* is 0.99. Cosine similarity between the two *Ws* is 0.99.

Table 2: Non-overlapping Results: Cosine Similarity

Ν	S_w	μ_b	ρ	W_{front}	A_{front}	W_{back}	A_{back}
10	1	0	0.5	0.93	0.99	0.94	0.98
10	1	0	1	0.99	0.99	0.98	0.99
10	1	5	0.5	0.99	0.99	0.98	0.99
10	2	0	0.5	0.95	0.98	0.88	0.97
20	1	0	0.5	0.91	0.99	0.95	0.99
30	1	0	0.5	0.94	0.99	0.93	0.99

trains with shape of 180, 000 * 120. For the Bayesian inference, we adjusted the prior hyper-parameters and verified that the results stayed consistent.

7.1 Comparison with neuron experts' labeled ground truth

For the real dataset, since we don't have the ground truth, we adopted neuroscience experts' labelling as our reference. The neuroscience experts labelled those connections by watching MEA viewer [6]. We compare our Bayesian analysis with the expert labeled result. As is shown in Fig. 9, 93.33% of all the connections detected by neuron experts are detected by our inference. We can



Figure 7: Comparison of network inferred from synthetic data with 10 electrodes and ground truth. Cosine similarity between the two *As* is 0.95. Cosine similarity between the two *Ws* is 0.99.

Table 3:	Overlappi	ing Results:	Cosine	Similarity
		. /		

Ν	No	S_w	μ_b	ρ	W_{front}	A_{front}	W_{back}	A_{back}	W_o
10	2	1	0	0.5	0.97	0.9	0.95	0.99	0.98
10	2	1	0	1	0.99	0.99	0.99	0.99	0.99
10	2	1	5	0.5	0.99	0.99	0.99	0.99	0.99
10	4	2	0	0.5	0.92	0.98	0.95	0.99	0.93
20	4	1	0	0.5	0.94	0.99	0.97	0.99	0.96
30	4	1	0	0.5	0.90	0.99	0.93	0.99	0.97

see that the real recording has some regional patterns, which satisfies the intuition of our splitting strategy. Because of the difficulty in detecting those connections in neuroscience experiments, our neuroscience experts can find some connections that they are confident in but cannot guarantee to find all the connections. However, in our Bayesian inference, we gave the overall possible connections based on our probabilistic model. Thus we found more connections using Bayesian inference compared with the neuroscience experts' result. We also calculated the latency between electrodes using CCG based on the neuroscience expertsâĂŹ labeled ground truth connectivity. The latency shows a similar pattern as our inferred weights.



Figure 8: Regional connections in non-overlapping split. "Front" indicates the region composed of electrodes 0 and 1 in the ground truth figure, while "back" indicates the region of electrode 2 and 3. Each element in the inferred regional connectivity matrix W reflects the elements of the corresponding regions in ground truth W altogether. Regional inferences after split precisely reveal the overall connectivity between regions.

7.2 Cadmium concentration vs. controlled experiment

Cadmium is a toxic heavy metal that accumulates in living systems and is currently one of the most important occupational and environmental pollutants [5]. Cells have a calcium signalling toolkit whose components can be mixed and matched to create various spatial and temporal signals [3]. Cadmium can change the intracellular concentration of calcium, which is a universal and versatile intracellular messenger [9].

We reduced the release probability of presynaptic vesicles by titrating in increasing amounts of cadmium in order to decrease the influx of calcium current into the presynaptic terminal. Baseline recordings were performed 5 minutes prior to any cadmium addition. Following the baseline recording, the cadmium concentration was brought to $1\mu M$ and the solution was mixed gently. The MEA was placed on the recording headstage and allowed to equilibrate for 3 minutes, followed by a 3-minute recording. The same procedure was performed for $5\mu M$ cadmium in the same MEA.

To analyze how cadmium affects the structure of neuronal networks, we compared the obtained posterior matrices A and W for



Figure 9: Comparison between our Bayesian inference and ground truth detected by neuroscience experts. Green represents the connections found by Bayesian inference, while red denotes the connections not found. 93.33% of the connections detected by the neuroscience experts are detected by Bayesian inference.

a set of MEA recordings, comprised of one control with no cadmium introduced and 2 others with $1\mu M$ and $5\mu M$ of cadmium. We used a threshold of 0.05 (the value is chosen based on the outlier of the weight distribution) to eliminate those tiny weights in the inferred weight matrix, which is noise in biology. To understand the topology, we look at the following graph metrics for each posterior sample: number of connections, average clustering coefficient, and average path length. Overall, we find that Cadmium does change the topology of the estimated networks with statistical significance in posterior regions. As is shown in Fig. 10, the mean connection counts decrease with introduction of cadmium to the culture. Higher concentrations of cadmium lead to a drastic decrease in the number of connections in the functional networks.

Clustering coefficient is the number of edges between a electrodeâĂŹs immediate neighbors divided by all possible connections that could exist among them. It measures the level of local connectivity between electrodes. In Fig. 11, average clustering coefficient decreases with statistical significance when cadmium is in the culture, indicating that cadmium cultures are less likely to have connections within tightly connected groups or communities.

Average path length is the average shortest path length for all possible pairs of electrodes. Fig. 12 shows that the average path length increases significantly with the increase of cadmium concentration. That validates the idea that cadmium can impede neuron signal transmissions by changing the intracellular concentration of calcium.





Figure 10: Posterior of number of connections with Cadmium concentration. Cd is short for Cadmium. The mean connection counts decrease with introduction of cadmium to the culture. Higher concentrations of cadmium lead to a further decrease in the number of connections in the functional networks.

Distribution of Average Clustering Coefficient



Figure 11: Posterior of average clustering coefficient with Cadmium concentration. Average clustering coefficient decreases with statistical significance with the introduction of Cadmium into the culture.

In Fig. 13, we compare the indegree of each electrode for control and $5\mu M$ Cadmium concentration. We can see that the overall connection pattern maintains but the indegree decreases prominently with the introduction of Cadmium into the culture. Similar phenomenon can be detected for outdegree of each electrode for control and $5\mu M$ Cadmium concentration in Fig. 14.





8 RELATED WORK

MEAs provide opportunities for researchers to understand neuronal connectivity, by recording spike trains but also presents severe data analysis challenges. Inferring functional connectivity networks is critical to many applications in neuroscience. To analyze highdimensional spike trains, there exist several methods to measure connection weights. For example, CCG [12, 27] is a metric that was built from two electrodesâĂŹ spike trains. It presents the probability of an electrode firing to a spike in a τ milliseconds time window before or after another electrode fires. MIC has been widely used to identify known and novel relationships for data sets in gene expression and global health [25]. MIC is a two-variable dependent measure that captures functional relationship, by providing a score that roughly equals the coefficient of determination of the data relative to the regression function. Biophysically-inspired metrics extracts a directed functional connectivity matrix, based on the spike train [4]. It uses exponential decay property in axon potential signals to quantify the connection coefficients. However, despite the popularity of applying those metrics on MEA recordings, we can not rely on them to get reliable functional networks because they are deterministic and there is no model for the data.

GLM, a commonly used modeling framework [8, 23], can model the binary fire/not (1/0) and spike train recordings. We can use the *logit* link function to model the binomial distribution. GLMs are often fit by maximum a posteriori (MAP) estimation [21, 29]. However, when it comes to high-density recordings, MAP cannot fully convey the information in the posterior with a point estimate. Bayesian inference is a process using Bayes' theorem based on the prior beliefs about the probabilistic data generation process [22]. It updates the posterior for parameters in data generation model, as more data becomes available. GLM and graph-based models have been combined to infer the neuron connectivity patterns in a Bayesian way [15]. However, the Bayesian techniques have a



Figure 13: Incoming connections for control and Cadmium $5\mu M$. Nodes are sized and colored according to the number of incoming connections. The overall connection pattern remains but the indegree decreases prominently with the introduction of Cadmium into the culture.

considerable downside of the increased computation time, especially when the number of electrodes is large. In this paper, we introduced a scalable Bayesian inference framework on large scale MEA recordings. The functional connectivities are inferred from a joint probabilistic model of GLM and networking.

9 DISCUSSION

In this paper we have focused on inferring functional neuronal network connections using MEA recordings. Specifically, our goal has been to infer the functional connectivity for MEAs with large numbers of probes. Along this line, we proposed a scalable Bayesian inference framework. Our framework makes use of the hierarchical structure of networks of neurons and splits the whole array into



Cd 5µM Outdegree



Figure 14: Outgoing connections for control and Cadmium $5\mu M$. Nodes are sized and colored according to the number of outgoing connections. The overall connection pattern remains but the outdegree decreases prominently with the introduction of Cadmium into the culture.

smaller local networks for network inference. The splitting strategy decreases the average sampling time quadratically with the number of sub-regions. We also provide a strategy for inferring the connectivity between local networks. By comparing with ground truth for both synthetic data and real world human expert labeled MEA recordings, our experimental results provide compelling evidence that our framework can infer the underlying functional connectivity. Furthermore, we applied the proposed framework to a controlled cadmium dataset, and the results confirm its utility. As the density of the MEA continues to increase, our method will become more valuable to be able to infer the neuronal network structure efficiently.

Our experimental results demonstrate the usefulness of this framework. Here we suggest some avenues for future work. First,

the model could be extended to account for more detailed biological knowledge. For example, there are different types of neurons such as motor neurons and interneurons, which exhibit different types of behaviors in response to incoming signals. The auto-regressive propensity model, which is currently linear, can also be improved to incorporate nonlinear effects or other mechanistic firing models. Secondly, developing a method to determine when and how to split can allow for the framework to be used more robustly. Moreover, using the framework to study different factors instead of Cadmium, such as the presence of certain genes, may produce a greater understanding of the biologically complex systems.

REFERENCES

- Antonello Baldassarre, Christopher M Lewis, Giorgia Committeri, Abraham Z Snyder, Gian Luca Romani, and Maurizio Corbetta. 2012. Individual variability in functional connectivity predicts performance of a perceptual task. *Proceedings* of the National Academy of Sciences 109, 9 (2012), 3516–3521.
- [2] Mark F Bear, Barry W Connors, and Michael A Paradiso. 2007. Neuroscience. Vol. 2. Lippincott Williams & Wilkins.
- [3] Michael J Berridge, Peter Lipp, and Martin D Bootman. 2000. The versatility and universality of calcium signalling. *Nature reviews Molecular cell biology* 1, 1 (2000), 11–21.
- [4] Yazan N Billeh, Michael T Schaub, Costas A Anastassiou, Mauricio Barahona, and Christof Koch. 2014. Revealing cell assemblies at multiple levels of granularity. *Journal of neuroscience methods* 236 (2014), 92–106.
- [5] Jacopo Junio Valerio Branca, Gabriele Morucci, and Alessandra Pacini. 2018. Cadmium-induced neurotoxicity: still much ado. *Neural regeneration research* 13, 11 (2018), 1879.
- [6] Daniel C Bridges, Kenneth R Tovar, Bian Wu, Paul K Hansma, and Kenneth S Kosik. 2018. MEA Viewer: A high-performance interactive application for visualizing electrophysiological data. *PloS one* 13, 2 (2018), e0192477.
- [7] Fernando J Chaure, Hernan G Rey, and Rodrigo Quian Quiroga. 2018. A novel and fully automatic spike-sorting implementation with variable number of features. *Journal of neurophysiology* 120, 4 (2018), 1859–1871.
- [8] Zhe Chen. 2013. An overview of bayesian methods for neural spike train analysis. Computational intelligence and neuroscience 2013 (2013).
- [9] Grace Choong, Ying Liu, and Douglas M Templeton. 2014. Interplay of calcium and cadmium in mediating cadmium toxicity. *Chemico-biological interactions* 211 (2014), 54–65.
- [10] Jason E Chung, Jeremy F Magland, Alex H Barnett, Vanessa M Tolosa, Angela C Tooker, Kye Y Lee, Kedar G Shah, Sarah H Felix, Loren M Frank, and Leslie F Greengard. 2017. A fully automated approach to spike sorting. *Neuron* 95, 6 (2017), 1381–1394.
- [11] Karl J Friston. 1994. Functional and effective connectivity in neuroimaging: a synthesis. Human brain mapping 2, 1-2 (1994), 56–78.
- [12] Matteo Garofalo, Thierry Nieus, Paolo Massobrio, and Sergio Martinoia. 2009. Evaluation of the performance of information theory-based methods and crosscorrelation to estimate the functional connectivity in cortical networks. *PloS one* 4, 8 (2009), e6482.
- [13] Adam Gazzaley, Jesse Rissman, and Mark DàĂŹesposito. 2004. Functional connectivity during working memory maintenance. *Cognitive, Affective, & Behavioral Neuroscience* 4, 4 (2004), 580–599.
- [14] Christopher M Lewis, Conrado A Bosman, and Pascal Fries. 2015. Recording of brain activity across spatial scales. *Current opinion in neurobiology* 32 (2015), 68–77.
- [15] Scott Linderman, Ryan P Adams, and Jonathan W Pillow. 2016. Bayesian latent structure discovery from multi-neuron recordings. In Advances in neural information processing systems. 2002–2010.
- [16] Honglei Liu and Bian Wu. 2017. Active learning of functional networks from spike trains. In Proceedings of the 2017 SIAM International Conference on Data Mining. SIAM, 81–89.
- [17] Ming-Gang Liu, Xue-Feng Chen, Ting He, Zhen Li, and Jun Chen. 2012. Use of multi-electrode array recordings in studies of network synaptic plasticity in both time and space. *Neuroscience bulletin* 28, 4 (2012), 409–422.
- [18] Ming-Gang Liu, Xue-Feng Chen, Ting He, Zhen Li, and Jun Chen. 2012. Use of multi-electrode array recordings in studies of network synaptic plasticity in both time and space. *Neuroscience bulletin* 28, 4 (2012), 409–422.
- [19] Chunjie Luo, Jianfeng Zhan, Xiaohe Xue, Lei Wang, Rui Ren, and Qiang Yang. 2018. Cosine normalization: Using cosine similarity instead of dot product in neural networks. In *International Conference on Artificial Neural Networks*. Springer, 382–391.
- [20] Ralph-Axel Müller, Patricia Shih, Brandon Keehn, Janae R Deyoe, Kelly M Leyden, and Dinesh K Shukla. 2011. Underconnected, but how? A survey of functional

Yun Zhao, Richard Jiang, Zhenni Xu, Elmer Guzman, Paul K. Hansma, Linda Petzold

connectivity MRI studies in autism spectrum disorders. *Cerebral cortex* 21, 10 (2011), 2233–2243.

- [21] Liam Paninski. 2004. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems* 15, 4 (2004), 243–262.
- [22] Thomas Parr, Geraint Rees, and Karl J Friston. 2018. Computational neuropsychology and Bayesian inference. Frontiers in human neuroscience 12 (2018), 61.
- [23] Jonathan W Pillow, Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M Litke, EJ Chichilnisky, and Eero P Simoncelli. 2008. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454, 7207 (2008), 995–999.
- [24] Robert Prevedel, Young-Gyu Yoon, Maximilian Hoffmann, Nikita Pak, Gordon Wetzstein, Saul Kato, Tina Schrödel, Ramesh Raskar, Manuel Zimmer, Edward S Boyden, et al. 2014. Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy. *Nature methods* 11, 7 (2014), 727.
- [25] David N Reshef, Yakir A Reshef, Hilary K Finucane, Sharon R Grossman, Gilean McVean, Peter J Turnbaugh, Eric S Lander, Michael Mitzenmacher, and Pardis C Sabeti. 2011. Detecting novel associations in large data sets. *science* 334, 6062 (2011), 1518–1524.
- [26] Isabel Delgado Ruz and Simon R Schultz. 2014. Localising and classifying neurons from high density MEA recordings. *Journal of neuroscience methods* 233 (2014), 115–128.
- [27] Emilio Salinas and Terrence J Sejnowski. 2001. Correlated neuronal activity and the flow of neural information. Nature reviews neuroscience 2, 8 (2001), 539–550.
- [28] Kenneth R Tovar, Daniel C Bridges, Bian Wu, Connor Randall, Morgane Audouard, Jiwon Jang, Paul K Hansma, and Kenneth S Kosik. 2018. Action potential propagation recorded from single axonal arbors using multielectrode arrays. *Journal of neurophysiology* 120, 7 (2018), 306–320.
- [29] Wilson Truccolo, Uri T Eden, Matthew R Fellows, John P Donoghue, and Emery N Brown. 2005. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of neurophysiology* 93, 2 (2005), 1074–1089.

Table 4: Hyper-parameters for synthetic experiments

Т	μ_{w_n}	S_{b_n}		
100	1	1		

Table 5: Hyper-parameters for real data experiments

Т	μ_{w_n}	S_{b_n}	S_{w_n}	μ_b	ρ
100	1	1	1	-2	0.1

Appendix A APPENDIX FOR REPRODUCIBILITY

To support the reproducibility of the results in this paper, we provide the details in our experiments. We use 24 CPUs with model of Intel(R) Xeon(R) CPU E5-2670 v3 @ 2.30GHz. We perform parallel sampling with OpenMP and the average sampling time scales with the number of CPUs [15].

The network prior does affect the inference of A and W but in an indirect way, which is small compared to the effect of the data. For all the results in this paper, we ran the Gibbs sampler for 1000 iterations and verified that the results stayed consistent. Hyperparameters we used for synthetic and real data in Tab. 4 and Tab. 5 respectively.

A.1 Hyper-parameters for synthetic experiments

The hyper-parameters in Tab. 4 are fixed for synthetic experimental results in Fig. 6, Fig. 7, Fig. 8 and Tab. 2 and Tab. 3.

A.2 Hyper-parameters for real data experiments

The hyper-parameters we used for inference on real MEA recordings in Fig. 9, Fig. 10, Fig. 11, Fig. 12, Fig. 13 and Fig. 14 are listed in Tab. 5. We chose the hyper-parameters based on the choice in [15] and the observed firing probability in biological experiments.