Published in final edited form as:

Neural Comput. 2009 February ; 21(2): 450–477.

Identifying Functional Connectivity in Large-Scale Neural Ensemble Recordings: A Multiscale Data Mining Approach

Seif Eldawlatly,

Computer Science and Engineering, Michigan State University, East Lansing, MI 48824, U.S.A

Rong Jin, and

Neuroscience Program, Michigan State University, East Lansing, MI 48824, U.S.A

Karim G. Oweiss

Electrical and Computer Engineering, and Neuroscience Program, Michigan State University, East Lansing, MI 48824, U.S.A

Seif Eldawlatly: eldawlat@egr.msu.edu; Rong Jin: rongjin@msu.edu; Karim G. Oweiss: koweiss@msu.edu

Abstract

Identifying functional connectivity between neuronal elements is an essential first step toward understanding how the brain orchestrates information processing at the single-cell and population levels to carry out biological computations. This letter suggests a new approach to identify functional connectivity between neuronal elements from their simultaneously recorded spike trains. In particular, we identify clusters of neurons that exhibit functional interdependency over variable spatial and temporal patterns of interaction. We represent neurons as objects in a graph and connect them using arbitrarily defined similarity measures calculated across multiple timescales. We then use a probabilistic spectral clustering algorithm to cluster the neurons in the graph by solving a minimum graph cut optimization problem. Using point process theory to model population activity, we demonstrate the robustness of the approach in tracking a broad spectrum of neuronal interaction, from synchrony to rate co-modulation, by systematically varying the length of the firing history interval and the strength of the connecting synapses that govern the discharge pattern of each neuron. We also demonstrate how activity-dependent plasticity can be tracked and quantified in multiple network topologies built to mimic distinct behavioral contexts. We compare the performance to classical approaches to illustrate the substantial gain in performance.

1 Introduction

Monitoring the activity of neurons adjacent to a recording electrode is an essential first step in understanding the computational principles of biological neural systems. As tools for simultaneous recording of multiple neuronal ensembles continue to excel (Normann, Maynard, Rousche, & Warren, 1999; Wise, Anderson, Hetke, Kipke, & Najafi, 2004), identifying neuronal elements that collectively and cooperatively respond to external stimuli becomes of utmost importance and can substantially help recognize the topology of neural circuits that may be momentarily configured to process and store information, particularly during learning and behavior.

Extensive research in the past few decades has discovered functional relationships between simultaneously recorded neurons from their spike train activity. Existing techniques, however, are limited to resolving the correlation of very few recorded neurons (Gerstein & Clark, 1964; Gerstein & Perkel, 1969; Abeles & Goldstein, 1977; Abeles & Gerstein, 1988; Rosenberg, Amjad, Breeze, Brillinger, & Halliday, 1989; Brillinger, 1992; Martignou, Deco, Laskey, Diamond, Freiwald, & Vaadia, 2000; Jarvis & Mitra, 2001; Lindsey & Gerstein,

2006). In addition, a common assumption among all of these techniques is that the interval over the functional relationship between any neuron pair remains fixed for both neurons. While this may be the case in local populations, such a simplistic assumption does not take into account multiple factors governing the discharge probability of a neuron. This probability can vary considerably over short latencies of a few milliseconds depending on the strength of synaptic connectivity to other adjacent neurons (Maass & Natschlager, 2000). It can also vary over relatively longer latencies, possibly hundreds of milliseconds (Fellous, Tiesinga, Thomas, & Sejnowski, 2004), depending on the selective activation of the highly distributed neuronal elements that continuously interact to simultaneously encode multiple behavioral parameters.

Many recent studies on spike-timing-dependent plasticity (STDP) have also demonstrated that neurochemical receptors can generate excitatory post synaptic potentials (EPSPs) at resting potential that can last a few hundred milliseconds, which elongates the time window required for synaptic integration and eventually alters the temporal length for neuronal excitability (Markram, Lubke, Frotscher, & Sakmann, 1997; Tsodyks & Markram, 1997; Bi & Poo, 1998; Roberts & Bell, 2002; Destexhe & Marder, 2004; Dzakpasu & Zochowski, 2005; Frerking, Schulte, Wiebe, & Stäubli, 2005). In turn, this may trigger slow modulation of the discharge probability of distal neurons over much longer latencies. More recently, information content in spike trains was shown to vary considerably depending on the timescale chosen for the analysis, and, depending on a particular choice of that timescale, strongly correlated network states can result even though correlations over a fixed bin width may be statistically insignificant (Averbeck, Latham, & Pouget, 2006; Schneidman, Berry, Segev, & Bialek, 2006). The implications of these findings seem to reinforce Hebb's original hypothesis that each neuron can participate in different cell assemblies at different times (Hebb, 1949), indicating variable degrees of involvement in encoding behavioral parameters that is a function of the behavioral state, the particular brain area from which they are recorded, and the subject's level of attention and experience with the behavioral task. Thus, there is an urgent need to assess the interaction between simultaneously recorded neuronal elements over a multitude of timescales to better understand the nature of neural computations that give rise to behavioral events.

The objective of this letter is to introduce a new approach for identifying functionally interdependent neurons regardless of the interval length over which this functional relationship exists. The approach is specifically tailored to large-scale ensemble recordings obtained across multiple cortical areas where identifying the dynamics of neural circuits is highly desirable. These dynamics are governed by many temporal and spatial variations in neuronal interaction that form the core of a highly plastic cortical information processing mechanism.

2 Theory

Suppose that in a P neuron population, we want to approximate the probability of the p th neuron firing within an interval of length T . If the interval is binned to a very small bin width Δ chosen to account for the refractory period (i.e., the digital form of the spike train can accommodate consecutive 1 bins, and no more than one event can occur in any given bin), then the neural discharge pattern during T will be a vector of length $N = T/\Delta$, and the spikes will occupy some fraction N_p/N of these N bins, where N_p denotes the number of events within T . This simple calculation gives an empirical estimate of the firing rate of the neuron within T . Alternatively, if $T = [0, t)$, this probability can also be calculated with knowledge of the conditional intensity function $\lambda_p(t | H_p(t))$, where $H_p(t)$ denotes the history of all the processes that affect the firing probability of the neuron p up to time t (e.g., stimulus feature, other neurons' interaction, the neuron's own firing history). In this model, considering the history, $H_p(t)$ models the effects caused by postsynaptic potentials. Specifically, the probability of a spike event for small Δ is (Brillinger, 1975)

$$\Pr\{\text{spike from neuron } p \text{ in } (t, t+\Delta] | H_p(t)\} \cong \lambda_p(t)H_p(t)\Delta. \quad (2.1)$$

The spike train s_p can be modeled as a conditionally independent Bernoulli random process with probability mass function (Brown, 2005):

$$f(s_p(t)) = \begin{cases} \lambda_p(t)H_p(t)\Delta, & \text{if } s_p(t)=1 \\ 1 - \lambda_p(t)H_p(t)\Delta, & \text{if } s_p(t)=0. \end{cases} \quad (2.2)$$

Here the success probability of the Bernoulli random variable is time varying, given its dependence on the history interval. This interval is likely to be of variable duration when the intrinsic properties of the synapses undergo substantial changes, for example, as a result of Hebbian learning.

3 Two-Neuron Model

Assume that we have two spike trains from two neurons p and q , both binned with the same Δ , and we are interested in the joint probability of firing at two time instants t_i and t_j , where $\{t_i, t_j\} \in [0, t)$, $t_i \neq t_j$. In this case, the joint probability is expressed as

$$\Pr\{\text{spike from neuron } p \text{ in } [t_i, t_i+\Delta), \text{ spike from neuron } q \text{ in } [t_j, t_j+\Delta) | H_p(t_i), H_q(t_j)\} = \rho_{pq}(t_i, t_j)\Delta^2, \quad (3.1)$$

where $\rho_{pq}(t_i, t_j)$ represents the second-order cross-product density between neurons p and q . This density can be estimated from knowledge of the cross-covariance density $\sigma_{pq}(t_i, t_j)$ as (Brillinger, 1975)

$$\sigma_{pq}(t_i, t_j) = \rho_{pq}(t_i, t_j) - \lambda_p(t_i)\lambda_q(t_j), \quad (3.2)$$

where $\lambda_p(t_i)$ and $\lambda_q(t_j)$ are the mean firing intensities of neurons p and q , respectively.

As $|t_i - t_j| \rightarrow \infty$, $\sigma_{pq}(t_i, t_j) \rightarrow 0$, the processes p and q become independent. However, when $|t_i - t_j|$ is finite and the points in both processes are isolated (i.e., describe orderly point processes; Brillinger, 1976), we can replace the joint probability in equation 3.1 with expectations and substitute in equation 3.2:

$$\sigma_{pq}(t_i, t_j) = E[s_p(t_i)s_q(t_j)]/\Delta^2 - \lambda_p(t_i)\lambda_q(t_j). \quad (3.3)$$

Using equation 2.2, the cross-covariance can be completely expressed in terms of the mean and conditional intensity functions of the processes p and q as

$$\sigma_{pq}(t_i, t_j) = \lambda_p(t_i)H_p(t_i)\lambda_q(t_j)H_q(t_j) - \lambda_p(t_i)\lambda_q(t_j), \quad (3.4)$$

where we have used the fact that the expectation in equation 3.3 reduces to the product of the conditional intensity functions given that they are independent conditioned on their joint firing history. This includes the firing history of neuron p as well as that of neuron q . Equation 3.4

implies that the cross-covariance between two neurons at instants t_i and t_j depends not only on the instantaneous probability of firing of the two neurons, but also on their history of firing. The dependence on the precise times t_i and t_j , and not on $|t_i - t_j|$, reflects the nonstationarity inherent in this correlation that results from biophysical mechanisms widely believed to underlie synaptic plasticity (Froemke & Dan, 2002).

Consider, for example, a modified version of the conditional intensity of an inhomogeneous Poisson model (Brillinger, 1975),

$$\lambda_p(t_i|\alpha_p, H_p(t_i)) = \exp \left(\beta_p + \sum_{q=1}^P \sum_{m=1}^{M_{pq}} \alpha_{pq}(m\Delta) I_q(t_i - m\Delta) \right), \quad (3.5)$$

where α_{pq} denotes the synaptic coupling between neurons p and q in the m th time bin; M_{pq} is the length (in bins) of the history interval that relates neuron p 's firing probability to activity from neuron q and can vary as q varies; $I_q(t_i - m\Delta)$ is an indicator function that takes the value of 1 if neuron q fired a spike in the m th time bin and 0 otherwise; and β_p is the log of the background rate of neuron p . The α_{pq} 's are time-varying weight factors governing how much the occurrence of a spike from neuron q , going back M_{pq} bins in time, influences the probability of firing of neuron p at time t_i . The model in equation 3.5 has been previously demonstrated to be powerful in fitting data from the motor cortex during goal-directed arm reach movements (Truccolo, Eden, Fellows, Donoghue, & Brown, 2005; Srinivasan, Eden, Willsky, & Brown, 2006), as well as data from place cells in rat hippocampus area CA1 during foraging in a linear track (Brown, Nguyen, Frank, Wilson, & Solo, 2001).

A related model in Okatan, Wilson, and Brown (2005) was restricted to a fixed length history interval among all neurons in the ensemble. The model in equation 3.5 is more realistic for two reasons: first, it takes into account the variable length of the history interval among any given neuron pair consistent with the expected variations in the postsynaptic potentials duration depending on the cell morphology, and, second, it accounts for temporal precision in spike occurrence when Δ is too small (e.g., 1 ms), which is important for identifying STDP in synaptic connectivity, as will be detailed later.

Consider the weight function α_{pq} given by

$$\alpha_{pq}(t) = A_{pq} \sin(F\pi t/M_{pq}) \exp(-Bt/M_{pq}), \quad (3.6)$$

where F and B are constants and A_{pq} is a magnitude controlled by synaptic modification. The damped sinusoidal form in this expression implies that a spike from a presynaptic neuron occurring toward the end of the history interval has a smaller contribution to the firing probability of the postsynaptic neuron than if it were to occur at an earlier time. This weight function mimics the change in the membrane potential of postsynaptic neurons caused by an EPSP when A_{pq} is positive or an inhibitory postsynaptic potential (IPSP) when A_{pq} is negative. The damped sinusoidal form of α_{pq} models a synchronized oscillatory impulse response of the effect of potentially many unobserved bidirectional, excitatory, and inhibitory neurons that affect the discharge pattern of neurons p and q . The time-varying characteristics of α_{pq} are controlled by the sinusoid frequency and the time constant of the exponential term, which is equivalent to the bandpass filtering effect of synaptic transmission (Armstrong-Gold & Rieke, 2003). Other functions can also be used to mimic different characteristics. A sample function is shown in Figure 1 for two history intervals while other constants are fixed. A narrow

timescale will be suitable to track the dynamics of α_{pq} when τ_{pq} is small, and a larger timescale will be needed if τ_{pq} is large. It is very likely that the interaction history intervals between multiple neurons in an ensemble will be distinct, given the possibility of a neuron to participate simultaneously in multiple neuronal clusters consistent with Hebbian learning.

4 Population Model

We assume that each of the P neurons may belong to any of K clusters. A cluster is defined as a group of neuronal elements that exhibit all types of functional interdependency within an analysis interval of length T . This definition is consistent with the notion of a neural assembly originally proposed by Hebb (1949). In our context, it is assumed that within a given cluster, the probability of a spike occurring in any given bin is statistically dependent on the activity of the individual neurons in that cluster yet independent of the activity of neurons outside that cluster. Therefore, the dependence of the correlation on the history of firing within a cluster as derived in equation 3.4 plays an important role in the clustering step of the algorithm. Typically, knowledge about the length of these history terms is absent, as well as the degree of participation of any given neuron in the cluster. We show below that the proposed approach can identify these functional relationships without knowledge of these parameters.

Let us reconsider model 3.5. If neuron p belongs to cluster k , then all the α_{pl} 's for $l \notin \{k\}$ are essentially zero. Therefore, the conditional intensity function is expressed as

$$\lambda_p(t_i|\alpha_p, H_p(t_i)) = \exp \left(\beta_p + \sum_{l \in \{k\}} \sum_{m=1}^{M_{pl}} \alpha_{pl}(m\Delta) I_l(t_i - m\Delta) \right). \quad (4.1)$$

If the background rate β_p is stationary (i.e., $\lambda_p(t_i) = \lambda_p = \exp(\beta_p)$), then the normalized cross-covariance takes the form

$$\begin{aligned} \sigma'_{pq}(t_i, t_j) = & \exp \left(\sum_{l \in \{k\}} \sum_{m=1}^{M_{pl}} \alpha_{pl}(m\Delta) I_l(t_i - m\Delta) \right. \\ & \left. + \sum_{l \in \{k\}} \sum_{m=1}^{M_{ql}} \alpha_{ql}(m\Delta) I_l(t_j - m\Delta) \right) - 1, \end{aligned} \quad (4.2)$$

where we have normalized by $\exp(\beta_p + \beta_q)$ for convenience. Note here that we consider the more general case where $\alpha_{pl} \neq \alpha_{lp}$ and $M_{pl} \neq M_{lp}$. This accounts for variable interaction lengths and connection strengths depending on the pre- and postsynaptic activation expressed by α_{pl} and α_{lp} . Equation 4.2 reveals that for the case where $l = q$ in the first term and $l = p$ in the second term (i.e., cluster k consists of neurons p and q only), the correlation coefficient will be essentially zero if neuron p excites or inhibits neuron q with the same degree that neuron q inhibits or excites neuron p over the same interval.

This analysis emphasizes that correlations are highly dependent on the timescale that governs the functional connectivity among neurons, which depends on the strengths and directions of the synaptic coupling, as well as the particular elements of a given cluster. In an actual situation, specific information about these elements and the associated history is not available to the experimenter, and therefore functional interdependence between any neuron pair has to be identified without knowledge of these parameters.

5 Scale Space Representation

The dependence of the correlation on the variable-duration history terms necessitates a preprocessing step to identify any intricate variations in neuronal discharge probability across multiple timescales. This step consists of projecting the spike train onto a scale space using a variable-duration Haar rectangular basis (Mallat, 1998). At any given timescale j , the spike train can be expressed as

$$s_p^j(t) = \mathbf{W}^{(j)} s_p(t) \quad j=0, 1, \dots, J, \quad (5.1)$$

where $\mathbf{W}^{(j)}$ is a lumped matrix representing the Haar wavelet transform and J is the total number of timescales (i.e., depth of temporal analysis). This is a parameter to be chosen by the user. Figure 2 shows an example of a scale space representation of a spike train snippet up to two timescales. The Haar transformation operator $\mathbf{W}^{(j)}$ for the first two timescales takes the form

$$\mathbf{W}^{(1)} = \sqrt{2} \begin{bmatrix} 1/2 & 1/2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1/2 & 1/2 \\ 1/2 & -1/2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1/2 & -1/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/2 & -1/2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1/2 & -1/2 \end{bmatrix}$$

$$\mathbf{W}^{(2)} = 2 \begin{bmatrix} 1/4 & 1/4 & 1/4 & 1/4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/4 & 1/4 & 1/4 & 1/4 \\ 1/4 & 1/4 & -1/4 & -1/4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/4 & 1/4 & -1/4 & -1/4 \end{bmatrix}, \quad (5.2)$$

where multiplying the top half of $\mathbf{W}^{(1)}$ by the spike train s_p gives the probability of firing in successive bins that are two samples in size. These represent the approximation coefficients a_p^1 . Multiplying the bottom half of $\mathbf{W}^{(1)}$ by the spike train gives the relative change in the firing pattern in successive bins that are two samples in size. These represent the detail coefficients d_p^1 . Similarly, a_p^2 and d_p^2 can be obtained by using $\mathbf{W}^{(2)}$, which is equivalent to computing the probability of firing and the relative change in successive bins that are four samples in size, and so on.

Two important features can be extracted from the scale space representation: the probability of firing at timescale j indicated by the sequence $2^{-j/2} a_p^j(t)$ along the transform approximation path (left branch), and the relative change in firing probability across successive bins indicated by the sequence $2^{-j/2} d_p^j(t)$ along the details path (right branch), where positive or negative values indicate a decrease or increase in the probability of firing relative to the previous bin, respectively. At this stage, the scale space representation incorporates all the desired information about the dynamics of the discharge patterns.

6 Similarity Measures

Once the scale space representation is obtained for every spike train, a similarity measure is computed between pairwise neurons in each timescale up to a temporal depth set by the user. The choice of this depth largely depends on the expected latency of interaction between neurons, which typically depends on the brain regions from which these neurons are recorded.

The similarity measure used throughout this letter is the sample cross-correlation between neurons p and q at timescale j and can be expressed as

$$\sigma_{pq}^j(t_i, t_j) = E[s_p^j(t_i)s_q^j(t_j)]. \quad (6.1)$$

The full correlation matrix at timescale j , $\Sigma^{(j)} \in \mathbb{R}^{P \times P}$, can be formed using equation 6.1 for all $p, q \in \{1, \dots, P\}$. Other similarity measures can also be used (e.g., mutual information, coherence, Granger causality).

We now define the total instantaneous energy of a given cluster, $C_k \in \mathbb{R}^{P \times P}$, as the algebraic sum of all instantaneous pairwise correlations between the elements of that cluster, that is,

$C_k = \sum_j \sum_{p, q \in \{k\}} \sigma_{pq}^j$, where the time indices have been dropped for notational simplicity. We note here that all the off-diagonal entries $C_k(p, q)$ are zero for $p, q \notin \{k\}$. The objective is to determine the nonzero entries $C_k(p, q)$, $\forall p, q \in \{k\}$ that can serve as similarity measures between pairwise neurons, regardless of the length of the history interval governing the interaction between the two neurons. These entries can be obtained by simultaneously diagonalizing the correlation matrices $\Sigma^{(j)}$ computed from equation 6.1 for $j = 0, 1, \dots, J$. Specifically, a block diagonal matrix $R \in \mathbb{R}^{(P \times J) \times (P \times J)}$ that has the j th $P \times P$ block as $\Sigma^{(j)}$ is formed. This matrix is diagonalized using singular value decomposition (SVD) as

$$R = \begin{bmatrix} \Sigma^{(0)} & & & \\ & \Sigma^{(1)} & & \\ & & \ddots & \\ & & & \Sigma^{(J)} \end{bmatrix} \cong \sum_{i=1}^Q \lambda_i \mathbf{u}_i \mathbf{u}_i^T, \quad (6.2)$$

where $\lambda_i/\mathbf{u}_i$ denote the eigenvalue-eigenvector pair associated with the i th dominant mode of the block diagonal matrix R .

To compute an instantaneous similarity measure between pairwise neurons, the dominant Q modes of the decomposition in equation 6.2 are used to approximate the augmented correlation matrix R . In practice, this is achieved by first vectorizing the J matrices $\Sigma^{(j)}$ to $P^2 \times 1$ -dimensional vectors, concatenating the J vectors obtained in a matrix of dimension $P^2 \times J$, performing SVD on the resulting matrix and keeping the first Q modes in a matrix of size $P^2 \times Q$, and then finally reshaping the $P^2 \times Q$ matrix to Q matrices of size $P \times P$. SVD is used here as an efficient way to fuse the correlation matrices obtained at different timescales into one similarity matrix while taking into account the expected variance in neuronal temporal interactions. SVD captures second-order statistics while simultaneously decorrelating the information across multiple timescales, thereby eliminating any redundancy in the cross-correlations obtained.

7 Spectral Clustering

The next step relies on clustering a graphical representation of the population. Graph theory is used, and each neuron is represented as an object. Each object p —in our case, a neuronal spike train—is represented as a vertex in the graph. Any two objects, p and q , are connected by a directed edge whose weight w_{pq} is the similarity between the two objects. An example of an undirected graph is shown in Figure 3. In our case, the weight w_{pq} is the sum of the (p, q) th entry of the Q principal matrices weighted by the corresponding eigenvalue obtained from equation 6.2.

The task here is to cluster the neurons in this graphical representation by identifying the groups of strongly similar objects. In our case, we use a spectral clustering approach drawn from machine learning theory (Chung, 1994; Ding, He, Zha, Gu, & Simon, 2001). Briefly, a spectral clustering algorithm performs graph partitioning using a minimum cut algorithm. The connected graph is divided into multiple disconnected subgraphs such that only edges of small weights are removed. In the example shown in Figure 3, by removing the edges with small weights, we are able to identify two clusters of strongly correlated objects, with objects 1 and 2 in the first cluster and objects 3 and 4 in the second cluster. Since a typical graph cut problem is NP-hard, some approximation methods exist to efficiently solve the problem using “soft” memberships to reduce the computational complexity (Jin, Ding, & Kang, 2005).

We denote by a_{pk} the probability of the p th neuron to be in the k th cluster and solve the minimum cut problem by finding the set of probabilities $\{a_{pk}\}$ that maximizes the following objective function:

$$l = \sum_{k=1}^K \frac{\sum_{p=1}^P \sum_{q=1}^P a_{pk} a_{qk} w_{pq}}{\sum_{p=1}^P \sum_{q=1}^P a_{pk} w_{pq}} \quad (7.1)$$

The final cluster memberships are derived from $\{a_{pk}\}$ by assigning each object to the cluster with the largest probability: $k_p^* = \arg \max_{k \in \{1, \dots, K\}} a_{pk}$. Note that although many clustering algorithms exist, a number of studies have shown that spectral clustering is able to outperform many well-known clustering algorithms, particularly when cluster boundaries are complex to define (Ng, Jordan, & Weiss, 2001). Our results demonstrate that this is indeed the case.

Here are the steps of the entire algorithm:

Pseudocode of the Proposed Method

- For $j = 1$ to J (timescale index)
 - For $p = 1$ to P (neuron index)
 - Obtain the wavelet representation s_p^j of each spike train s_p
 - End
 - Compute the $P \times P$ correlation matrix $\Sigma^{(j)}$ of the spike trains at timescale j
 - End
- Form the block diagonal correlation matrix R from the matrices $\Sigma^{(j)}$ and obtain its spectral decomposition using SVD
- Obtain the pairwise similarity of neurons using the weighted sum of the first few dominant modes of R
- Compute the cluster membership of each neuron by maximizing the objective function 7.1

8 Results

At this stage of development, it was essential to test the algorithm on spike train data for which we know the absolute truth. Therefore, we extensively tested the algorithm on data simulated using the point process model in equation 3.5 with known relationships among neuronal elements. We also propose metrics to assess performance when the ground truth is not available when analyzing real data.

8.1 Clustering Synaptically Coupled Populations

For comparison, we used synaptic coupling in equation according to the following expressions (Okatan et al., 2005):

$$\begin{aligned}\alpha_{pp}^0(t) &= -2\sin(3000\pi t/M_{pp})\exp(-3000t/M_{pp}), \\ \alpha_{pq}^+(t) &= 2.5\sin(4000\pi t/M_{pq})\exp(-3000t/M_{pq}), \\ \alpha_{pq}^-(t) &= -3\sin(2000\pi t/M_{pq})\exp(-3000t/M_{pq}),\end{aligned}\quad (8.1)$$

where t is in seconds, α_{pp}^0 and $\alpha_{pq}^\pm$ are the autoinhibitory and excitatory or cross-inhibitory interactions, respectively. We generated data sets ($n = 25$) containing 16 neurons, each with a total duration of 98 seconds in each trial. In each data set, the 16 neurons were divided into four clusters with 4 neurons each.

8.1.1 Population with Fixed History Interval—Figure 4a shows the network topology of the 16-neuron population with no across-cluster connectivity. Within a cluster, each neuron inhibits itself as well as its neighbor in the clockwise direction, while exciting the other neighbor in the counterclockwise direction. In this data set, the parameter M_{pq} was set to 120 for all interactions. Figure 4b shows a 3-second raster of the generated spike trains. A sample interspike interval (ISI) histogram of a neuron in the population and that of an independent Poisson model for two choices of M_{pq} (4 and 120 bins) is shown in Figure 4c, demonstrating the effect of inhibition (first 40 ms) and the excitation (40–100 ms) intervals resulting from interaction with other neurons when M_{pq} is large, while this interaction effect diminishes when M_{pq} decreases. We also computed the coefficient of variation (CV) of the ISI histograms ($CV = \text{std}(\text{ISI})/\text{mean}(\text{ISI})$) to compare the statistics of the data generated by the model to those typically observed in cortical neurons (Shadlen & Newsome, 1998; Stevens & Zador, 1998). The right-most plot in Figure 4c illustrates that as the history interval increases, the CV decreases below 1 ($CV = 1$ is equivalent to a Poisson neuron), consistent with characteristics of many types of cortical neurons (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997).

In each trial, four across-cluster interactions pairs were randomly added to simulate noise or potential error margins that may be encountered, for example, during spike sorting of simultaneously recorded neurons (Lewicki, 1998). Each pair consisted of an excitation in one direction and an inhibition in the other using the same expressions in equation 8.10. Figure 4d shows the augmented R matrix obtained by considering correlation matrices up to a timescale of 96 ms, illustrating the computed similarity between pairwise neurons averaged across trials. Neurons were ordered such that neurons 1 to 4 represent cluster 1, neurons 5 to 8 represent cluster 2, and so on. Neurons in the same cluster are observed to share significantly larger similarity than those across clusters. The average clustering accuracy using the proposed approach shown in Figure 4e (the solid curve) was above 96% for timescales in the range of approximately 200 ms to 1.5 seconds and peaks around 360 ms, consistent with the fixed history interval chosen in the model ($M_{pq} \Delta = 360$ ms); then it declines below 80% for smaller (below 10 ms) as well as larger timescales (more than 10 sec). Since we did not account for synchrony explicitly in this data set, these results confirm our expectation that the clustering accuracy is maximized at temporal resolution matching the network properties.

To assess the effect of each of the computational steps in the algorithm on the clustering accuracy, we first examined the contribution of the scale space step. When the similarity matrix was computed using a fixed bin width (i.e., without the scale space representation prior to clustering), an average accuracy of only about 51% was achieved. We then compared the performance when the k -means was used as opposed to spectral clustering. Figure 4e shows

that the clustering accuracy deteriorates (maximum accuracy in this case is 61%) compared to the case when spectral clustering is used. Therefore, we concluded that both steps are important for superior performance.

8.1.2 Population with Variable History Interval—The interaction history interval with neighboring neurons plays a fundamental role on the probability of firing of the neuron under consideration. For example, if the time constant of an excitatory synapse between neurons p and q is small (e.g., $M_{pq} = 4$ bins, as in Figure 4c), neuron p is likely to fire M_{pq} times the likelihood of firing independently. In contrast, when the time constant is large (e.g., $M_{pq} = 120$ bins), neuron p incorporates the history of neuron q 's firing over a much longer timescale and therefore ignores the precise timing of neuron q 's firing and emphasizes firing rate information. Thus, the former case represents synchrony effects governing the probability of firing, while the latter represents the response to rate modulation effects, both believed to be major constituents of the neural code (Sjöström, Turrigiano, & Nelson, 2001; Nelson, Sjöström, & Turrigiano, 2002).

We varied the history parameter M_{pq} and examined the algorithm performance as illustrated in Figure 5a. The performance systematically peaked around the specific history interval chosen in the model. This is indicated by observing the monotonic relationship in Figure 5b between the timescale where maximum accuracy was attained and the history interval in the population model. In our model, the postsynaptic neuron firing pattern is a highly nonlinear function of the individual presynaptic neurons' firing pattern as indicated by the exponential term and the nonlinear relation between M_{pq} and the synaptic strength in equation 3.6. Therefore, we did not expect the relationship between the timescale of maximum accuracy and the history interval to be linear. It is also noteworthy that we are not interested in estimating the history intervals of interaction, but rather identify the clusters of functionally interdependent neurons.

8.1.3 Clusters with Distinct Temporal Interaction—We further examined the effect of varying history intervals across clusters, while maintaining it fixed among neurons within a given cluster. We used the same topology in Figure 4a with $M^1 = M^2 = 20$, and we varied both M^3 and M^4 , where M^k denotes the history interval (in bins) of interaction between neurons in cluster k . Figure 6a shows the clustering accuracy for different settings of H^3 and H^4 ($H^k = M^k \times \Delta$) with a peak accuracy of more than 90% in all cases. As can be seen from Figure 6b, the monotonic relationship indicates that the accuracy is not diminished as a result of the distinct patterns of interactions expressed by the variable history intervals.

8.1.4 Distinct Within-Cluster Temporal Interactions—We also examined the performance of the approach when variable history intervals of interaction occur within a given cluster. This case is relatively challenging since the dependence between any pair of neurons' firing pattern is governed only by the presence of a nonzero α_{pq} , while dependence on similar M_{pq} is not present. The same topology of Figure 4a was used with M_{1i} , M_{5i} , M_{9i} , and M_{13i} set to 10; M_{2i} , M_{6i} , M_{10i} , and M_{14i} set to 20; M_{3i} , M_{7i} , M_{11i} , and M_{15i} set to 110; and M_{4i} , M_{8i} , M_{12i} , and M_{16i} set to 120, where $i \in \{1, 2, \dots, 16\}$. Figure 7a demonstrates that the average clustering accuracy reaches approximately 97% around the 100 ms timescale showing that the neuronal elements in each cluster are correctly identified despite the variability in the history intervals within the cluster.

For comparison, we assessed the performance to that obtained using two standard techniques: the crosscorrelogram and the joint poststimulus time histogram (JPSTH) (Gerstein & Perkel, 1969). Both methods rely on using a fixed bin width correlation analysis and therefore are not expected to be robust to variations in the history interval. We tested window sizes of various numbers of bins (each bin was four samples). A correlation matrix was formed in each case,

where each entry corresponded to the peak of the cross-correlogram or JPSTH obtained for the corresponding pair of neurons that exceeded a predefined significance level. Figures 7b and 7c show the average clustering accuracy for different window sizes, obtained using the cross-correlogram- and JPSTH-based similarity matrices as input to a standard hierarchical clustering. In this analysis, we used the average linkage hierarchical clustering algorithm in which the distance between two clusters is computed as the average of the pairwise distance between the members of the two clusters (Jain, Murty, & Flynn, 1999). In a comparison of Figure 7a to 7b and 7c, better clustering accuracy is obtained using the proposed approach.

8.2 Variability in Network Topology

So far we have examined the performance of the algorithm under systematic variation in temporal dependence to illustrate its ability to track nonstationarity in the temporal interaction between neurons within a fixed network structure. Here we examine the ability of the algorithm to track variations in spatial interactions. This was investigated in two ways: by manually inducing changes in the structure of the network in which new connections randomly appear and by gradually varying the connectivity strength based on an activity-dependent STDP model, widely believed to underlie Hebbian learning.

8.2.1 Structural Plasticity—To invoke structural plasticity in the model, new connections were instantaneously created in the population between subsets of neurons from different clusters, as shown in Figure 8a. This can mimic changes in neuronal circuitry resulting from instantaneous activation of distinct pathways, for example, when cortical rewiring takes place after traumatic brain injury or stroke (Winship & Murphy, 2008). These structural changes are long thought to accompany functional plasticity across multiple timescales, even under normal conditions (Fusi, Asaad, Miller, & Wang, 2007; Gross, 2006). Another example of structural changes triggered by functional plasticity is the sudden transition in behavioral goal during a motor task (Sakamoto et al., 2008). This may trigger an observed neuron, previously not involved in encoding behavioral task parameters, to be instantaneously “recruited” based on the correlation between its tuning characteristics and the new task demands. In such a case, the synaptic connections linking this neuron to other across-cluster neurons can be modeled as a step function centered at the time marker of the response trigger, and weighted by an excitatory or inhibitory coupling α_{pq} with appropriate strength.

For illustration purposes, the third neuron in each cluster now excites the first neuron from the neighboring cluster in the clockwise direction, while the fourth neuron in a cluster inhibits the second neuron from the neighboring cluster in the counter-clockwise direction. The synaptic coupling was introduced over distinct time markers and over multiple timescales. Noise was added in the same way as previously described. The across-cluster interaction expressions used for synaptic coupling are

$$\begin{aligned}\alpha_{pq}^+(t) &= 1.25 \sin(4000\pi t / M_{pq}) \exp(-3000t / M_{pq}) \\ \alpha_{pq}^-(t) &= -1.5 \sin(2000\pi t / M_{pq}) \exp(-3000t / M_{pq}).\end{aligned}\tag{8.2}$$

The above expressions differ from equation 8.1 by the constant amplitude term A_{pq} , which indicates relatively weaker coupling between clusters. The history interval was fixed to 120 bins for all interactions.

The similarity matrix shown in Figure 8b shows significant correlation appearing between neuron 5 and neurons 2, 3, and 4, consistent with the model, since neuron 5 is excited by neuron 3, and neuron 3 is coupled to neurons 2 and 4. Similarly, significant correlation appears between neuron 8 and neurons 1, 2, and 3. These changes in the correlation structure lead to a slightly

diminished clustering accuracy (about a 4% decrease) compared to the isolated cluster topology illustrated in Figure 8c. These results demonstrate that despite the presence of weak interactions in the similarity matrix, the clustering accuracy was not significantly affected, indicating the ability of the algorithm to detect subtle differences between weak across-cluster interactions relative to stronger within-cluster ones.

8.2.2 Spike-Timing-Dependent Plasticity—Here, we investigated an activity-dependent model of plasticity by modifying synaptic strength based on the timing of pre- and postsynaptic spike occurrences (Froemke & Dan, 2002; Dan & Poo, 2004). Figure 9a shows the topology used where STDP was induced by simultaneously stimulating two neurons from each cluster. This was achieved by simultaneously increasing the background rate β_p in equation 5.1 for the paired neurons at specific phases. By doing so, we increase the likelihood of observing two spikes from the pre- and the postsynaptic neurons within a shorter interval, thereby potentiating their synaptic connections. We carried out 120 stimulation phases at 0.2 Hz with 12 ms duration in each. The parameter M_{pq} was set to 40 bins (120 ms) for all synaptic couplings. The percentage change in synaptic strength was modeled using the relation $d\alpha_{pq} = A \exp(-|dt|/\tau)$, where dt is the pre/post-interspike interval. The parameters A and τ were chosen as 1.01 and 14.8 ms, respectively, for long-term potentiation (LTP) and -0.52 and 33.8 ms for long-term depression (LTD), respectively. These parameters were chosen according to the model described in Froemke and Dan (2002) in which the measured synaptic changes were fitted with exponential functions similar to Figure 9b.

Increasing stimulation time results in LTP, and this means that the post-synaptic neuron should gradually “merge” into the presynaptic neuron’s cluster. For example, neuron pair (2, 5) was weakly connected before stimulation. During each stimulation phase, these two neurons were stimulated simultaneously. This repetitive pairing resulted in LTP, and the same process was applied to neuron pairs (3, 8), (9, 14), and (12, 15). Thus, the clustering accuracy should decrease if the number of clusters is held equal to the number of clusters prior to the paired stimulation, which is indeed the case in Figure 9c. Figures 9d and 9e show the similarity matrix obtained using timescales up to 96 ms after 10 and 120 stimulation phases, respectively. The results demonstrate that the algorithm is able to track the effect of naturally occurring spikes as a result of STDP and the corresponding changes in the network topology.

8.3 Clustering Large Populations

We examined the performance of the proposed approach when applied to a larger population of 120 neurons. The population consisted of four clusters of 30 neurons each, where each cluster is formed of a 2D grid of neurons. The strength of the connections between clusters 1 and 2 was manually varied similar to the 16-neuron population above. Figure 10b shows the average pairwise distance between neurons of clusters 1 and 2 for variable connectivity strengths between the clusters. These distances were computed using the similarity matrix prior to probabilistic spectral clustering. As expected, the distance decreases as the connectivity strength between the two clusters increases, which is reflected on the average clustering accuracy in Figure 10c. As the across-clusters connections strengthen, the clustering accuracy decreases when the number of clusters is fixed to be 4 and increases when the number of clusters is fixed to be 3.

8.4 Estimating the Timescale of Maximum Accuracy

We have demonstrated that a functional relationship exists between the timescale at which maximum accuracy occurs and the history interval length chosen to model the population interaction. When we are dealing with real data, prior knowledge of this history interval is not available, and therefore this relationship cannot be exactly discovered. For that purpose, we computed a clustering validity index, the Dunn index (DI) (Bezdek & Pal, 1998), defined as

$$DI = \min_{\substack{1 \leq i \leq N, 1 \leq j \leq N \\ i \neq j}} d(C_i, C_j) / \max_{1 \leq k \leq N} \text{diam}(C_k), \quad (8.3)$$

where $d(C_i, C_j)$ is the average distance between spike trains of pairwise neurons in clusters i and j ($i \neq j$), $\text{diam}(C_k)$ is the maximum pairwise distance between spike trains of neurons in cluster k , and N is the total number of clusters. This index will be maximized when the across-clusters distance (the numerator) is maximized, while the within-cluster distance (the denominator) is minimized. Thus, it can identify sets of clusters that are compactly separated.

Figure 11 shows the average Dunn index at different timescales for the data sets analyzed in section 8.1.2. The values shown were computed directly using the clustering outcome from the proposed approach. A comparison of Figure 11 with the clustering accuracy shown previously in Figure 5a shows that for each choice of history interval H , the Dunn index is maximized around the same timescale where the clustering accuracy peaks. In practice, the Dunn index is computed using entries of the augmented R matrix by normalizing the entries in the R matrix by the maximum value, then subtracting each entry in the normalized R from one. The resulting matrix thus comprises the distance between all pairwise neurons and is eventually used in the computation of equation 8.3.

9 Discussion

Two major features of the proposed approach are specifically tailored to the type of neuronal interaction likely to be observed in cortical data: (1) the projection of the spike trains onto a scale space to account for any synergistic interplay between temporal and rate codes that may be simultaneously present in a given ensemble activity pattern, and (2) the probabilistic spectral clustering to overcome any complexity of clustering data with arbitrary structures. This second feature yields two additional advantages. First, it is nonparametric and therefore can track a wide variety of cluster shapes, and second, its probabilistic nature is useful in quantifying the degree of uncertainty in the membership of a given neuron to each cluster. This implies more robustness to plastic changes in cortical circuitry within and across multiple trials.

A typical problem that arises in the use of clustering algorithms in general is the knowledge of the exact number of independent clusters prior to the analysis. While a few heuristic approaches have been proposed in the literature for many types of data, this problem can be quite challenging to the application at hand given the ever-present cortical adaptation caused by synaptic plasticity, cellular conditioning, or representational plasticity (Buonomano & Merzenich, 1998). One popular approach for estimating the number of clusters to look for in the data is the elbow criterion (Jain et al., 1999), which was improved by the introduction of the gap statistic (Tibshirani, Walther, & Hastie, 2001). Specifically, the gap statistic criterion is based on the observation that the within-cluster dispersion, defined as the overall distance between pairwise elements in a cluster, tends to decrease significantly when the number of clusters exceeds the true number. Therefore, it is possible to estimate the correct number of clusters by identifying the elbow (or the “knee”) in the graph of within-cluster dispersion and comparing it to the reduction observed in the within-cluster dispersion of a reference distribution. In a previous study, we reported on the use of the gap statistic applied to a population of leaky integrate-and-fire (LIF) neurons (Chen, El-Dawlatly, Jin, & Oweiss, 2007). We demonstrated that the gap statistic approach was capable of detecting the correct number of independent neuronal clusters in over 90% of the experiments. Similar to the Dunn index used in this letter to determine the maximum clustering accuracy, the gap statistic can also be considered a clustering validity index. However, the main difference between both

measures is that the gap statistic is used prior to the clustering to estimate the number of clusters to look for in the data, while the Dunn index is used after clustering to choose the clustering result corresponding to the timescale that best represents the temporal interaction between neuronal elements in the observed activity patterns.

In addition to the numerous systems neuroscience applications that may benefit from the proposed approach, a fundamental objective in many recent brain-machine interface (BMI) demonstrations is to decode ensemble spike trains to better characterize the nature of sensorimotor transformations associated with distinct motor behaviors (Moran & Schwartz, 1999; Serruya, Hatsopoulos, Paninski, Fellows, & Donoghue, 2002; Taylor, Tillery, & Schwartz, 2002; Hochberg et al., 2006; Lebedev & Nicolelis, 2006; Lebedev, O'Doherty, & Nicolelis, 2008). Many of these studies have examined the utility of multiple single unit recordings in the primary motor, premotor, and supplementary motor areas of the cortex in controlling artificial devices. In the overwhelming majority of these studies, neuronal firing rates estimated over a fixed bin width (typically 50–100 ms) are believed to contain all the information needed to estimate the neural correlate of behavioral. However, the large variability in individual firing rates typically observed across multiple repeated trials results in unreliable decoding. As might be expected, decoding poorly generalizes to nonstereotypical tasks. These algorithms require extensive periodic “calibration,” even when recordings are stable and the subject’s performance remains steady (Serruya, Hatsopoulos, Fellows, Paninski, & Donoghue, 2003).

Faithful decoding of motor cortex activity patterns when the subject is freely behaving and naturally interacting with the surrounding requires dynamic identification of clusters of neurons involved in encoding behavioral parameters. We hypothesize that the inherently varying spectrum of temporal and spatial interactions among cortical neurons motivates the need to assess the functional connectivity between these neurons to improve decoding performance under continuous cortical adaptation. Clustering techniques such as the one reported in this letter could yield substantial improvement in decoding performance by restricting the input to the decoder to a spatiotemporal neural subspace where the relevant task information resides. Thus, they may constitute an important predecoding step that continuously assesses potential variations in network states reminiscent of any plastic changes in the cortical circuitry. Our most recent results demonstrate that a substantial gain in performance can be achieved when side information about these variations is continuously supplied to the decoder compared to a “completely blind” decoder operation (Aghagolzadeh & Oweiss, forthcoming). This may have a potential impact on characterizing how cortical adaptation plays a role in the adaptive control of neuroprosthetic devices.

10 Conclusion

As tools and techniques for recording large neuronal ensembles continue to accelerate, unveiling the functional relationship between recorded neurons becomes naturally motivated. We have proposed a new approach to identify functionally interdependent neurons in a recorded ensemble from their observed spike train activity. Two important features, scale-space projection and spectral clustering steps, enable the algorithm to track spatial and temporal variability in neuronal discharge patterns that may underlie cortical plasticity. The technique also introduced the use of graph theory as an efficient mean to represent large populations of neurons and opens up the opportunity to apply a multitude of other clustering algorithms to discover any inherent structure in the patterns of interaction under various experimental conditions. We also demonstrated that the technique can efficiently track the dynamics of STDP derived from experimental data models reported in the literature. This feature augments the capability of the algorithm to be used in quantifying plasticity of cortical circuits, particularly during learning and memory.

The approach can serve as a preprocessing step to reduce the dimensionality of the neural space after which statistical inference about the functional relationship between the neurons can be obtained. Reverse-engineering the cortical circuit that gave rise to the observed data becomes an ultimate goal in that respect. For that, our current work aims at using the outcome of this algorithm to infer the network topology and validate the results with anatomical observations on real electrophysiological data.

Acknowledgments

We sincerely thank the anonymous reviewers for their helpful suggestions and comments that helped improve the quality of this letter. This work was supported by NINDS grant number NS054148.

References

- Abeles M, Gerstein G. Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *Journal of Neurophysiology* 1988;60(3):909–924. [PubMed: 3171666]
- Abeles M, Goldstein MH. Multiple spike train analysis. *Proceedings of the IEEE* 1977;65:762–773.
- Aghagolzadeh M, Oweiss KG. Identification and tracking of functional plasticity in motor neuronal ensembles improves the performance of population decoders. *Society for Neuroscience Abstracts*. (forthcoming).
- Armstrong-Gold CE, Rieke F. Bandpass filtering at the rod to second-order cell synapse in salamander (*Ambystoma tigrinum*) retina. *J Neurosci* 2003;23(9):3796–3806. [PubMed: 12736350]
- Averbeck BB, Latham P, Pouget A. Neural correlations, population coding and computation. *Nature Reviews Neuroscience* 2006;7:358–366.
- Bezdek JC, Pal NR. Some new indexes of cluster validity. *IEEE Transactions on Systems, Man, and Cybernetics, Part B* 1998;28(3):301–315.
- Bi GQ, Poo MM. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 1998;18:10464–10472. [PubMed: 9852584]
- Brillinger D. The identification of point process systems. *Annals of Probability* 1975;3(6):909–924.
- Brillinger DR. Measuring the association of point processes: A case history. *American Mathematical Monthly* 1976;83(1):16–22.
- Brillinger D. Nerve cell spike train data analysis. *J Am Stat Assoc* 1992;87(418):260–271.
- Brown, EN., editor. *Theory of point processes for neural systems: Methods and models in neurophysics*. Paris: Elsevier; 2005.
- Brown EN, Nguyen DP, Frank LM, Wilson MA, Solo V. An analysis of neural receptive field plasticity by point process adaptive filtering. *PNAS* 2001;98(21):12261–12266. [PubMed: 11593043]
- Buonomano DV, Merzenich MM. Cortical plasticity: From synapses to maps. *Annual Review of Neuroscience* 1998;21(1):149–186.
- Chen, F.; El-Dawlatly, S.; Jin, R.; Oweiss, K. Identifying and tracking the number of independent clusters of functionally interdependent neurons from biophysical models of population activity. *Proc. of 3rd Int. IEEE EMBS Conf. on Neural Engineering*; Los Alamitos, CA: IEEE; 2007.
- Chung, FRK. *Spectral graph theory*. Providence, RI: American Mathematical Society; 1994.
- Dan Y, Poo MM. Spike timing-dependent plasticity of neural circuits. *Neuron* 2004;44:23–30. [PubMed: 15450157]
- Destexhe A, Marder E. Plasticity in single neuron and circuit computations. *Nature* 2004;431:789–795. [PubMed: 15483600]
- Ding, CHQ.; He, X.; Zha, H.; Gu, M.; Simon, HD. A min-max cut algorithm for graph partitioning and data clustering. *Proceedings of the 2001 IEEE International Conference on Data Mining*; Washington, DC: IEEE Computer Society; 2001.
- Dzakpasu R, Zochowski M. Discriminating different types of synchrony in neural systems. *Physica D* 2005;208:115–122.

- Fellous JM, Tiesinga PHE, Thomas PJ, Sejnowski TJ. Discovering spike patterns in neuronal responses. *Journal of Neuroscience* 2004;24(12):2989–3001. [PubMed: 15044538]
- Frerking M, Schulte J, Wiebe SP, Stäubli U. Spike timing in CA3 pyramidal cells during behavior: Implications for synaptic transmission. *J Neurophysiol* 2005;94(2):1528–1540. [PubMed: 15872069]
- Froemke RC, Dan Y. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* 2002;416(6879):433–438. [PubMed: 11919633]
- Fusi S, Asaad WF, Miller EK, Wang XJ. A neural circuit model of flexible sensorimotor mapping: Learning and forgetting on multiple timescales. *Neuron* 2007;54(2):319–333. [PubMed: 17442251]
- Gerstein G, Clark W. Simultaneous studies of firing patterns in several neurons. *Science* 1964;143:1325–1327. [PubMed: 17799237]
- Gerstein GL, Perkel DH. Simultaneously recorded trains of action potentials: Analysis and functional interpretation. *Science* 1969;164:828–830. [PubMed: 5767782]
- Gross L. A new window into structural plasticity in the adult visual cortex. *PLoS Biology* 2006;4(2):e42. [PubMed: 20076528]
- Hebb, DO. *The organization of behavior: A neuropsychological theory*. Hoboken, NJ: Wiley; 1949.
- Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, et al. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 2006;442:164–171. [PubMed: 16838014]
- Jain AK, Murty MN, Flynn PJ. Data clustering: A review. *ACM Comp Surv* 1999;31(3):264–323.
- Jarvis MR, Mitra PP. Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Comput* 2001;13(4):717–749. [PubMed: 11255566]
- Jin, R.; Ding, C.; Kang, F. A probabilistic approach for optimizing spectral clustering. In: Weiss, Y.; Schölkopf, B.; Platt, J., editors. *Advances in neural information processing systems*. Vol. 18. Cambridge, MA: MIT Press; 2005.
- Lebedev MA, Nicolelis MA. Brain-machine interfaces: Past, present and future. *Trends Neurosci* 2006;29:536–546. [PubMed: 16859758]
- Lebedev MA, O'Doherty JE, Nicolelis MAL. Decoding of temporal intervals from cortical ensemble activity. *J Neurophysiol* 2008;99(1):166–186. [PubMed: 18003881]
- Lewicki MS. A review of methods for spike sorting: The detection and classification of neural action potentials. *Network: Computation in Neural Systems* 1998;9(4):53–78.
- Lindsey B, Gerstein G. Two enhancements of the gravity algorithm for multiple spike train analysis. *J Neurosci Methods* 2006;150(1):116–127. [PubMed: 16105685]
- Maass W, Natschlager T. A model for fast analog computation based on unreliable synapses. *Neural Computation* 2000;12(7):1679–1704. [PubMed: 10935922]
- Mallat, S. *A wavelet tour of signal processing*. Norwood, MA: Kluwer; 1998.
- Markram H, Lubke J, Frotscher &, Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 1997;275:213–215. [PubMed: 8985014]
- Martignon LG, Deco G, Laskey K, Diamond M, Freiwald W, Vaadia E. Neural coding: Higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comp* 2000;12:2621–2653.
- Moran DW, Schwartz AB. Motor cortical activity during drawing movements: Population representation during spiral tracing. *Journal of Neurophysiology* 1999;82(5):2693–2704. [PubMed: 10561438]
- Nelson SB, Sjöström PJ, Turrigiano GG. Rate and timing in cortical synaptic plasticity. *Philos Trans R Soc Lond B Biol Sci* 2002;357(1428):1851–1857. [PubMed: 12626018]
- Ng, A.; Jordan, M.; Weiss, A. On spectral clustering: Analysis and an algorithm. In: Dietterich, TG.; Becker, S.; Ghahramani, Z., editors. *Advances in neural information processing systems*. Vol. 14. Cambridge, MA: MIT Press; 2001.
- Normann R, Maynard EM, Rousche PJ, Warren DJ. A neural interface for a cortical vision prosthesis. *Vision Research* 1999;39(15):2577–2587. [PubMed: 10396626]
- Okatan M, Wilson MA, Brown EN. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Computation* 2005;17(9):1927–1961. [PubMed: 15992486]
- Rieke, F.; Warland, D.; de Ruyter van Steveninck, RR.; Bialek, W. *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press; 1997.

- Roberts PD, Bell CC. Spike-timing dependent synaptic plasticity in biological systems. *Biological Cybernetics* 2002;87:392–403. [PubMed: 12461629]
- Rosenberg JR, Amjad AM, Breeze P, Brillinger DR, Halliday DM. The Fourier approach to the identification of functional coupling between neuronal spike trains. *Prog Biophys Mol Biol* 1989;53:1–31. [PubMed: 2682781]
- Sakamoto K, Mushiaki H, Saito N, Aihara K, Yano M, Tanji J. Discharge synchrony during the transition of behavioral goal representations encoded by discharge rates of prefrontal neurons. *Cereb Cortex*. 2008 February 9; bhm234.
- Schneidman E, Berry MJ, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 2006;440(7087):1007. [PubMed: 16625187]
- Serruya M, Hatsopoulos N, Fellows M, Paninski L, Donoghue J. Robustness of neuroprosthetic decoding algorithms. *Biological Cybernetics* 2003;88(3):219–228. [PubMed: 12647229]
- Serruya MD, Hatsopoulos NG, Paninski L, Fellows M, Donoghue J. Brain-machine interface: Instant neural control of a movement signal. *Nature* 2002;416(6877):141–142. [PubMed: 11894084]
- Shadlen MN, Newsome WT. The variable discharge of cortical neurons implications for connectivity, computation, and information coding. *J Neurosci* 1998;18:3870–3896. [PubMed: 9570816]
- Sjöström PJ, Turrigiano GG, Nelson SB. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 2001;32(6):1149–1164. [PubMed: 11754844]
- Srinivasan L, Eden UT, Willsky AS, Brown EN. A state-space analysis for reconstruction of goal-directed movements using neural signals. *Neural Comp* 2006;18(10):2465–2494.
- Stevens CF, Zador AM. Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci* 1998;1:210–217. [PubMed: 10195145]
- Taylor DM, Tillery SIH, Schwartz AB. Direct cortical control of 3D neuroprosthetic devices. *Science* 2002;296(5574):1829. [PubMed: 12052948]
- Tibshirani R, Walther G, Hastie T. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society B* 2001;63:411–423.
- Truccolo W, Eden UT, Fellows MR, Donoghue J, Brown E. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol* 2005;93(2):1074–1089. [PubMed: 15356183]
- Tsodyks MV, Markram H. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *PNAS* 1997;94:719–723. [PubMed: 9012851]
- Winship IR, Murphy TH. In vivo calcium imaging reveals functional rewiring of single somatosensory neurons after stroke. *J Neurosci* 2008;28(26):6592–6606. [PubMed: 18579732]
- Wise K, Anderson D, Hetke J, Kipke D, Najafi K. Wireless implantable microsystems: High-density electronic interfaces to the nervous system. *Proceedings of the IEEE* 2004;92:76–97.

Figure 1.

Synaptic weights for (a) $M_{pq} = 40$ bins ($H_{pq} = 120$ ms) and (b) $M_{pq} = 120$ bins ($H_{pq} = 360$ ms). In both cases, $A_{pq} = 2.5$, $B = 2000$, and $F = 3000$. When the history interval is limited to 40 bins, the function exhibits a smaller time constant within the short time support where the function is nonzero. In contrast, the function exhibits larger time constant when $M_{pq} = 120$.

Figure 2.

Two-level Haar scale space representation of a sample spike train. See the text for an interpretation of this representation.

Figure 3.

A sample undirected graph with four objects and arbitrary edge weights for which the spectral clustering algorithm clusters objects with strong similarity (a small edge weight indicates stronger similarity between objects).

Figure 4.

(a) Network topology for the independent clusters population (solid lines), *I* = inhibition, *E* = excitation. (b) Sample 3-sec trace of the obtained spike trains. The initialization period is indicated by a black bar. (c) Sample ISI histogram for a neuron in the population and an independent neuron at 12 ms and 360 ms history intervals, respectively. Coefficient of variation of the ISI histogram. (d) Augmented correlation matrix using timescales up to 96 ms. (e) Average clustering accuracy versus timescale for a fixed 360 ms history interval obtained using probabilistic spectral clustering and *k*-means clustering.

Figure 5.

(a) Clustering accuracy versus timescale fitted with third-order polynomial for the network in Figure 4a for variable history intervals (b) Gray dots: Timescale of maximum clustering accuracy obtained for the 25 trials versus the population true history interval. Black dots: Average of maximum accuracy timescales.

Figure 6.

(a) Average clustering accuracy $\pm$ std versus timescale fitted with third-order polynomial for different settings of M^3 and M^4 , keeping M^1 and M^2 fixed to 20 bins. Note that $H^k = M^k \times \Delta$ in ms where $\Delta = 3$ ms. (b) Gray dots: Timescale of maximum clustering accuracy versus history interval of clusters 3 and 4 for each of the 25 trials. Black dots: Average of maximum accuracy timescales.

Figure 7.

(a) Average clustering accuracy $\pm$ std over all trials versus timescale. (b) Average clustering accuracy using cross-correlograms for different window sizes. (c) Average clustering accuracy using JPSTH for different window sizes.

Figure 8.

(a) Network topology with across-cluster synaptic coupling (dashed lines). (b) Average clustering accuracy versus timescale. (c) Augmented R using timescales up to 96 ms.

Figure 9.

(a) Network topology with paired stimulation to induce STDP. (b) EPSP slope as a function of pre-post-synaptic spikes amplitude (A) and time constant (τ) were chosen as 1.01 and 14.8 ms for LTP and -0.52 and 33.8 ms for LTD, respectively. (c) Average clustering accuracy poststimulation phases assuming four clusters and two clusters. (d) Similarity matrix using timescales up to 96 ms after 10 stimulation phases and (e) after 120 stimulation phases.

Figure 10.

(a) Network topology for a population of 120 neurons, where each cluster has 30 neurons. The bidirectional connections refer to an excitation connection in one direction and an inhibition in the opposite one. Black connections represent the within-cluster connections, while gray connections represent the across-clusters 1 and 2 connections that undergo functional plasticity. (b) Average pairwise distance between neurons in clusters 1 and 2 for variable connectivity strengths. A connectivity strength of 1 indicates an across-clusters connection of the same strength as the within-clusters connections. (c) Average clustering accuracy for variable across-clusters connectivity strengths (clusters 1 and 2).

Figure 11.

Average Dunn index $\pm$ std over all trials versus timescale fitted with third-order polynomial for the network in Figure 4a for different history intervals.