

Detecting Non-Redundant Collective Activity of Neurons

Medorian D. Gheorghiu
Transylvanian Institute of
Neuroscience
Cluj-Napoca, Romania
Technical University of Cluj-Napoca
Email: medorian@tins.ro

Adriana Nagy-Dăbâcan
Transylvanian Institute of
Neuroscience
Cluj-Napoca, Romania
Email: dabacan@tins.ro

Raul C. Mureșan
Transylvanian Institute of
Neuroscience
Cluj-Napoca, Romania
Email: muresan@tins.ro

Abstract—Brain activity is characterized by the activity of huge number of neurons that are densely connected in complex networks. Detecting how neural activity is coordinated across time and space over brain circuits is a very challenging task. In particular, ensembles of neurons may co-activate in specific patterns whose expression correlates to various stimuli, cognitive states, or behavioral outcomes. These patterns are expressed on a variety of timescales, from milliseconds to hundreds of milliseconds or seconds. Here, we extend a recently-introduced method for extracting stereotypical firing patterns by applying several post-processing steps that enable the precise estimation of pattern wavefronts irrespective of the integration timescale used to detect patterns. Using electrophysiology data recorded in mouse visual cortex, we show that this approach enables a more precise estimation of the relation between firing patterns and meso-scale dynamics, reflected in the local-field potential. The method removes the redundancy and blurring generated by the convolution of spike-trains with exponentially-decaying kernels and offers a sharper representation of the expression of firing patterns in brain data.

Keywords: multi-neuronal activity, population dynamics, ensembles, cross-correlation, auto-correlation, visual cortex, data reduction

I. INTRODUCTION

The analysis of neuronal data poses great challenges for several reasons. One of the greatest difficulties comes from the fact that the brain contains large numbers of interconnected neurons, each firing at various time delays relative to other neurons. In response to a visual, auditory or olfactory stimulus, information about that stimulus is encoded in highly distributed neuron firing patterns expressed across the cortical network [1]. Therefore, multi-electrode recordings are necessary to properly characterize the dynamics of brain circuits, but this results in large amounts of multidimensional data that are difficult to analyze. The obvious thing to do is to find techniques able to cope with multi-dimensional data. It was shown in [2] and [3] that the collective firing of the neurons can be characterized using multidimensional vectors (patterns). Stereotypical patterns expressed in response to different stimuli can be further identified by using various clustering algorithms, such as K-means [3] [4] or Kohonen mapping [2] [5] [6]. However, such clustering methods require a priori knowledge of the number of clusters in the data, which is not always obvious. In addition, they are also sensitive to the initial conditions

and thus they may generate a sub-optimal approximation of the data.

Here, we propose a method for “correcting” the misrepresentation of multi-dimensional pattern data by integrating the information obtained about the patterns with their moment of occurrence with respect to the neighboring vectors, while preserving the information about the collective neural activity. By using this simple approach of formatting the patterns, we extend a recent method that we have developed for the identification and characterization of neural patterns. We show that with the proper “correction” applied to pattern expression, one can use classical analysis methods, such as cross-correlation, to study the temporal coordination between neurons on longer timescales. This is obtained by reducing the blurring induced by convolution, thus enabling one to determine exactly where the pattern wavefronts emerge and to express pattern initiations as discrete events. The technique facilitates the subsequent application of classical spike-train analysis methods on these discrete events [3].

II. METHODS

A. Experimental Procedures

Neuronal data was recorded from the visual cortex of a head-fixed, anesthetized mouse, using Neuronexus silicon probes organized in tetrode configuration (4x2 tetrodes). Moving gratings with variable drifting directions (0 to 315°, in 45° increments) and fixed contrast were randomly presented on a display placed in front of the animal. Each of the 8 stimuli was presented 30 times in trials of 2.6 seconds duration. The vitals of the animals, such as heart rate, temperature, and respiration rate were monitored during the entire duration of the recording session. The parameters of the trials are presented in the following table (from [16]):

TABLE I: Drifting Gratings Visual Stimulation

Trial duration	2600 ms
Velocity of gratings	1 cycle/s
Focal distance of mouse eyes	15 cm
Mouse field of view	120 degrees
Optimal illumination range	10 - 60 lux
Minimum necessary frame rate	30 Hz
Optimal spatial frequency	0.1 cycle/degree

Visual responses were recorded from 8 tetrodes (32 channels) with high resolution sampling at 32 kSamples/s. Signals were filtered between 500 Hz and 3.5 kHz to

extract the multi-unit (MU) activity and their corresponding waveforms. Separately, they were filtered between 1 Hz and 300 Hz and down-sampled 32:1 to obtain the local-field potential (LFP). An off line spike-sorting method was finally used to extract 14 single units (SU) from these waveforms, together with other statistical data such as width and amplitude of cells, firing rate and standard deviation of waveforms. All the experimental procedures have been performed under the Society for Neuroscience guidelines, Romanian law, and the European Communities Council Directive 2010/63/EU regarding the care and use of animals for experimental procedures. Experiments were approved by the Local Ethics Committee (3/CE/02.11.2018) and the National Veterinary Authority (ANSVSA; 147/04.12.2018).

B. Data Analysis

The electrophysiological data recorded from the cortex using multi-electrode recording can be used to extract the simultaneous activity of multiple neurons. There are a variety of techniques employed to study the activity of neurons within a predefined window of time, such as the binning of spikes as in [11], [12], [13], or the convolution of the spikes with exponentially decaying kernels followed by the sampling of the resulting continuous signals as in [2], [5], [3], [10], [14] and [15]. Here, we employed the methods and the findings from [3], which, combined with a method of extracting the most important pattern from a range of consecutive patterns, would allow us to use cross-correlation analysis techniques beyond the range of kernels reflecting mere synchrony of spikes.

1) *From Spikes to Patterns*: As described in [3], spike-trains obtained from multiple neurons are stacked together, and converted into continuous signals by low-pass filtering the dichotomous signals using convolution with an exponentially decaying kernel (see Fig. 1). The time constants (τ) used by the kernel can range from small values ($\tau=1-5$ ms) reflecting the coincidence of cellular activation (synchrony between spikes), to large values ($\tau>100$ ms) corresponding to the firing-rate co-variation across neurons. The resulting continuous signals were then sampled in n -dimensional state vectors (where n is the number of neurons) at 1 kHz and passed through a K-Means clustering algorithm using standard Euclidian distance, initialized with $K=1000$ clusters and with a stopping criterion for the clustering iterations defined as an error $\epsilon < 0.01$. This clustering step was necessary in order to identify a finite, stereotypical set of model vectors that represent the expressed classes of patterns. Due to the convolution process, each element of the vector became a real number and, therefore the resulting activity vectors can take on an infinite number of states. Clustering solves this problem by mapping these states onto a finite alphabet of representative model patterns.

2) *Complexity of a Pattern*: One pattern can be viewed as a snapshot of the integrated activity of all the recorded neurons at a given moment in time. Each element of the vector represents the integrated (low-pass filtered) activity of a single neuron at that particular time instance. Fig. 1 presents the patterns as n -dimension vectors of gray scale elements ranging from white to black, where n is the number of neurons contributing to the pattern: white color corresponds to a value of zero - indicating no spikes in the

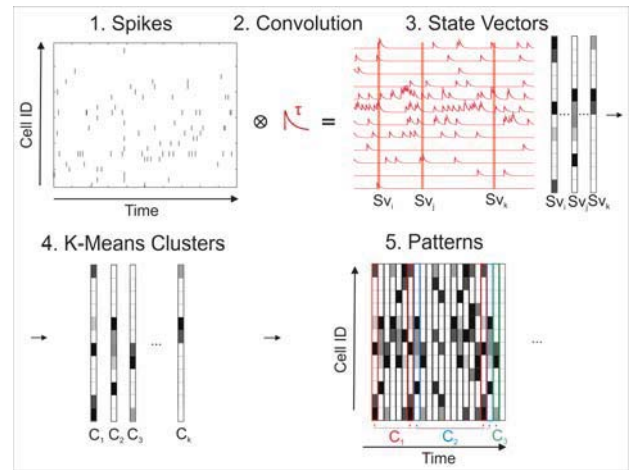


Fig. 1: From spikes to patterns: neuronal spikes obtained from offline spike-sorting (step 1) are converted to continuous signal by convolution (step 2). The resulted signals are sampled at 1 kHz (step 3) to obtain state vectors. A K-Means algorithm is used in step 4 to cluster the state vectors into patterns, which are replacing the original state vectors in dataset (step 5).

recent past, and black color corresponds to a large value - indicating a recent spike (the memory of the signal is given by the integration time-constant τ). The complexity of a pattern is given by the number of elements with values greater than a threshold of 0.36 ($1/e$), representing the ending value of a decaying signal of value 1.0 after τ ms. A neuron was considered *active* if it had a value greater than 0.36, and *silent* if its activation in the pattern was less than 0.36 [3]. Here, we considered only the timestamps of patterns with at least one active neuron in their components (*complexity* ≥ 1), and disregarded the patterns where all the neurons were silent.

3) *Auto and Cross-Correlation Histograms*: The cross-correlation between the spikes of two cells can be used to identify synchrony and oscillation events in neural signals [7]. This is done by measuring the dependence or association between two signals: a signal $g(t)$ is time shifted by different time lag values δ and compared with another fixed signal $f(t)$. The cross-correlation function describes the degree of correlation, can reveal oscillatory coupling, and evidences the delay between the signals. While the cross-correlation function is used for comparing two continuous signals, the cross-correlation histogram can be used to investigate the relationship between two binary signals, such as spikes or pattern's timestamps. Consider two binary signals $F(t)$ and $G(t)$ that can take values of 1 (when there is an event/spike/pattern at time t) or 0 (there is no event/spike/pattern at time t). The cross-correlation histogram (CCH) is then computed using the equation:

$$CCH_{X,Y}(\delta) = \sum_t^L X(t)Y(t + \delta)$$

where L is the length of the signal. The auto-correlation histogram (ACH) of a signal is computed by comparing the train of timestamps to a copy of itself, shifted with various time lags δ .

C. CCH/ACH on pattern expression is hindered by large time-constants τ

Classical analysis methods typically used in neuroscience can be applied to binary signals to characterize spike trains. A pattern representing the status of n neurons at an instance t of time has also discrete timestamps, so it can be classified as a binary event. That is to say, the analysis techniques that are used for spikes can also be applied to characterize the expression of patterns [3]. One of the most widely used methods for detecting synchrony and coupling between binary signals is the cross-correlation [8]. This technique works well for events that have a time lag between two consecutive timestamps, a condition that usually occurs, for example, when a pattern is obtained from neuronal spikes convolved with small time constants (Figure 2.A, top and bottom). On the other hand, if the spikes or patterns are firing one after the other at very small time intervals, the central peak of the cross-correlation histogram will have a "blinding" effect as the smearing can hide side-peaks reflecting oscillations [9]. In Figure 2.B, top and bottom, we show examples obtained on patterns computed with $\tau \geq 20$ ms. To understand why the blinding effect happens, we first superimposed the computed patterns over the signals of origin (see Figure 3.A): it can be observed that the same pattern C_i can be placed multiple times in consecutive order due to a small difference in the convolved signals from one timestamp to the next. The binary signal representing the timestamps of these similar patterns will have multiple consecutive expressions ("firings"), which translates in multiple consecutive 1's at times when the patterns were expressed, which generates wide, blurred peaks in CCH and ACH. This phenomenon results from the relatively slow decay of the convolution kernel. Another consequence of the slow decay of the kernel can be illustrated by two consecutive patterns C_k and C_{k+1} that share almost the same values for their active neurons, but the clustering algorithm had separated the two state vectors into different clusters based on the small distance between activity vectors. Between the two patterns, C_{k+1} is a less active version of C_k and, therefore, it contains no new information. However, when the CCH is computed between these two patterns (data not shown), the central peak will show an offset, which can be misinterpreted as neuronal interactions.

This "avalanche" of redundant patterns can be viewed as a limitation for slow changing signals, and it limits the usable time constants of the kernel to small values.

III. RESULTS

Here, we propose a method of going around this limitation by extracting the most significant pattern from trains of similar or redundant consecutive patterns. This can be simply done by combining the information already known about the patterns, such as their location, and the convolution values associated to their constitutive neurons.

Considering two consecutive patterns, we started by comparing the values that correspond to the same active neurons. The difference can help us determine whether we are moving away from a recent spike. The goal is to reconstruct the dataset of patterns keeping only the patterns that have in their components recently activated neurons. Figure 3.A presents the train of patterns originated from the clustering of convolved neuronal spikes using a large time

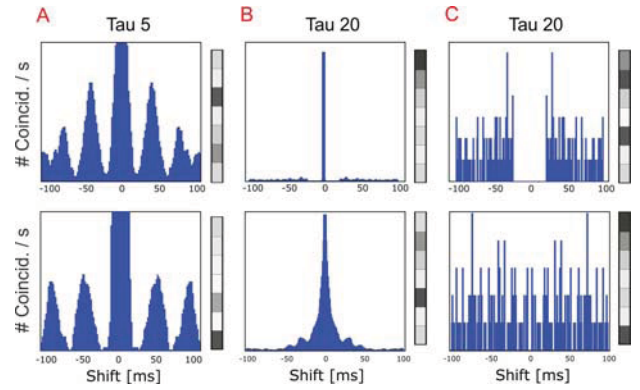


Fig. 2: Auto-correlations computed using short and long decay time constants τ . A: ACH for patterns obtained from convolution of spikes with $\tau = 5$ ms. B: ACH of patterns resulted from $\tau = 20$ ms. C: Two examples of ACH obtained from convolved spikes with $\tau = 20$ ms, after the trimming of redundant information

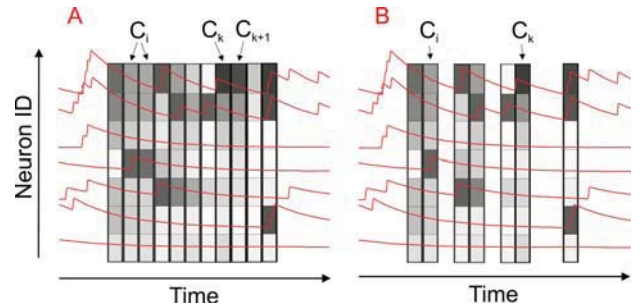


Fig. 3: Trimming down the redundant patterns. A: Detection of the most important/informative patterns and detection of redundant information. B: Train of patterns resulted after the extraction of redundant information. It mostly contains patterns with freshly fired neurons

constant τ . The "decimation" process of undesired patterns should start from left to right, that is, following the natural evolution of spikes and their corresponding decaying signals along the time axis. The next step is to find the identical patterns that sit one next to the other, similar to C_i in the figure, where i was the active neuron in both patterns, and keep only the first instance of such identical patterns. Then we compare the corresponding neuron dimensions of any two consecutive patterns C_k and C_{k+1} , where C_k is different than C_{k+1} and i is the active neuron in both patterns: having $C_{k_i} \geq C_{(k+1)_i}$ is an indication that we are moving away from a recently fired neuron, where C_{k+1} is the less active of the patterns and can be deleted.

After the last two patterns of the dataset are inspected, the result is a train of patterns that represents a clean version of the original (Figure 3.B): most of the redundant information is deleted, and the timestamps that are preserved contain only the time-flags of patterns with recently fired neurons.

After obtaining the clean dataset, the next step was to compare the original auto-correlation histogram of patterns to the auto-correlation histogram calculated using the trimming approach. We should expect a reduction of the central peak amplitude as a result of removal of redundant information. Figure 2.C shows such an example: the top row image unveiled a central gap in the auto-correlation

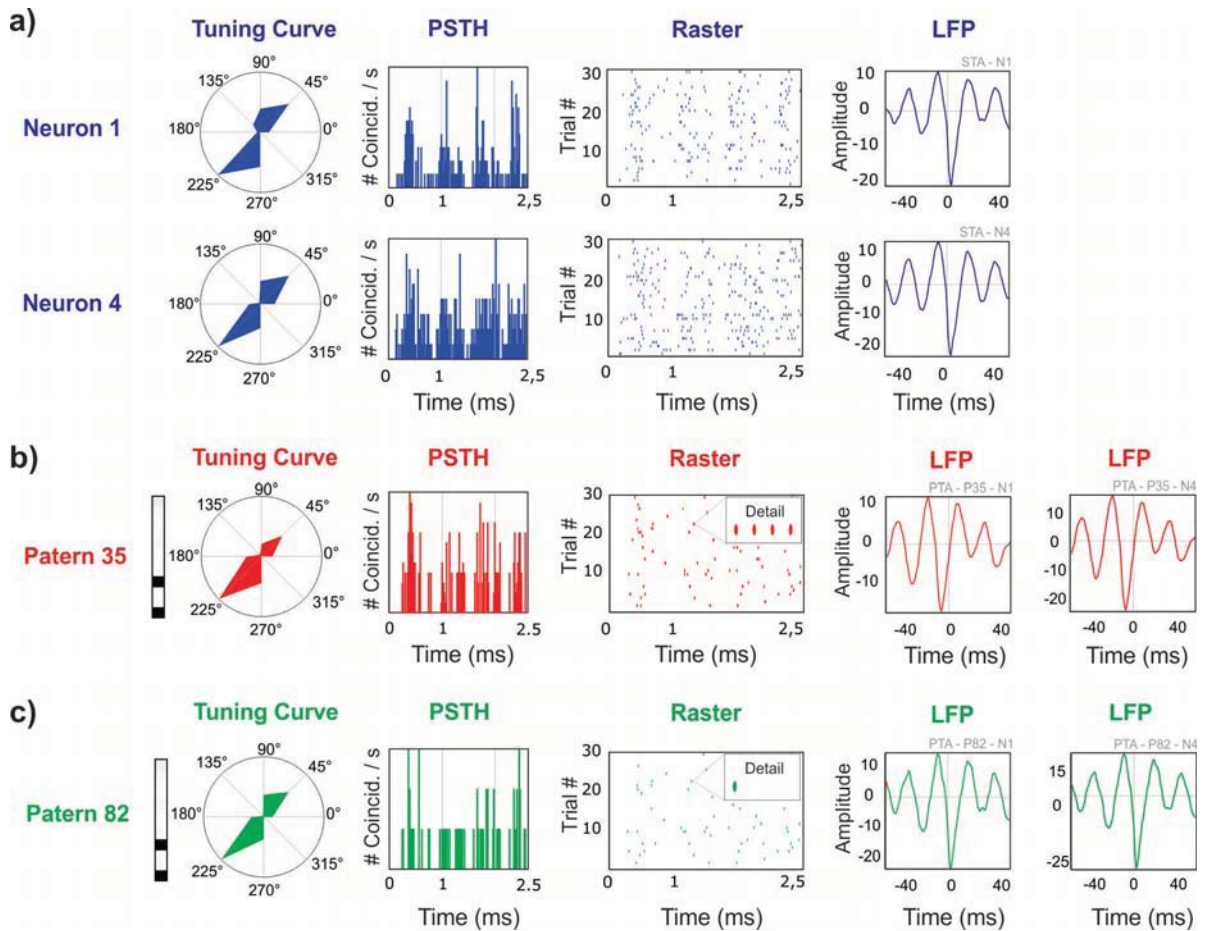


Fig. 4: Properties of neurons and patterns. a) In blue are two example of neurons and their stimulus tuning curves, peri-stimulus time histogram (PSTH), raster plot, and the associated local-field potential spike triggered average (STA). b) In red is a pattern with the two blue neurons from a) firing simultaneously, the PSTH, raster plot, and STA. c) In green we show the properties of the pattern from b) after removing its redundant expression with the algorithm described above.

function suggesting a refractory period on the activity of the patterns, and the bottom row presents a histogram that might indicate a bursting activity of the active neurons composing the pattern. Results suggest that the proposed technique is indeed useful for the study of temporal coordination between neurons at when larger timescales are used to integrate neural spikes.

As a control, we verified whether the removal of redundancy introduced any artificial or unexpected disturbance in the expression of patterns. Under certain conditions, patterns can inherit their tuning curves from the tuning curves of the composing neurons. Figure 4 shows in the left column (Tuning Curve) examples of neurons and their corresponding firing pattern before (Figure 4.b, Pattern 35, Tuning Curve) and after (Figure 4.c, Pattern 82, Tuning Curve) data manipulation using the proposed method. The tuning curves of the patterns recapitulate the tuning curves of component active neurons, indicating that the proposed data manipulation does not remove essential information about the properties of patterns.

Finally, we wanted to estimate the relation between firing patterns and the population dynamics reflected in the local-field potential (LFP), and to compare the results across the “raw”, the “trimmed” patterns, and their constituent neurons. In this particular case, similarly to tuning curves, the patterns should inherit the same time relation to the

LFP as that of their composing neurons. To this end, we computed the spike/pattern triggered average (STA) of the local field potential, thus quantifying their relation to the expression of neuron spikes/pattern timestamps (Figure 4.LFP). As shown in Figure 4.b.LFP - where an unprocessed or “raw” pattern is depicted in red, the redundancy of the pattern changed its temporal relation to the local field potential by slightly shifting the STA of the pattern to the left compared to the STA of the composing neuron. The “trimmed” pattern showed in green in Figure 4.c.LFP corrected this inconsistency and the STA became similar to that of the component neurons, indicating that the proposed method is also useful in correcting the STA shift due to pattern redundancy.

IV. COMMENTS AND CONCLUSION

Firing patterns of neurons can be identified reliably using clustering methods that determine stereotypical combinations of neuronal activations. To define the relation of neurons’ activities on different timescales (from fast, millisecond-, to slow, second-timescales), low-pass filtering is applied to spikes by convolution with exponentially decaying kernels. The latter process enables one to examine how information is encoded in multiple neural spike trains on various timescales.

One serious problem arises when slower convolution timescales are used, because resulting patterns become

redundant in time, i.e., they are expressed at multiple successive timescales due to the slow decay of the convolution kernel. This can be exacerbated by clustering procedures which assign the same identity to patterns whose distance in Euclidean space is small. The redundant expression of patterns "blind" the estimation of temporal structure in the data by smearing the side-peaks of the auto-correlograms. In addition, redundancy also has the effect of shifting the estimate of the time-relation between the LFP and the patterns, leading to misleading results.

Here, we devised a procedure which compares successive activity vectors and identifies active neurons in order to determine redundant pattern expression due to time-smearing induced by convolution with slow-decaying exponential kernels. We consider that the important information is stored in the patterns and the timestamps where the neurons fire, and the rest of the patterns are just diluted information about the former and can be eliminated. We showed that eliminating redundancy reveals the hidden information about the temporal coordination between neural populations and LFP and enables reducing the time-smearing in cross- and auto-correlation histograms, thus helping to identify time coordination between / within pattern expression.

We would like to stress out that in many cases when encoding on large timescales is investigated, reducing the integration time-constants to identify patterns is not an option. In such cases, integrating more information from the past to correlate it with a new firing of a component neuron requires larger integration time-constants. This has the effect of introducing redundancy in pattern expression estimation. The present method presents a solution to this problem.

A downside of the method is that the elimination of redundant information reduces the number of patterns on which the cross-/auto-correlation is applied (a statistical method), resulting in what appears as noisier correlation histograms. To compensate for this effect, other methods such as curve fitting should be used to better estimate time-relations using correlograms. Our results suggest that the major benefit of the proposed algorithm is that it allows a reliable estimate of the time relation between pattern expression and meso-scale circuit dynamics, reflected by the LFP.

ACKNOWLEDGMENT

This work was supported by: two grants from the Romanian National Authority for Scientific Research and Innovation, CNCS-UEFISCDI (Project Numbers PN-III-P4-ID-PCE-2016-0010 and COFUND-NEURON-NMDAR-PSY), and a National Science Foundation grant NSF-IOS-1656830 funded by the U.S. Government.

REFERENCES

- [1] Douglas and Martin (2004), *Neuronal circuits of the neocortex*. Annu. Rev. Neurosci. Vol. 27. <https://doi.org/10.1146/annurev.neuro.27.070203.144152>
- [2] Jurjut and Raul (2019), *A Color-Based Visualization Technique for Multielectrode Spike Trains*. Journal of Neurophysiology Vol 102. <https://doi.org/10.1152/jn.00758.2009>
- [3] Jurjut and Raul (2019), *Hold Your Methods! How Multineuronal Firing Ensembles Can Be Studied Using Classical Spike-Train Analysis Techniques*. Front. Syst. Neurosci. <https://doi.org/10.3389/fnsys.2019.00021>
- [4] Lloyd, S. P. (1982). *Least squares quantization in PCM*. IEEE Trans. Inf. Theory 28. <https://doi.org/10.1109/TIT.1982.1056489>
- [5] Jurjut, O. F., Nikolic, D., Singer, W., Yu, S., Havenith, M. N., and Muresan, R. C. (2011). *Timescales of multineuronal activity patterns reflect temporal structure of visual stimuli*. PLoS One. <https://doi.org/10.1371/journal.pone.0016758>
- [6] Kohonen, T. and Schroeder, M. R. and Huang, T. S. (2001). *Self-Organizing Maps*. Springer-Verlag, 3rd edition.
- [7] Raul C. Mureşan, Ovidiu F. Jurjut, Vasile V. Moca, Wolf Singer, and Danko Nikolic. (2008). *The Oscillation Score: An Efficient Method for Estimating Oscillation Strength in Neuronal Activity*, Journal of Neurophysiology. <https://doi.org/10.1152/jn.00772.2007>
- [8] D. Nikolic, R.C. Muresan, W. Feng, W. Singer (2012). *Scaled correlation analysis: a better way to compute a cross-correlogram*. European Journal of Neuroscience 35(5). <https://doi.org/10.1111/j.1460-9568.2011.07987.x>
- [9] D. Nikolic, V.V. Moca, W. Singer, R.C. Muresan. (2008). *Properties of multivariate data investigated by fractal dimensionality*. Journal of Neuroscience Methods 172(1). <https://doi.org/10.1016/j.jneumeth.2008.04.007>
- [10] Nikolic, D., Häusler, S., Singer, W., and Maass, W. (2007). *Temporal dynamics of information content carried by neurons in the primary visual cortex*. Proceedings of the NIPS 2006, Advances in Neural Information Processing Systems, Vol. 19, (Cambridge, MA: MIT Press).
- [11] Grün, S., Diesmann, M., and Aertsen, A. (2002). *'Unitary Events' in multiple single-neuron activity*. I. Detection and significance. Neural Comput. 14. <https://doi.org/10.1162/089976602753284455>
- [12] Schneidman, E., Berry, M. J. II, Segev, R., and Bialek, W. (2006). *Weak pairwise correlations imply strongly correlated network states in a neural population*. Nature 440.
- [13] Osborne, L. C., Palmer, S. E., Lisberger, S. G., and Bialek, W. (2008). *The neural basis for combinatorial coding in a cortical population response*. J. Neurosci. 28. <http://www.jneurosci.org/content/28/50/13522>
- [14] Gerstein, G. L., and Aertsen, A. M. (1985). *Representation of cooperative firing activity among simultaneously recorded neurons*. J. Neurophysiol. 54. <https://doi.org/10.1152/jn.1985.54.6.1513>
- [15] Häusler, S., and Maass, W. (2007). *A statistical analysis of information processing properties of lamina-specific cortical microcircuit models*. Cereb. Cortex 17. <https://doi.org/10.1093/cercor/bhj132>
- [16] Christopher M. Niell and Michael P. Stryker (2008), *Highly Selective Receptive Fields in Mouse Visual Cortex*. J Neurosci. 23. <https://doi.org/10.1523/JNEUROSCI.0623-08.2008>