

Received: 1 June 2019, Accepted: 8 June 2020

Edited by: C. Muravchik

Reviewed by: R. Cofre

Licence: Creative Commons Attribution 4.0

DOI: <https://doi.org/10.4279/PIP.120003>

ISSN 1852-4249

## Further results on why a point process is effective for estimating correlation between brain regions

I Cifre<sup>1,2\*</sup>, M Zarepour<sup>3,4</sup>, S G Horovitz<sup>5</sup>, S A Cannas<sup>3,4†</sup>, D R Chialvo<sup>2,4,6</sup>

Signals from brain functional magnetic resonance imaging (fMRI) can be efficiently represented by a sparse spatiotemporal point process, according to a recently introduced heuristic signal processing scheme. This approach has already been validated for relevant conditions, demonstrating that it preserves and compresses a surprisingly large fraction of the signal information. Here we investigated the conditions necessary for such an approach to succeed, as well as the underlying reasons, using real fMRI data and a simulated dataset. The results show that the key lies in the temporal correlation properties of the time series under consideration. It was found that signals with slowly decaying autocorrelations are particularly suitable for this type of compression, where inflection points contain most of the information.

---

\* [icifre@gmail.com](mailto:icifre@gmail.com)

† [sergiocannas@gmail.com](mailto:sergiocannas@gmail.com)

<sup>1</sup> Facultat de Psicologia, Ciències de l'educació i de l'Esport, Blanquerna, Universitat Ramon Llull, C. Císter 34, Barcelona, (08022), Spain.

<sup>2</sup> Center for Complex Systems & Brain Sciences (CEMSC<sup>3</sup>), Universidad Nacional de San Martín, 25 de Mayo 1169, San Martín, (1650), Buenos Aires, Argentina.

<sup>3</sup> Instituto de Física Enrique Gaviola (IFEG), Facultad de Matemática, Astronomía, Física y Computación, Universidad Nacional de Córdoba, Ciudad Universitaria, (5000), Córdoba, Argentina.

<sup>4</sup> Consejo Nacional de Investigaciones Científicas y Tecnológicas (CONICET), Godoy Cruz 2290, Buenos Aires, Argentina.

<sup>5</sup> National Institute of Neurological Disorders and Stroke, National Institutes of Health, Bethesda, MD, USA.

<sup>6</sup> Instituto de Ciencias Físicas (ICIFI). Escuela de Ciencia y Tecnología, Universidad Nacional de San Martín, 25 de Mayo 1169, San Martín, (1650), Buenos Aires, Argentina.

### I. Introduction

The large-scale dynamics of the brain exhibit a plethora of spatiotemporal patterns. An important methodological challenge is to define adequate coarse-graining of the brain imaging data which comprises thousands of the so-called BOLD (“blood oxygen level dependent”) time series. The usual analysis aims at identification of bursts of correlated activity across certain regions, which requires extensive computations, complicated in part by the large size of the data sets.

A decade ago it was discovered that this type of problem can be simplified efficiently by using only the timings of the peak amplitude signal events; i.e., converting the raw continuous signal into a point process (PP) [1–7]. Subsequent work using similar approaches [8–14] further confirmed that the method entails large compression of the original signals without significant loss of information. Overall, these findings not only suggest a way to speed up computations, but most importantly highlight the need to clarify which aspects or features of the brain imaging signals contain the most relevant information.

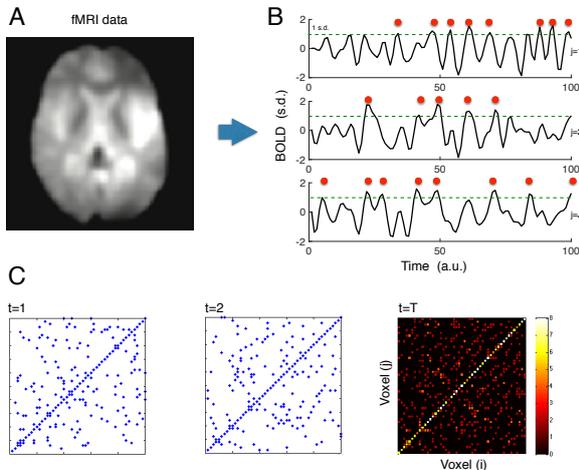


Figure 1: The basic aspects to consider in defining the point process of the BOLD signal. The traces on panel B are examples of raw time series ( $j$ ) of fMRI BOLD signals at three brain locations (called “voxels”). Time points are selected at the upward threshold crossings or the peaks of the signal (filled circles). The temporal co-occurrence of these points defines co-activation matrices for different lengths of time (graphs in C), which can be further averaged to estimate the correlation matrix for the entire time  $T$  of the system under study.

The present work is dedicated to identifying the reasons underlying the effectiveness of this approach. The results will show that the key lies in the temporal correlation properties of the time series under consideration: signals with temporal correlations are particularly suitable for this type of compression because inflection points contain most of the information.

The paper is organized as follows: in the next section the point process is defined and a simple example is presented. Section 3 discusses the main reason the point process works, emphasizing the relevance of the BOLD autocorrelations. This result is further tested in Section 4 using ground-truth simulated BOLD data in which the autocorrelation is altered. The paper closes with a brief discussion to summarize the relevance of the main conclusion.

## II. Definitions and examples

The basic steps that have been used [1–4] to define the point process in brain signals are summarized in Fig. 1. The raw data consist of time series recorded from the brain using functional magnetic resonance imaging (fMRI) corresponding to the activity of one of many thousands of small brain regions. It is accepted that this imaging technique measures a “blood oxygen level-dependent” signal (i.e., “BOLD”) in each small region, giving an estimation of the blood oxygen saturation, which itself is proportional to local neuronal activity.

The point process can be defined in different ways, but for the reasons discussed later, the end results are equivalent. As shown in Fig. 1, time points can be selected at the upward threshold crossings (here at unity) of the signal (filled circles). A second approach is to construct the point process by selecting the local peaks and/or valleys of the BOLD time series. For ease of discussion, we will deal with the second option in this paper. The temporal co-occurrence of the points defines the co-activation matrix (bottom graphs), which can be further averaged to estimate the correlation matrix of the system under study. Figure 2 illustrates, for those unfamiliar to the subject, three examples of typical BOLD time series that are usually recorded from the brain. From visual inspection it is already apparent that they are smooth traces, exhibiting temporal correlations, as will be further discussed later. There are also spatial correlations, for instance between the top two traces, which is evidenced in the images’ heat map between counter-lateral regions.

*A qualitative comparison of how well it works:* It has already been established, in different circumstances [2–4], that the co-activation matrix obtained with the point process method is very similar to the correlation matrix computed from the full (i.e., continuous) BOLD signal. Figure 3 shows an example of a correlation matrix constructed from the point process computed from a subject while resting (data fully described in [4]). The results demonstrate that as few as 4 points are already sufficient to define clusters of co-activation, as demonstrated previously in [2–4]. In addition, the results here show how de-activations (i.e., blueish colors) can also be evidenced by the PP approach.

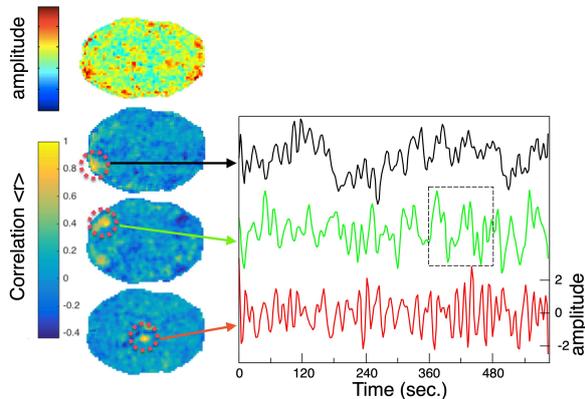


Figure 2: Right: Examples of typical BOLD time series from three selected sites in the brain. The dashed box in the middle time series indicates a portion of the signal used later in the analysis presented in Fig. 4. Left: Images correspond to a snapshot of activity amplitudes (top), the correlations between the three selected seeds (red dotted circles) and the rest of the brain (middle three images), and the corresponding brain structural slice at MNI coordinate  $z=10$  (bottom). The MNI  $x, y, z$  coordinates are  $-31 -95 10$ , respectively, for the top trace,  $43 -73 10$  for the middle trace and  $-10 -31 10$  for the bottom trace.

### III. Why does it work? A simple theory

As discussed previously, the example in Fig. 3 implies a large compression; the question then is *why a few points are enough to compute results similar to those obtained with the full signal*. A simple visual inspection of the BOLD traces reveals that the type of signals we are dealing with are temporally correlated. This is very well known; the neuronal activity is temporally and spatially correlated, and furthermore, the activity is convoluted by the hemodynamic transfer function which in itself introduces additional temporal correlations.

Therefore, for any time series with these properties, it seems natural to think that the most informative points are those in which its derivative changes sign. The other points are redundant since they can be predicted, to a certain degree, by a linear estimator. This is illustrated in Fig. 4, using as an example two minutes of BOLD recording (normalized by its standard deviation  $\sigma$ ). After setting a threshold  $\nu$ , the inflection points larger than a given  $\nu$  value are identified. These points constitute the marked point process in question.

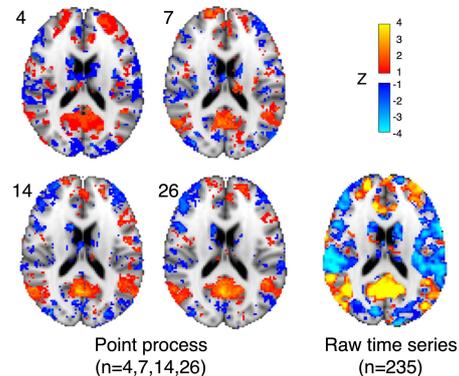


Figure 3: Example of correlation maps obtained from the raw BOLD time series of length  $n=235$  (right panel) and from the derived point process (left panels) for different numbers of points ( $n=4,7,14,26$ ). The left images represent, as “heat maps”, the co-activation of the seed (located at MNI coordinates  $x=4, y=-60, z=18$ ) with respect to each voxel. Note that a few points already suffice to identify well-defined clusters that are 1-4 standard deviations away from chance co-activations. The right panel corresponds to the Pearson correlation computed from the entire length of the raw BOLD time series. Red/blue colors label positive/negative point co-activations in the case of the left maps and positive/negative correlations in the case of the right map.

Now we ask how much of the raw signal is left out if these inflection points are used to extrapolate a piece-wise linear time-series. To answer this we analyze BOLD time series from the brain of a subject during an experiment in which fMRI data are collected at rest [3]. We proceed to compute the linear correlation between the two time series, the raw and the piece-wise linear one. In panels D and E the results are shown for different values of threshold  $\nu$  (in units of  $\sigma$ ) as well as for the correlation of the time series, estimated by the value of the first autocorrelation coefficient  $\gamma$ . Panel D shows that as the BOLD signal autocorrelation increases, the similarity between the piece-wise linear and the raw signals increases, evaluated in two ways: by the root mean squared error (**rmse**) and by the linear correlation  $\langle r \rangle$  between the two time series. As expected, raising the threshold  $\nu$  above zero produces an increasing loss of information about the signal, which is reflected in a monotonic increase in the **rmse** and a decrease in the  $\langle r \rangle$  values (see Panel E).

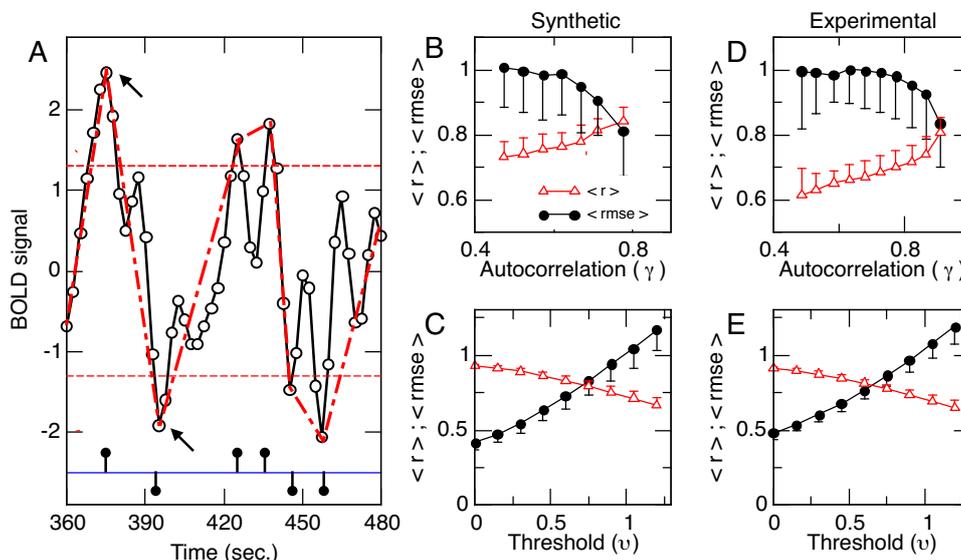


Figure 4: Why it works: The trace in Panel A is an example of a two-minute recording of a BOLD brain signal during rest (denoted by the dotted box in Fig. 2). The point process is defined by the timing of the peaks and valleys larger than a given threshold, (two are indicated here by arrows). The points, in this case, are only six (dots depicted in the bottom trace) out of the 120 samples of the original time series. The ability of these six points to preserve information about the original BOLD signal can be estimated by its similarity with a piece-wise linear time series (dashed red lines) constructed by joining the peaks and valleys. Horizontal dotted lines in Panel A denote the threshold for the example. Panels D and E correspond to the computed similarity between the BOLD signals and the piece-wise linear signals (evaluated by the correlation  $\langle r \rangle$  and  $\text{rmse}$  values) for different autocorrelation  $\gamma$  and threshold  $\nu$  values. Panels B and C correspond to similar calculations using synthetic time series. For Panels B and D  $\nu$  was fixed at 1. For Panels C and E  $\gamma$  was 0.85.

According to the present hypothesis, the functional dependence shown by the BOLD signals in Panels D and E will be replicated using synthetic signals with similar autocorrelation properties. To this end, we generate artificial time series with autocorrelation values identical to those of the BOLD signals, using the MATLAB routine `f_alpha_gaussian.m` from [16] (see also source codes at [17]). Panels B and C show that the behavior with respect to the threshold  $\nu$  and  $\gamma$  for the synthetic and empirical data are very similar.

The results show that the key to understanding why the approach works lies in the correlation properties of the time series under consideration. In synthesis, it is found that signals with long-range correlations are particularly suitable for this type of compression, where inflection points contain most of the information. The results also apply to other signals from any origin, as long as their autocorrelation features are similar.

#### IV. Further testing using synthetic ground-truth data.

The results in the previous section emphasize the relevance of the individual signal's autocorrelation as the main property related to the ability of a point process to preserve information about the original signal, and consequently, about functional connectivity between signals. We can test this by manipulating the autocorrelation in any given system for which the “ground-truth” crosscorrelations are known. The data reported by Smith et al. [19] can be used for our purpose.

The authors in [19] reviewed and compared the available fMRI analysis methods ranging from simple measures of pair-wise linear correlations to sophisticated multivariate approaches. In the process, they generated diverse and realistic simulated fMRI data sets describing different underlying networks. These simulations were based on the dy-

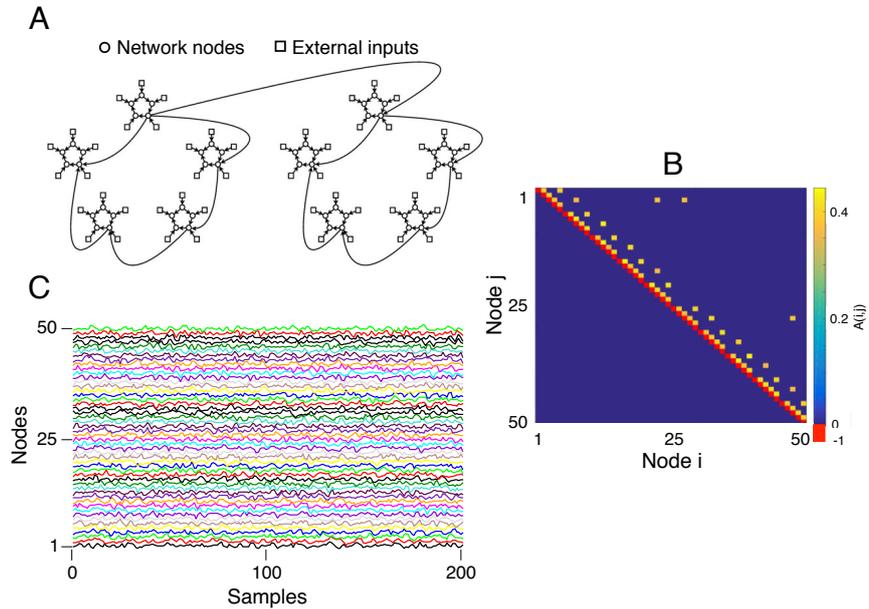


Figure 5: Simulated fMRI network from Ref. [19]. Panel A depicts the topology and Panel B the adjacency matrix of the interactions of the nodes, where negative values (labeled red) correspond to self-interactions. Connections are directed: a node in the upper diagonal of the matrix denotes a directed connection from a lower-numbered node to a higher-numbered one. Panel C shows an example of the BOLD time series simulated on this network.

dynamic causal modelling (DCM) [20] fMRI forward model (see [19] for full details). We used these simulated BOLD time series (downloaded from the authors' site [22]) to test the point process approach in comparison with standard correlation methods. First, for completeness, we briefly describe the essence of the model used in these simulations.

*The model:* Smith et al. simulations used a neural network model coupled with Buxton's nonlinear balloon model [21] for the vascular dynamics. Each neural network node has a binary external input (square symbols in the top diagram of Fig. 5). The state of the inputs (i.e., active or inactive) is given by a Poisson process which controls the probability of switching states, and can be seen as external signals or as noise at the neural level. Subsequently, the neural signals propagate across the network according to the DCM neural network model, where node interactions are defined by the  $A$  network matrix:

$$\dot{z} = \sigma Az + Cu \quad (1)$$

where  $z$  is the neural time series,  $\dot{z}$  is its rate of change,  $u$  are the external inputs and  $C$  the weights

controlling how the external inputs feed into the network (here just through the identity matrix). The off-diagonal terms in  $A$  determine the network connections between nodes (arrows in Fig. 5A), and the diagonal elements are all set to -1 to model within-node temporal decay; in this way,  $\sigma$  controls both the within-node (neural) temporal inertia and the temporal lag between nodes.

The time series of the activity for each node is then fed to a nonlinear balloon model for vascular dynamics [21] output, which is a function of the changing neural activity. The parameters were adjusted by the authors in order to match BOLD time series seen for typical data of brain resting activity recorded with 3 Tesla fMRI technology. Finally, various sources of variability were added to account for realistic expectations. The BOLD time series and the underlying ground-truth network matrices are accessible on the authors' site [22]. In the following paragraphs, we will explore the ability of the point process to extract the underlying network and to compare it with the commonly-used correlation method.

Figure 5 shows the ground-truth network con-

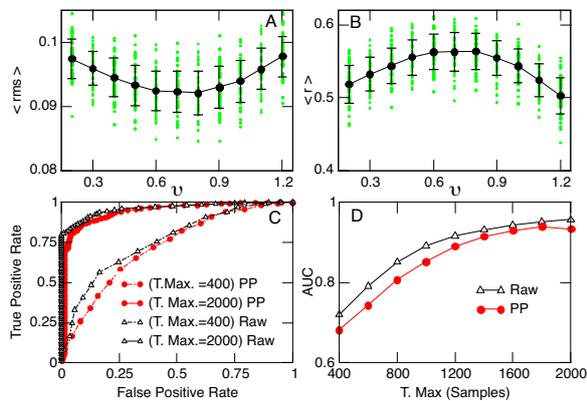


Figure 6: Panels A and B: Comparison of the partial correlation matrices –one calculated with the original BOLD data (Raw) and the other from its derived point process (PP)– as a function of the threshold  $\nu$  used to define the point process. Panel A corresponds to the root mean squared values of the differences, and panel B to the correlation coefficients. Dots correspond to individual results while averages (filled circles) and error bars correspond to the mean and standard deviation of the 50 simulated BOLD records (time series length  $T=200$  samples) Panel C shows the Receiver Operating Characteristic curve obtained for both methods in detecting the presence of links ( $\nu = 0.7$  and  $\sigma = 2.2$ ) computed for two time series of maximum length (T. Max. indicated in the legend). Panel D corresponds to the area under the ROC curve, computed for both methods as a function of the increasing maximum length, T. Max., of the time series considered.

sidered here (file `sim4.mat` downloaded from the site [22]). Panels A and B illustrate the topology and the adjacency matrix of the network ( $A$  in Eq. 1). It comprises 10 regular modules interconnected by a few links, a typical small-world graph; a total of 50 nodes interconnected by 61 positive off-diagonal interactions (40 of which correspond to nearest neighbors), as well as 50 negative (diagonal) self-interactions. Panel C shows typical traces of the computed BOLD time series recorded at the 50 nodes, simulating data from a fMRI typical session. The data set contains 50 stochastic realizations representing simulated fMRI records of different human subjects.

*Extracting the correlation graph:* To benchmark the relative merits of the point process approach, we first extracted the point process from the BOLD

time series, and then computed the covariance matrices for both the point process and the raw BOLD dataset. Following this, the partial correlation from both matrices was calculated (as in [19]), and their differences compared for various levels of the threshold  $\nu$  used to define the point process. As seen in Figs. 6A and 6B, there was an optimum threshold (for this dataset  $\sim \nu = 0.7$ ) at which the correlation matrices became more similar.

We then proceeded to check how well the method performed in predicting the underlying connectivity graph from the time series. Specifically, we checked how well the correlation matrices from both the point process and the raw BOLD dataset described the off-diagonal elements depicted in Fig. 5B, i.e., the synthetic ground-truth network. We used the receiver operating characteristic curve (ROC) [23], which benchmarks specificity and sensitivity as a function of a given parameter. To determine whether a connection is predicted or not between two nodes, we chose a threshold of  $2\sigma$  (corresponding to  $P = 95\%$ ) at the  $(i, j)$  partial correlation matrices entry. In general, to obtain the ROC curve a given range of relevant parameters must be explored while the true/false positive/negative predictions are counted. Here, for convenience, we chose to explore a range of time series lengths from relatively very short (set to 20 samples here) up to a variable maximum length, T.Max. For the examples presented, the values of T.Max. ranged from 400 to 2000 samples.

Figure 6C illustrates the results, where the family of curves (triangles for raw data and circles for the point process) corresponds to time series of various maximum lengths (T.Max.). As expected, the shortest T.Max (400 samples) gave the lowest confident results, while the longest T.Max. (2000 samples) resulted in a very good estimation of the true network connections. The area under the curve, plotted in Fig. 6D, is a good estimation of the relative goodness of the prediction, where a value of 1 corresponds to a perfect prediction and a value of 0.5 is equivalent to chance. Note that the main motivation for these numerical simulations is to demonstrate that there is *close similarity between the PP and Raw ROC curves* in Fig. 6D. This similarity is indicative of the good performance of the point process approach compared with the use of the raw time series.

## V. Discussion and conclusion

In this note, we revisited the heuristic point process approach originally introduced by Tagliazucchi *et al.* [1–3] to represent brain spatiotemporal dynamics in terms of the relatively high-amplitude inflection points of the BOLD signal. At the same time Caballero and colleagues [5–7] independently reported similar results, but these were based on de-convolution techniques. Later some extensions of the approach were presented by several authors [8–11].

*Why it works:* The present results show that the PP approach works due to a rather trivial fact: in any case of rather strongly autocorrelated signals, the most informative points are the inflection points; the remaining samples are more or less interpolated by straight lines (see Fig.4A) which can in principle, and for certain applications, be ignored. Thus, in general, it is expected that time series that exhibit slowly decaying temporal correlations will be particularly suitable for this type of approach.

*Relevance of the current results for functional connectivity of the brain:* Since its introduction, it has been suggested that the PP (or its variants) contains dynamical information, in the sense that it is potentially able to identify the timing and the location of fluctuating epochs of high correlations among brain regions. This identification has recently acquired relevance in the context of what is now dubbed “dynamical functional connectivity”, a very active area of research in the neuro-imaging community (see for instance the reviews by Keilholz *et al.* [25] and Iraj *et al.* [27]. In line with this, the recent report of Esfahlani *et al.* [26] emphasizes the fact that few *events* of co-activation can estimate the functional connectivity architecture of a system, a finding which is in full agreement with our arguments. Thus, it is very important to understand that behind all these reports there is a basic reason why these few points contain most of the information.

There is a lot of room for further investigation based on estimation of the correlation between these relatively large-amplitude inflection points. In particular, it seems a promising approach for inspection of non-stationarities in fMRI BOLD data, in certain pathologies that are known to exhibit bursts of non-stationarity, such as in Parkin-

son Disease syndrome and Tourette Disease syndrome. Typical of both cases is the existence of few epochs of coherent brain activity, which can be blurred if only the average functional connectivity is computed. Finally, it seems reasonable that the approach can be applied beyond brain research, to inspect similar problems in other fields.

*Acknowledgements* - The authors thank Steve M. Smith and Mark Woolrich (Imperial College, London) for sharing information on the Netsim package. This work was partially supported by NIH (U.S.) Grant No. 1U19NS107464-01. IC was supported by Ministerio de Economía, Industria y Competitividad (Spain) grant PSI2017-82397-R. MZ, DRC, & SAC were supported in part by CONICET (Argentina). DRC is grateful for the support of the Escuela de Ciencia y Tecnología, UNSAM.

- 
- [1] E Tagliazucchi, P Balenzuela, D Fraiman, D R Chialvo, *Brain resting state is disrupted in chronic back pain patients*, *Neurosci. Lett.* **485**, 26 (2010).
  - [2] E Tagliazucchi, P Balenzuela, D Fraiman, P Montoya, D R Chialvo, *Spontaneous BOLD event triggered averages for estimating functional connectivity at resting state*, *Neurosci. Lett.* **488**, 158 (2011).
  - [3] E Tagliazucchi, P Balenzuela, D Fraiman, D R Chialvo, *Criticality in large-scale brain fMRI dynamics unveiled by a novel point process analysis*, *Front. Physiol.* **3**, 15 (2012).
  - [4] E Tagliazucchi, M Siniatchkin, H Laufs, D R Chialvo, *The voxel-wise functional connectome can be efficiently derived from co-activations in a sparse spatio-temporal point-process*, *Front. Neurosci.* **10**, 381 (2016).
  - [5] N Petridou, C C Gaudes, L L Dryden, S T Francis, P A Gowland, *Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity*, *Hum. Brain Mapp.* **34**, 1319 (2013).

- [6] T W Allan, S T Francis, C Caballero-Gaudes, P G Morris, E B Liddle, P F Liddle, M J Brookes, P A Gowland, *Functional connectivity in MRI is driven by spontaneous BOLD events*, *PLoS ONE* **10**, 4 (2015).
- [7] F I Karahanoglu, D Van De Ville, *Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks*, *Nat. Commun.* **6**, 7751, (2015).
- [8] W Li, Y Li, C Hu, X Chen, H Dai, *Point process analysis in brain networks of patients with diabetes*, *Neurocomputing* **145**, 182 (2014).
- [9] X Liu, J H Duyn, *Time-varying functional network information extracted from brief instances of spontaneous brain activity*, *P. Natl. Acad. Sci. USA.* **110**, 4392 (2013).
- [10] X Liu, C Chang, J H Duyn, *Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns*, *Front. Syst. Neurosci.* **7**, 10 (2013).
- [11] J E Chen, C Chang, M D Greicius, G H Glover, *Introducing co-activation pattern metrics to quantify spontaneous brain network dynamics*, *NeuroImage* **111**, 476 (2015).
- [12] X Jiang, J Lv, D Zhu, T Zhang, X Hu, L Guo, T Liu, *Integrating group-wise functional brain activities via point processes*, *IEEE 11th International Symposium on Biomedical Imaging (ISBI)*, 669 (2014).
- [13] E Amico et al, *Posterior cingulate cortex-related co-activation patterns: a resting state fMRI study in propofol-induced loss of consciousness*, *PLoS ONE* **9**, 6 (2014).
- [14] G R Wu, W Liao, S Stramaglia, D R Ding, H Chen, D Marinazzo, *A blind deconvolution approach to recover effective connectivity brain networks from resting state fMRI data*, *Med. Image. Anal.* **17**, 365 (2013).
- [15] D Cordes, V Haughton, J D Carew, K Arfanakis, K Maravilla, *Hierarchical clustering to measure connectivity in fMRI resting-state data*, *Magn. Reson. Imaging* **20**, 305 (2002).
- [16] M Stoyanov, M Gunzburger, J Burkardt, *Pink noise,  $1/f^\alpha$  noise, and their effect on solutions of differential equations*, *Int. J. Uncertain. Quan.* **1**, 257 (2011).
- [17] <https://people.sc.fsu.edu/~jburkardt>
- [18] V M Eguiluz, D R Chialvo, G A Cecchi, M Baliki, A V Apkarian, *Scale-free brain functional networks*, *Phys. Rev. Lett.* **94**, 018102 (2005).
- [19] S M Smith, K L Miller, G Salimi-Khorshidi, M Webster, C F Beckmann, T E Nichols, J D Ramsey, M W Woolrich, *Network modelling methods for FMRI*, *NeuroImage*, **54**, 875 (2011).
- [20] K J Friston, L Harrison, W Penny, *Dynamic causal modelling*, *Neuroimage* **19**, 1273 (2003).
- [21] R Buxton, E Wong, L Frank, *Dynamics of blood flow and oxygenation changes during brain activation: the balloon model*, *Magn. Reson. Med.* **39**, 855 (1998).
- [22] <http://www.fmrib.ox.ac.uk/datasets/netsim/>
- [23] T Fawcett, *An introduction to ROC analysis*, *Pattern Recogn. Lett.* **27**, 861 (2006).
- [24] D Gutierrez-Barragan, M A Basson, S Panzeri, A Gozzi, *Infraslow State Fluctuations Govern Spontaneous fMRI Network Dynamics*, *Curr. Biol.* **29**, 2295 (2019).
- [25] S Keilholz, C Caballero-Gaudes, P Bandettini, G Deco, V Calhoun, *Time-Resolved Resting-State Functional Magnetic Resonance Imaging Analysis: Current Status, Challenges, and New Directions*, *Brain Connectivity*, **7**, 465 (2017).
- [26] F Z Esfahlani, Y Jo, J Faskowitz, L Byrge, D P Kennedy, O Sporns, R F Betzel, *High-amplitude co-fluctuations in cortical activity drive functional connectivity*, *bioRxiv* **800045**, (2020).
- [27] A Iraj, A Faghiri, N Lewis, Z Fu, S Rachakonda, V D Calhoun, *Tools of the trade: Estimating time-varying connectivity patterns from fMRI data*, *PsyArXiv* **1** (2020).