Limitations of principal component analysis as a method to detect neuronal assemblies

Camila Sardeto Deolindo*, Ana Carolina Bione Kunicki *, Fabricio Lima Brasil*, Renan Cipriano Moioli*
*Edmond and Lily Safra International Institute of Neuroscience of Natal - Natal, Brazil
Email: camila.sardeto@natalneuro.org.br

***Abstract*—The anatomical and functional characterization of neuronal assemblies (NAs) is a major challenge in neuroscience. Principal component analysis (PCA) is a widely used method for feature detection, however, when dealing with neuronal data analysis, its limitations have not yet been fully understood. Our work complements previous PCA studies which, in general, characterise NAs based solely on excitatory neuronal interactions. We analysed the performance of PCA in two neglected scenarios: assemblies containing patterns of neural interactions 1) with inhibition and 2) with delays. The analyses considered two types of artificially generated data, one drawn from a traditional poissonian model, and the other drawn from a latent multivariate Gaussian model; in both models, data from a behaving Wistar rat was used for parameter tuning. Our results highlight scenarios in which neglecting complex interactions between neurons can lead to false conclusions when using PCA to detect NAs. Also, we reinforce the importance of more realistic simulations in the evaluation of neuronal signal processing algorithms.**

***Index Terms*—Neuronal Assemblies· Principal Component Analysis (PCA)· Neural Simulation·**

I. Introduction

In his book, Hebb [1] proposed that the interactions established between different neurons (synapses) were the basis of associative learning, and that modifications at cellular level would result in metabolic changes, leading to alterations in the activity pattern of a spatially distributed "assembly of nerve cells". The Hebbian synapse consists of a time-dependent, highly local and strongly interactive mechanism that increases synaptic efficiency as a function of the correlation between the presynaptic and the postsynaptic activities [2].

Hebb's description has led to the neuronal assembly (NA) hypothesis that defines a whole population - and not a single neuron- as the fundamental brain unit in functional signal processing [3, 4, 5]. The theory states that a single neuron, being naturally unstable due to the fluctuations of the thousands of inputs it receives per second, could not account for the complex computational processes in the brain and some of its phenomena, such as redundancy and learning[6, 7, 8].

The NA hypothesis has been extensively discussed in the scientific community, since it connects to a series of psychological and physiological phenomena [7]. However, the experimental confirmation of the existence and operation of assemblies is one of the major challenges of neuroscience [5].

To study NA, one first needs to simultaneously record from a large number of neurons, a challenging task which has progressed significantly in recent years [3, 5, 9]. This data is then processed taking into account a pre-established pattern of neuronal interactions, i.e., a metric based on a specific aspect of neuronal activity and which defines the formation or suppression of neuronal assemblies over time. Examples include covariance of mean firing rates [10, 11, 12, 13], coincident spiking times [4, 10, 11], and mutual information [10, 11, 14]. Finally, possible neuronal clusters are identified with respect to the chosen metric.

A popular method for assembly detection is principal component analysis (PCA) [15, 16, 17]. In previous studies employing neural simulation only synchronous excitatory interactions between neurons, as captured by the covariance of their mean firing rates, were considered [12, 13]. We complement these studies by analysing the performance of PCA in two, to the best of our knowledge, seldom explored scenarios: assemblies defined by patterns of neural interactions 1) with inhibition and 2) with delays. Likewise the excitatory case, inhibitory and time-delayed interactions also relate to Hebbian interactions [18, 19] and have been found in cortico-cortical and cortico-subcortical circuits [20, 21, 22].

We perform our analyses using two biologically plausible artificial spike train generation methods, one based on a homogeneous poissonian model [12, 13], the other relying on a latent multivariate Gaussian model [23]. Both models are tuned using data recorded from a behaving Wistar rat. We evaluate PCA response to simulated data that has a specified correlation structure, establishing a relation between its performance and real data statistics. In addition, our results contribute to define the conditions in which PCA is an appropriate algorithm for neuronal assembly detection.

This paper is organized as follows: Section II brings an overview of the PCA algorithm; Section III describes the experiments and their implementation procedures; Section IV shows the results and Section V presents the final discussion, conclusions, and further work proposals.

II. Theoretical background

A. Principal component analysis

PCA is one of the best known technique in multivariate analysis, tackling data compression and statistical pattern recognition in a dataset. The idea is to extract features, making a domain transformation of the input in a way that a reduced number of features retaining most of the explanatory capacity can be selected. The PCA algorithm is the optimum linear transformation in the mean square error sense, resulting in a set of linearly uncorrelated variables, the so called principal

978-1-4799-6644-8/14/$31.00 ©2014 IEEE

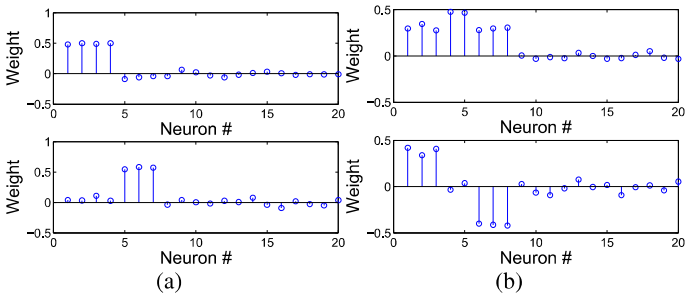

Fig. 1: Characteristic assembly firing patterns have similar weights for assembly members and null weights for other neurons. Figure 1a shows one example in which patterns can be successfully isolated and Figure 1b one example in which they can't. (Adapted from Lopes-dos Santos et al. [13])

components (PCs), representing the directions that maximize the rate of variance decrease [2].

Traditional PCA has limitations when dealing with high-order dependencies, since it's a linear method, is independent of the data source, having no tunable parameters, and, when applied to classification problems, assigns each new sample strictly based on what was observed on the training period, not updating statistical properties dynamically.

To overcome these limitations, several adaptations have been proposed, e.g. the Dynamic PCA (DPCA)[24] or the Recursive PCA [25], which deal with dynamic (non-stationary) processes; the Multi-Scale PCA [26], that aims to improve the performance when data have different time-scales; and, finally, the Non-Linear PCA [27], Kernel-PCA [28] which tackle non- linear relationships.

Specific to the NA detection problem, the PCs are calculated from multiple time series, each corresponding to a single neuron's action potential occurrences. Each temporal series is binned, resulting in a matrix ($M_{n_{Neurons}XT}$), where each entry $m_{n,t}$ denotes the number of spikes of a given neuron n in a given time bin t, $n_{Neurons}$ is the number of recorded neurons and T the total number of time bins. This matrix is then normalized by a z-score transformation. The PCs $\mathbf{p_i}$ are given by the set of eigenvectors $P = [\mathbf{p_1}, ..., \mathbf{p_{n_{Neurons}}}]$ of the covariance of the resultant normalized matrix.

The PC weights, given by the elements w_k of a PC $\mathbf{p}$, $k = 1, ...n_{Neurons}$, may indicate the neurons composing each NA, since a characteristic assembly firing pattern would have the same weights for assembly members and null weights for other neurons, as exemplified in Figure 1a.

However, some limitations exist when two assemblies share neurons: in this case, the NA shared neurons firing patterns, as captured by the covariance matrix, would result in PCs with diverse weight patterns associated to those neurons who constitute the assembly, thus impairing the NA membership identification [12, 13]. See Figure 1b. In this work, we explored the extent of these limitations in inhibitory interactions and in interactions with delays.

III. Methods

We analysed PCA performance when distinguishing between two NAs generated artificially by two biologically plausible methods, described in Section III-A. Twenty four single units were recorded from an awake behaving rat and the mean firing rate and the covariance matrix were used to tune our spike generation models and defined the statistical properties of our artificial populations.

To obtain the data, 32 tungsten microwires were surgically positioned within primary motor cortex (M1) of a Wistar rat. The microwires were distributed in 4X8 arrays, spaced $500\mu m$ in anteroposterior (AP) and mediolateral (ML) axis. The following coordinate relative to bregma in millimeters were used to center the arrays: +1.5 AP, +2.5 ML, 1.2 DV; [DV=dorsoventral]). The study was approved by the Ethics Committee from the University of São Paulo Medical School General Hospital number CAPPesq 0948/09.

Artificial assemblies were classified in three categories, as proposed by Lopes-dos Santos et al. [13]: disjoint assemblies, assemblies with at least one exclusive neuron, and assemblies without exclusive neurons. To determine if a group of neurons present correlated firing, thereby forming a NA, or if their firing activity is independent from each other, we follow the approach of Lopes-dos Santos et al. [12] and Peyrache et al. [17]. Basically, it has been shown that the eigenvalues of an autocorrelation matrix with statistically independent rows follow the Marčenko-Pastur distribution, hence this analytical distribution can be used as a statistical threshold to relate the PCs with correlated firing activity (assemblies). For example, if a given group of neurons form a NA, the associated PC eigenvalue of their firing activity will be greater than the threshold established by the Marčenko-Pastur distribution; conversely, if the neurons are independent, the PC eigenvalue will be located within this distribution.

A. Spike train generation

1) Method 1 - Spike train generation using a Poisson distribution: In the first method, the neuronal firing was treated as a homogeneous poissonian process, a highly employed method due to the time-independent simplification [4, 29, 30].Each simulated spike train was generated following a poissonian distribution. The activity of neurons selected to compose each assembly was modified by a altering the mean firing rate on a given number of randomly selected bins. This variation was positive when simulating excitatory interactions, and negative for inhibitions. In a second round of simulation, relative delays between the individual assembly members were defined, characterizing not only synchronous, but also non-zero phase differences.

2) Method 2 - Spike train generation using the cross-correlation between neurons: In the second method, we analysed a more complex neuronal group firing pattern, employing a richer temporal structure activity. Spike trains were simulated by a latent multivariate Gaussian model, following the method proposed by Macke et al. [23]. Such approach is important since a model with correlations accounts for real data better

than a model assuming independence [23]. The NAs were specified by the cross-covariance between neurons where non-zero values define the neurons belonging to each group. Thus, a functional connectivity is established between distinct units through the addition of statistical dependencies [5].

B. Experimental setup

The spiking activity extracted from the behaving rat data resulted in a mean firing rate of 5.12 spikes/bin. Our artificial spike trains had a length of 10^4 bins, each bin with $1ms$.

As mentioned in Section III, we analysed three categories of assemblies regarding their elements intersection. Thus, from our simulated neuronal pool, some neurons were chosen to participate in our artificial NAs A_i, i=1,2, according to the definition below. All other neurons remained with an independent activity.

- Disjoint assemblies: A_1={#1,#2,#3,#4}; A_2={#5,#6,#7}.
- Assemblies with at least one exclusive neuron: A_1={#1,#2,#3,#4,#5}; A_2={#4,#5,#6,#7,#8}.
- Assemblies without exclusive neurons: A_1={#1,#2,#3,#4,#5}; A_2={#3,#4,#5}.

For the spike trains generated according to the Poisson distribution, we have evaluated the outcome for N=300 coincident bins with a total of 20 neurons in our population. The performance variation with sparseness will not be analysed in this study.We analysed the performance change with respect to the variation in the randomly selected bins firing rate, gradually varying it (step 1%) from -95% up to 95% of the baseline value. Then, we repeated the analysis, introducing a one bin phase delay in the last neuron of each assembly.

For the spike trains using the cross-correlation model, our population was reduced to 10 neurons. In order to make our simulations more biologically plausible and avoid the use of positive definite covariance matrices that can't characterize a multivariate binary distribution covariance matrix [23], we extracted from real data 53 feasible values for positive and negative correlations to define the interactions between cells. Our goal was to determine how correlated a population must be in order to be identified by PCA. For each assembly, all elements covariance simultaneously assume the same value.

C. Assembly membership detection

As mentioned in Section II-A, the weights w_k of each PC may be an indicative of the associated neuron contribution to the assembly.

We have selected a set of "significant neurons" in our population, by establishing a threshold according to Equation 1 below. A neuron k was considered as part of an assembly if $|w_k|$ is greater than Thr, and not significant otherwise.

$$Thr = 2 \sum_{k=1}^{n_{Neurons}} \frac{|w_k|}{n_{Neurons}} \tag{1}$$

where $n_{Neurons}$ represent the total number of neurons in the analysed population.

The gain in Equation 1 was set based on several executions of the experiment with a disjoint assembly modelled following a poissonian process (Section III-A1), with a high increase of mean firing rate in the selected bins, until the relationship between false-positives and false-negatives was balanced.

D. Performance metrics

The definition of a metric for performance considered the following membership characteristics: 1) number of detected assemblies (Section III-D1) and 2) membership identification (Section III-D2). We used a different performance metric with respect to each characteristic to avoid misinterpretations. For example, the algorithm can be excellent at detecting how many assemblies exist, but lousy at detecting neurons which belong to that assembly.

1) Performance metrics - Number of detected assemblies: The performance metric for the number of detected assemblies ($P_{n_{Ass}}$) is given by:

$$P_{n_{Ass}} = \begin{cases} 1 - \frac{|nda - nra|}{nra}, \ if \quad 0 < nda < 2nra \\ 0, otherwise \end{cases} \tag{2}$$

where:

- nda is the number of assemblies detected by the algorithm, numerically equal to the number of eigenvalues greater than the maximum bound of the Marčenko Pasteur distribution [12].
- nra is the real number of assemblies pre-defined in the simulated data used as input to the algorithm.

This metric assumes a linearly decaying performance with respect to the number of incorrectly (miss)identified NAs. In other words, we equalize the effect of false positives and false negatives.

2) Performance metrics - Membership identification: A method for quantifying similarity between the significant neurons of each identified assembly and the pre-defined cells was established: it's performed by the creation of a so called "Similarity Matrix", ($S_{ndaXnra}$), using nda and nra as defined in Section III-D1.

Each element $s_{i,j}$ of S is computed as:

$$\begin{cases} s(i,j) = \begin{cases} \frac{diff(i,j)}{nReal(i)}, \quad if \quad diff(i,j) > 0 \\ 0, otherwise \end{cases} \\ \\ diff(i,j) = N_{Equal}(i,j) - N_{Diff}(i,j) \end{cases} \tag{3}$$

where:

- $N_{Equal}(i,j)$ is the number of neurons that belong both to the established assembly i and identified assembly j.
- $N_{Diff}(i,j)$is the number of neurons that are not in agreement between the established assembly i and identified assembly j (false positives).
- $nReal(i)$is the number of pre-defined neurons that compose the established assembly i.

Analysing matrix S, we can determine which of the encountered groups are more similar to the pre-defined groups. Thus, the performance for membership identification is calculated as:

$$P_{mem} = \sum_{k=1}^{nda} \frac{max\left(S(l,k)\Big|_{l=1}^{nra}\right)}{nra} \quad (4)$$

Such metric assumes that each false positive neuron will penalize the index by cancelling the contribution of a correctly identified neuron. Although Equation 4 apparently captures the effect of the total number of assemblies, it lacks a penalty for any false positive group of cells, and is, therefore, complemented by Equation 2.

IV. Results

A. Spike train generation using a Poisson distribution

The performance of the spike trains artificially generated using the Poisson distribution can be seen in Figure 2. Each column correspond to one of the defined membership categories (Section III). The first line corresponds to the number of detected assemblies metric (Equation 2) and the second to the membership identification metric (Equation 4). The case with delays follows in Figure 3.

Notice in Figures 2 and 3 the delimitation of performance "boundaries" and that all membership categories decreased performance (for both metrics) with lower firing rate variations in the selected bins. Also, there was a decay in performance seen in the diagonals of Figures 2d, 3d, 2e and 3e, more pronounced in the cases where assemblies shared neurons.

The scenarios portrayed in Figures 2f and 3f show that the membership identification metric also form a triangular "boundary" delimiting a slightly better performance region.

In Figure 3, we notice that the algorithm presented little robustness to delays. As soon as the perturbation was inserted, the corresponding neuron could no longer be detected.

B. Spike train generation using the cross correlation between neurons

The performance of the artificially generated spike trains with interactions defined based on the cross correlation between assembly members is presented in Figure 4.

Black dashed lines were inserted to explicit that our points are not equally distributed, since they were measured from real data. Thus, in order to establish general conclusions for the results in Figure 4, we restricted our analyses to the parts with more points (in the center), and excluded the borders. Spike trains with specified correlation coefficients submitted to our performance metrics reinforce the affirmation of Macke et al. [23], that even seemingly small pairwise correlations result in dramatic changes in the occurrence of firing patterns across many neurons. That is especially verified for negative correlation coefficients, when compared to poissonian models. We notice that the algorithm identifies poorly the number of NAs in those conditions, with a considerable amount of false positive results (not shown), leading to a poor performance.

V. Discussion

Modern equipments with several built-in functions for signal processing is a reality in neuroscience experiments. PCA is already a widely used method for spike sorting and is generally an embedded algorithm but experimenters are generally not concerned with its limitations. Because of its simplicity and suitability for pattern identification, the algorithm is a natural candidate for identifying NAs. This work contributed to elucidate the appropriate conditions for PCA usage.

We have proposed two metrics to analyse PCA performance for identifying NAs (Section III-D) and complemented previous studies considering other forms of interactions, inhibitory and time-delayed patterns of interactions. We performed our analysis using two types of plausible spike train generation methods (Sections III-A1 and III-A2). Following our approach, an experimenter can analyse the statistics of the measured spike trains and gain further insights as to decide whether or not PCA is an appropriate algorithm. The choice of the two metrics (Section III-D) was made to improve visualisation, but one could combine them into a single equation, obtaining a unique performance index.

Perturbations of similar magnitude applied to both NAs produce less distinguishable PCs weights due to our fixed membership threshold (Equation 1), which increase the number of false positives and false negatives. This may explain the performance decay in the scenario portrayed in the diagonals of Figure 2d and 3d (disjoint NAs), with more pronounced effects depicted by Figures 2e and 3e (NAs with common neurons).

The case where NAs have no exclusive neuron analysed through the membership identification metric (Figures 2f and 3f) shows that in the limit all cells are identified but it's not possible to separate them into two distinct populations. The PC weights of neurons in the assembly (w_k) tend to increase in magnitude the higher the variation of the firing rate, explaining the triangular performance "boundaries" of those figures.

Our results (Figures 2 and 3) confirm the statement of Lopes-dos Santos et al. [12], that a minimal activation firing rate is required for the proper detection of the number of NAs. We complemented this assertion by including inhibitory and delayed interactions. The performance decay is explainable since the feature extraction is less efficiently performed in low variance conditions. Statistically significant results are associated with an eigenvalue threshold (in this work, the maximum bound of the Marčenko- Pastur distribution, but one obtained from a percentile in surrogate methods can also be employed). The eigenvalue indicates the amount of variance in a given direction and, thus, there is always a compromise between significance and sensibility, justifying the poor performance in central regions of the aforementioned figures. PCA has a poor output in such conditions and different algorithms that are suited for identification of low variance patterns are required.

Also, as expected, PCA has limitations in the sense of extrapolating its analysis to delayed interactions. As soon

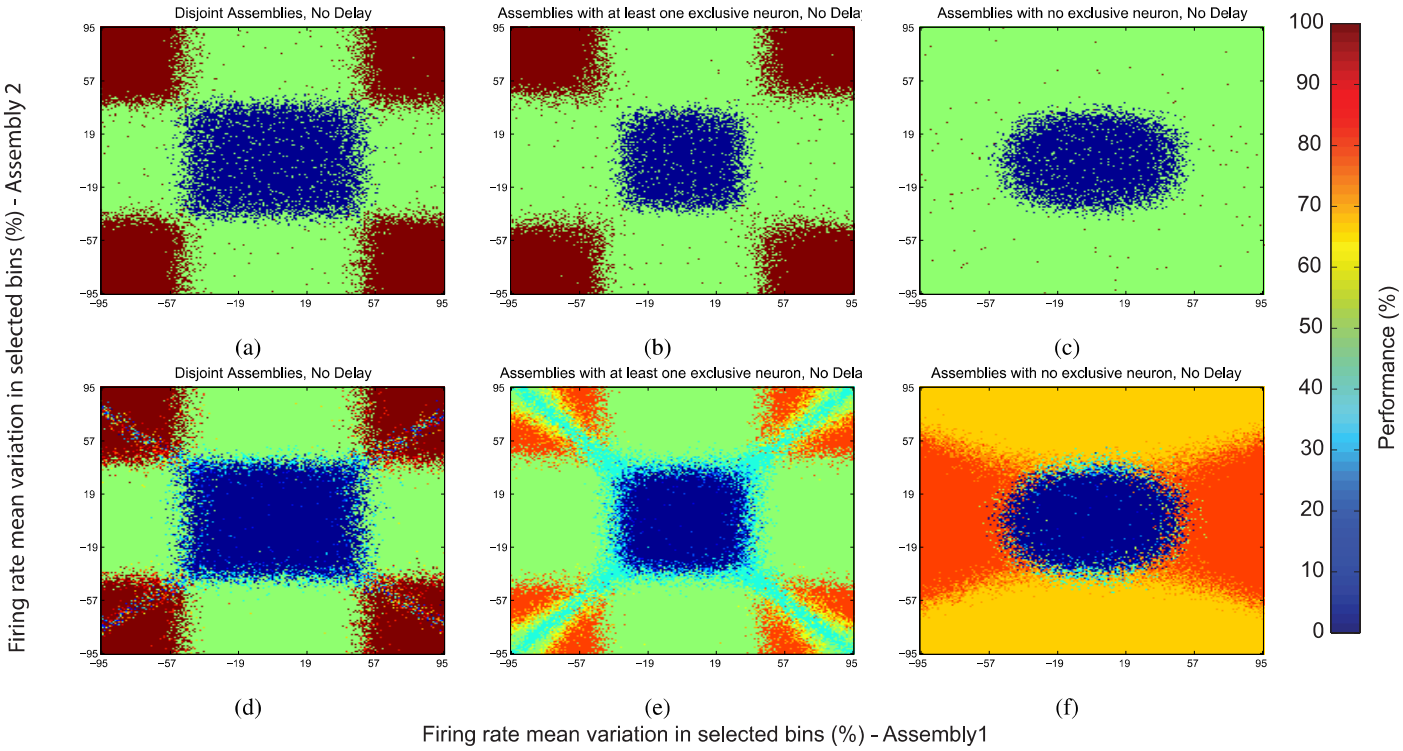

Fig. 2: PCA performance when applied to artificially generated Poisson Spike Trains. The first and the second lines correspond to the number of detected assemblies metric (Section III-C)and to the membership identification metric (Section III-D2), respectively. Each column correspond to one of the defined membership categories (disjoints, at least with one exclusive neuron and without exclusive neurons). In 10^4 samples, 300 bins were randomly selected, in which the mean firing rate of the poissonian model was modified by a scaling factor.

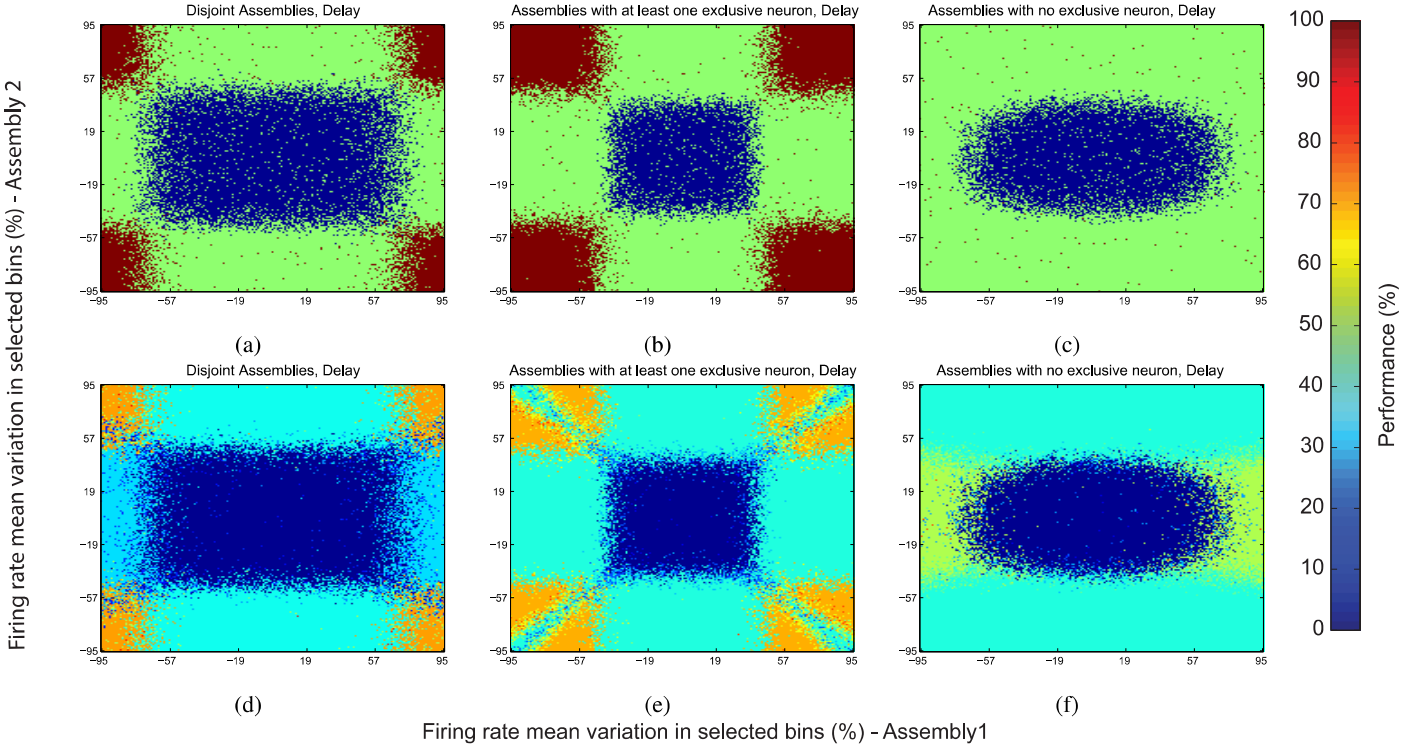

Fig. 3: PCA performance applied to the same assemblies indicated in Figure 2, but now with a one bin delay added to the last element of each assembly. The first and the second lines correspond to the number of detected assemblies metric (Section III-C)and to the membership identification metric (Section III-D2), respectively. Each column correspond to one of the defined membership categories (disjoints, at least with one exclusive neuron and without exclusive neurons). We can note by analysing the membership identification metric (Section III-D2) that the algorithm has little robustness to interactions with non-zero phase difference.

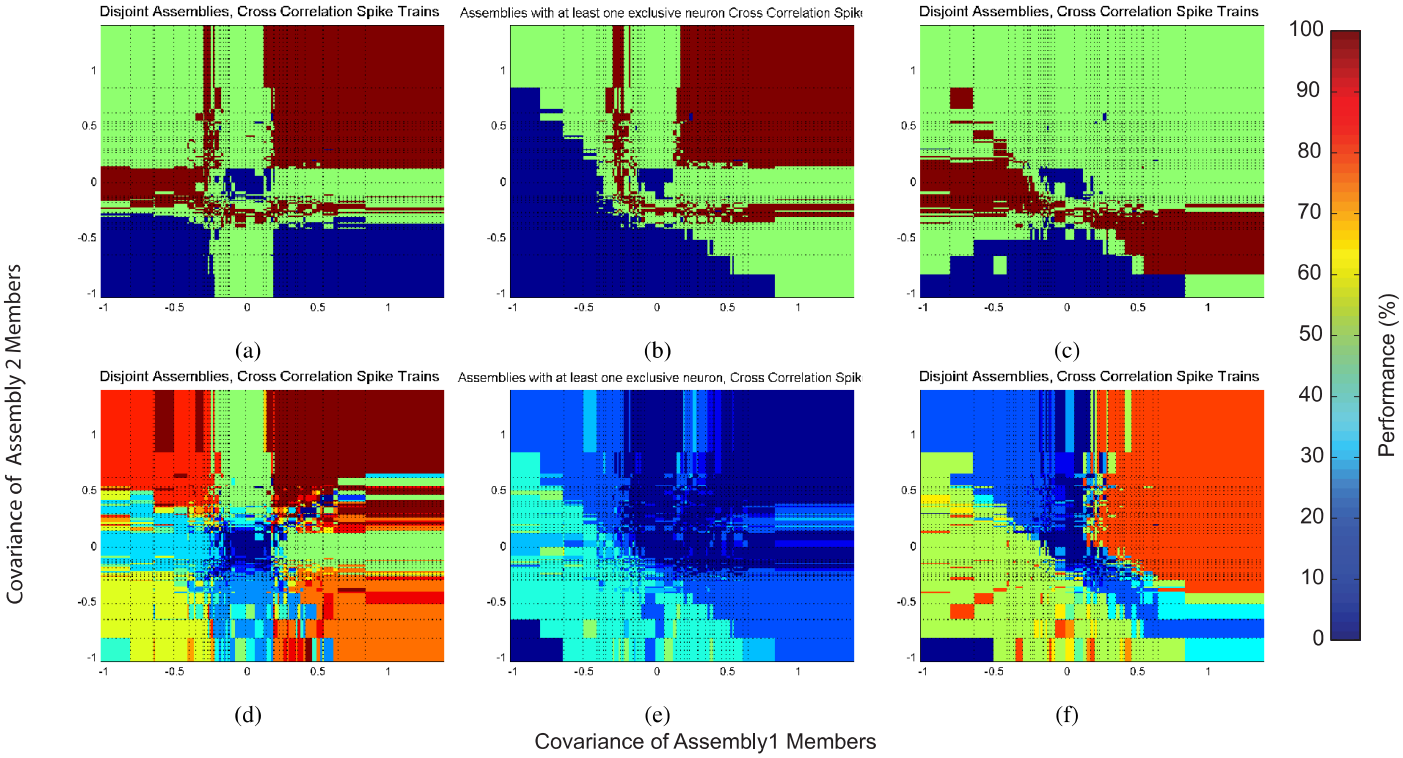

Fig. 4: PCA performance applied to artificial assemblies created using the algorithm proposed by Macke et al. [23], with correlation values obtained by real data statistics. Each column correspond to one of the defined membership categories (disjoints, at least with one exclusive neuron and without exclusive neurons). The first and second lines corresponds to the number of detected assemblies metric (Section III-C) and to the membership identification metric (Section III-D2), respectively. Black dashed lines delimit the 53 different cross correlation values for each axis (extracted from real data) and were inserted to highlight that the employed correlation values are not uniformly distributed.

as a delay is added, the corresponding neuron is no longer identified, which can be verified by comparing Figure 2 with Figure 3. Indeed, in real data, results are dependable on the time bin definition and the experimenter can modify its value to embrace delays. However, in huge populations, with great amounts of data, such conditions are not always verifiable by a person and one must be aware that the method has a limitation in identifying more complex synaptic relations, such as the ones established with Long-Term Synaptic Potentiation (LTP) and Long-Term Synaptic Depression(LTD).

As mentioned in Section III-A1, the modelling of spike trains as a Poisson process is widely used. However, such simplification does not match all neural response variability. By analysing the Fano factors, interspike interval distributions, and coefficients of variation, the veracity of such assumption can be verified. Here we complement the analysis by presenting responses to artificial spikes containing a richer temporal structure, using the method proposed by Macke et al. [23]. In future works, we propose analyses of other sorts of spike train models, e.g. the Poisson model with refractoriness; Krumin and Shoham [31]; Oweiss [32].

Contrasting to poissonian spike trains (Figures 2c and 3c), Figure 4c, employing a richer temporal structure, shows regions where assemblies without exclusive neurons can be correctly identified. Such observations highlight that the neglect of complex interactions between neurons can lead to false conclusions for real data and reinforce the need of more realistic simulations to test signal processing methods.

Finally, we notice that a poor performance in the identification of number of NAs does not automatically imply a similar result in membership. That is an important remark because, depending on the application, a penalized performance in this metric is acceptable, especially with the development of new techniques that increase simultaneous recording of neurons.

Acknowledgment

The authors would like to thank Vítor Lopes dos Santos, Jacob Macke and Phillip Berens for sharing know-how and information. This project had financial support from the National Institutes of Science and Technology program (INCTs) of the national Council for Scientific and Technological Development (CNPq/MCTI), from the Brazilian Innovation Agency (FINEP), from the Coordination for the Improvement of Higher Education Personnel (CAPES), from the Rio Grande do Norte Research Foundation (FAPERN) and from the Alberto Santos Dumont Association for Research Support (AASDAP).

References

[1] D. Hebb, *The Organization of Behavior*. Wiley, New York, 1949.

[2] S. Haykin, *Neural Networks and Learning Machines*, 3rd ed., 2009.

[3] M. A. Nicolelis, E. E. Fanselow, and A. A. Ghazanfar, "Hebb's Dream: The Resurgence of Cell Assemblies," *Neuron*, vol. 19, no. 2, pp. 219–221, Aug. 1997.

[4] K. D. Harris, "Neural signatures of cell assembly organization." *Nat. Rev. Neurosci.*, vol. 6, no. 5, pp. 399–407, May 2005.

[5] G. Buzsáki, "Neural syntax: cell assemblies, synapsembles, and readers." *Neuron*, vol. 68, no. 3, pp. 362–85, Nov. 2010.

[6] A. Paisley and A. Summerlee, "Relationships between behavioural states and activity of the cerebral cortex," *Prog. Neurobiol.*, vol. 22, pp. 155–184, 1984.

[7] G. Gerstein, "Neuronal assemblies," *IEEE Trans. Biomed. Eng.*, vol. 36, no. 1, pp. 4 – 14, 1989.

[8] M. Steriade, D. McCormick, and T. Sejnowski, "Thalamocortical oscillations in the sleeping and aroused brain," *Science*, vol. 262, no. 5134, pp. 679–685, 1993.

[9] D. A. Schwarz, M. A. Lebedev, T. L. Hanson, D. F. Dimitrov, G. Lehew, J. Meloy, S. Rajangam, V. Subramanian, P. J. Ifft, Z. Li, A. Ramakrishnan, A. Tate, K. Z. Zhuang, and M. A. L. Nicolelis, "Chronic, wireless recordings of large-scale brain activity in freely moving rhesus monkeys." *Nat. Methods*, vol. 11, no. 6, Apr. 2014.

[10] R. Quian Quiroga and S. Panzeri, "Extracting information from neuronal populations: information theory and decoding approaches." *Nat. Rev. Neurosci.*, vol. 10, no. 3, pp. 173–85, Mar. 2009.

[11] S. Panzeri, N. Brunel, N. K. Logothetis, and C. Kayser, "Sensory neural codes using multiplexed temporal scales." *Trends Neurosci.*, vol. 33, no. 3, pp. 111–20, Mar. 2010.

[12] V. Lopes-dos Santos, S. Conde-Ocazionez, M. A. L. Nicolelis, S. T. Ribeiro, and A. B. L. Tort, "Neuronal assembly detection and cell membership specification by principal component analysis." *PLoS One*, vol. 6, no. 6, Jan. 2011.

[13] V. Lopes-dos Santos, S. Ribeiro, and A. B. L. Tort, "Detecting cell assemblies in large neuronal populations." *J. Neurosci. Methods*, vol. 220, no. 2, pp. 149–66, Nov. 2013.

[14] J. D. Long and J. M. Carmena, "A statistical description of neural ensemble dynamics." *Front. Comput. Neurosci.*, vol. 5, no. November, p. 52, Jan. 2011.

[15] J. K. Chapin and M. A. Nicolelis, "Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations." *J. Neurosci. Methods*, vol. 94, no. 1, pp. 121–40, Dec. 1999.

[16] J. Hu, J. Si, B. Olson, and J. He, "Feature detection in motor cortical spikes by principal component analysis," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 13, no. 3, pp. 256–262, 2005.

[17] A. Peyrache, K. Benchenane, M. Khamassi, S. I. Wiener, and F. P. Battaglia, "Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution." *J. Comput. Neurosci.*, vol. 29, no. 1-2, pp. 309–25, Aug. 2010.

[18] T. Sejnowski, S. Chattarji, and P. Stanton, "Induction of synaptic plasticity by Hebbian covariance in the hippocampus," *Comput. neuron*, 1989.

[19] Y. Frégnac, "Hebbian Cell Assemblies," *Encycl. Cogn. Sci.*, 2003.

[20] X. L. Yang, D. V. Senthilkumar, and J. Kurths, "Impact of connection delays on noise-induced spatiotemporal patterns in neuronal networks." *Chaos An Interdiscip. J. Nonlinear Sci.*, vol. 22, no. 4, Dec. 2012.

[21] B. Xu, Y. B. Gong, L. Wang, and C. L. Yang, "Dependence of delay-induced coherence resonance on time-periodic coupling strength in Newman-Watts neuronal networks," *Eur. Phys. J. B*, vol. 85, no. 9, p. 299, Sep. 2012.

[22] Q. Wang, H. Zhang, M. Perc, and G. Chen, "Multiple firing coherence resonances in excitatory and inhibitory coupled neurons," *Commun. Nonlinear Sci.*, pp. 1–7, 2012.

[23] J. H. Macke, P. Berens, A. S. Ecker, A. S. Tolias, and M. Bethge, "Generating spike trains with specified correlation coefficients." *Neural Comput.*, vol. 21, no. 2, pp. 397–423, Feb. 2009.

[24] D. M. Fallahati, M. Backonja, H. Eghbalnia, and A. H. Assadi, "Dynamic PCA for network feature extraction in multi-electrode recording of neurophysiological data in cortical substrate of pain," *Neurocomputing*, vol. 44-46, pp. 401–405, Jun. 2002.

[25] T. Voegtlin, "Recursive principal components analysis." *Neural Netw.*, vol. 18, no. 8, pp. 1051–63, Oct. 2005.

[26] B. R. Bakshi, "Multiscale PCA with application to multivariate statistical process monitoring," *AIChE J.*, vol. 44, no. 7, pp. 1596–1610, Jul. 1998.

[27] D. Fotheringhame and R. Baddeley, "Nonlinear principal components analysis of neuronal spike train data." *Biol. Cybern.*, vol. 77, no. 4, pp. 283–8, Oct. 1997.

[28] B. Schölkopf, A. Smola, and K. Müller, "Nonlinear component analysis as a kernel eigenvalue problem," *Neural Comput.*, no. 44, 1998.

[29] M. Laubach, M. Shuler, and M. A. Nicolelis, "Independent component analyses for quantifying neuronal ensemble interactions." *J. Neurosci. Methods*, vol. 94, no. 1, pp. 141–54, Dec. 1999.

[30] A. Amarasingham, T.-L. Chen, S. Geman, M. T. Harrison, and D. L. Sheinberg, "Spike count reliability and the Poisson hypothesis," *J. Neurosci.*, vol. 26, no. 3, pp. 801–9, Jan. 2006.

[31] M. Krumin and S. Shoham, "Generation of spike trains with controlled auto- and cross-correlation functions." *Neural Comput.*, vol. 21, no. 6, pp. 1642–64, Jun. 2009.

[32] K. Oweiss, *Statistical signal processing for neuroscience and neurotechnology*, K. Oweiss, Ed., 2010.