

Bounded-observation Kalman Filtering of Correlation in Multivariate Neural Recordings

MohammadMehdi Kafashan¹, Ben J. Palanca² and ShiNung Ching^{1,3}

Abstract—A persistent question in multivariate neural signal processing is how best to characterize the statistical association between brain regions known as functional connectivity. Of the many metrics available for determining such association, the standard Pearson correlation coefficient (i.e., the zero-lag cross-correlation) remains widely used, particularly in neuroimaging. Generally, the cross-correlation is computed over an entire trial or recording session, with the assumption of within-trial stationarity. Increasingly, however, the length and complexity of neural data requires characterizing transient effects and/or non-stationarity in the temporal evolution of the correlation. That is, to estimate dynamics in the association between brain regions. Here, we present a simple, data-driven Kalman filter-based approach to tracking correlation dynamics. The filter explicitly accounts for the bounded nature of correlation measurements through the inclusion of a Fisher transform in the measurement equation. An output linearization facilitates a straightforward implementation of the standard recursive filter equations, including admittance of covariance identification via an autoregressive least squares method. We demonstrate the efficacy and utility of the approach in an example of multivariate neural functional magnetic resonance imaging data.

I. INTRODUCTION

In neural signal processing, functional connectivity refers to the characterization of statistical association between brain regions [1]. Several spectrotemporal metrics may be used to determine this association. Functional connectivity intends to describe the extent to which disparate brain regions exhibit synchronized activity. It may, or may not, correspond to anatomical connections [2]. The outcome of any application of functional connectivity analysis amounts to a weighted graph, where the recorded regions are the nodes and the statistical associations constitute the edge weights. For electrophysiological recordings, increasing effort has been directed at elucidating these weights in a directed fashion (see e.g., [3], [4] and the references therein). Despite these

advances, the Pearson correlation coefficient – the simple, zero-lag correlation between time series – remains a familiar, reliable and widely used association metric, particularly in neuroimaging [5].

The Pearson correlation (hereafter referred to as the cross-correlation) is usually obtained over a single trial. Implicit in this use is stationarity of the underlying data. When trials are long, such an assumption is problematic. For instance, a ‘medium’ correlation value (e.g., 0.5), may arise between regions that are consistently associated across the trial, or, between regions that alternate between strong and weak association. In order to disambiguate these scenarios, one requires a way to track correlations on a finer time scale. Computing correlations on shorter sub-trial windows (i.e., binning) can provide such resolution at the expense of increased noise susceptibility. In this note, we propose a simple Kalman-filter based method to estimate the underlying dynamic evolution of correlation structure. The filter operates on the time series constructed by computing correlations in successive bins. We note that, while similar approaches have been developed for this purpose [6], they do not constrain the observation (i.e., the cross-correlation). As a consequence, the resulting estimate may lie outside of the limits $[-1, 1]$. In the formulation presented herein, we assume that each successive correlation arises from a standard autoregressive model, transformed through a Fisher transform (a normalizing nonlinearity that constrains the random walk between $[-1, 1]$). The Fisher transform is commonly used for performing correlation statistical significance testing [7].

Due to the nature of the nonlinearity, the problem reduces to a straightforward Kalman approach after an inversion step. We proceed by presenting the filter equations and demonstrating their efficacy in tracking dynamic correlations in synthetic networks. We show that the covariance matrices needed for implementation can be successfully obtained in a data driven way through the use of an autoregressive least squares method (ARLS). Finally, we demonstrate proof-of-concept by applying the method to an example of a high-dimensional neural data obtained via functional magnetic resonance imaging (fMRI). The overall method can be used as a front-end for dimensionality reduction methods for the purposes of classification and/or denoising.

The remainder of this note is organized as follows. In Section II we formulate the model and filter. In Section III we provide the relevant equations and provide simulation results, including use of ARLS. The example involving multivariate fMRI recordings is presented in Section III-D. Conclusions and Future work are discussed in Section IV.

S. Ching Holds a Career Award at the Scientific Interface from the Burroughs-Wellcome Fund. This publication was supported by the Foundation for Anesthesia Education and Research (FAER MRTG-CT-02/15/2010) and the Washington University Institute of Clinical and Translational Sciences grant UL1 TR000448 from the National Center for Advancing Translational Sciences. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

¹M. Kafashan is with the Department of Electrical and Systems Engineering, Washington University in St. Louis, 1 brookings drive, campus box 1042, MO 63130, USA. kafashan at ese.wustl.edu

²B. Palanca is with the Department of Anesthesiology, Washington University School of Medicine, Campus Box 8054, 660 South Euclid, St. Louis, MO 63110, USA.

³S. Ching is with Faculty of Electrical Engineering & the Division of Biology and Biomedical Sciences, Washington University in St. Louis, 1 brookings drive, campus box 1042, MO 63130, USA. shinung at ese.wustl.edu

II. MODEL FORMULATION

A. Dynamic Correlation Model

Let \mathbf{z}_k , $k = 1, 2, \dots$ be a vector time series of correlation values between pairs of brain regions. We assume that \mathbf{z}_k arises from an underlying state $\mathbf{x}_k \in \mathbb{R}^N$ via

$$\mathbf{z}_k = \mathbf{F}^{-1}(\mathbf{x}_k), \quad (1)$$

where the function \mathbf{F}^{-1} is vector-valued inverse Fisher transform which can be written as

$$\mathbf{F}^{-1} \left(\begin{bmatrix} x^1 \\ \vdots \\ x^N \end{bmatrix} \right) = \begin{bmatrix} \tanh(x^1) \\ \vdots \\ \tanh(x^N) \end{bmatrix}. \quad (2)$$

Such a function bounds the state within $[-1, 1]$ and is the standard transformation of the cross-correlation into a normally distributed variable[7]. We assume that \mathbf{x}_k obeys a standard, linear state space model of the form

$$\begin{aligned} \mathbf{x}_k &= \mathbf{A}_k \mathbf{x}_{k-1} + \mathbf{w}_{k-1}, \\ \mathbf{y}_k &= \mathbf{F}^{-1}(\mathbf{C}_k \mathbf{x}_k + \mathbf{v}_k). \end{aligned} \quad (3)$$

Here, the state noise process $\mathbf{w}_k \in \mathbb{R}^N$ is an N -dimensional random vector with multivariate Gaussian distribution having zero mean and covariance matrix $\mathbf{Q}_w \in \mathbb{R}^{N \times N}$, the noise vector $\mathbf{v}_k \in \mathbb{R}^N$ is a zero-mean multivariate Gaussian random vector with covariance matrix $\mathbf{R}_v \in \mathbb{R}^{N \times N}$. Bounded nature of correlation coefficients result in a non-Gaussian noise for the observation vector \mathbf{y}_k . The covariance \mathbf{Q}_w determines the extent to which correlation can change in successive measurements, while \mathbf{R}_v simply characterizes measurement noise. In this sense, the vector \mathbf{y}_k consists of noisy observations of pairwise correlation values. We assume that state noise and measurement noise are uncorrelated.

We seek the optimal filter for obtaining the state estimate $\hat{\mathbf{z}}_k$ in the sense of minimum mean-squared error (MMSE), i.e.,

$$\min_{\hat{\mathbf{z}}_k} \mathbb{E} \left[\|\hat{\mathbf{z}}_k - \mathbf{z}_k\|_2^2 \right]. \quad (4)$$

Central to this problem is the calculation of the probability density function (p.d.f.) of the state vector at any given time k , conditioned on $\mathcal{Y}_k = \{\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_k\}$, the set of all the past observations. The presence of the nonlinearity (2) complicates this calculation, but only slightly since \mathbf{F} is smooth and invertible. Thus, it follows immediately that we can obtain a surrogate measurement

$$\mathbf{d}_k = \mathbf{F}(\mathbf{y}_k), \quad (5)$$

such that \mathbf{d}_k is linear in the state \mathbf{x}_k . Note that (1) directly yields the p.d.f. of \mathbf{z}_k from that of \mathbf{x}_k . Thus, the problem (4) reduces to the classical Kalman filter to obtain the (Gaussian) p.d.f. of \mathbf{x}_k given \mathcal{Y}_k [8], [9], based on the measurements \mathbf{d}_k .

III. RESULTS

A. Filter Equations

It follows directly from (3)-(5) that the p.d.f. $p(\mathbf{y}_k | \mathbf{x}_k)$ can be written as

$$p(\mathbf{y}_k | \mathbf{x}_k) = \frac{1}{|\mathbf{J}_k|} p(\mathbf{v}_k = \mathbf{F}(\mathbf{y}_k) - \mathbf{C}_k \mathbf{x}_k), \quad (6)$$

where $p(\mathbf{v}_k = \mathbf{F}(\mathbf{y}_k) - \mathbf{C}_k \mathbf{x}_k)$ is p.d.f. of \mathbf{v}_k , a zero-mean multivariate Gaussian with covariance matrix \mathbf{R}_v . A standard Bayesian approach thus yields the posterior density $p(\mathbf{x}_k | \mathcal{Y}_k)$, which is Gaussian with covariance matrix (Σ) and mean (μ) as

$$\begin{aligned} \Sigma &= \left(\mathbf{P}_{k|k-1}^{-1} + \mathbf{C}_k^T \mathbf{R}_v^{-1} \mathbf{C}_k \right)^{-1}, \\ \mu &= \Sigma \times \left(\mathbf{P}_{k|k-1}^{-1} \hat{\mathbf{x}}_{k|k-1} + \mathbf{C}_k^T \mathbf{R}_v^{-1} \mathbf{F}(\mathbf{y}_k) \right). \end{aligned} \quad (7)$$

The Kalman update equations are thus

$$\begin{aligned} \hat{\mathbf{x}}_{k|k-1} &= \mathbf{A}_k \hat{\mathbf{x}}_{k-1|k-1}, \\ \mathbf{P}_{k|k-1} &= \mathbf{A}_k \mathbf{P}_{k-1|k-1} \mathbf{A}_k^T + \mathbf{Q}_w, \\ \hat{\mathbf{x}}_{k|k} &= \left(\mathbf{P}_{k|k-1}^{-1} + \mathbf{C}_k^T \mathbf{R}_v^{-1} \mathbf{C}_k \right)^{-1} \\ &\quad \times \left(\mathbf{P}_{k|k-1}^{-1} \hat{\mathbf{x}}_{k|k-1} + \mathbf{C}_k^T \mathbf{R}_v^{-1} \mathbf{F}(\mathbf{y}_k) \right), \\ \mathbf{P}_{k|k} &= \left(\mathbf{P}_{k|k-1}^{-1} + \mathbf{C}_k^T \mathbf{R}_v^{-1} \mathbf{C}_k \right)^{-1}. \end{aligned} \quad (8)$$

In the subsequent results, we assume an initial multivariate Gaussian prior $\mathbf{x}_{0|-1}$ with mean $\hat{\mathbf{x}}_{0|-1}$ and covariance of $\mathbf{P}_{0|-1}$. We note that, for the problem (4), (8) returns the optimal estimate. After obtaining state estimate $\hat{\mathbf{x}}_{k|k}$, the correlation coefficients $\hat{\mathbf{z}}_k$ are approximated as $\hat{\mathbf{z}}_k = \mathbf{F}^{-1}(\hat{\mathbf{x}}_{k|k})$.

B. Example: Simulation

To illustrate the proposed filter, we simulate a univariate example of the system (3) with $A = C = 1$, $Q_w = 0.1$ and $R_v = 0.05$. Fig. 1 shows the estimate of the correlation, i.e., z_k , using the proposed filter, as well as the results of naive application of the Extended Kalman Filter (EKF) based on linearization of (2) [10], [11]. Clearly, state estimation using proposed optimal filter is closest to the true state of the system. This result is expected because EKF is simply a linearization of $\mathbf{F}^{-1}()$, which will produce poor results when the argument of the function is large (i.e., large correlation values).

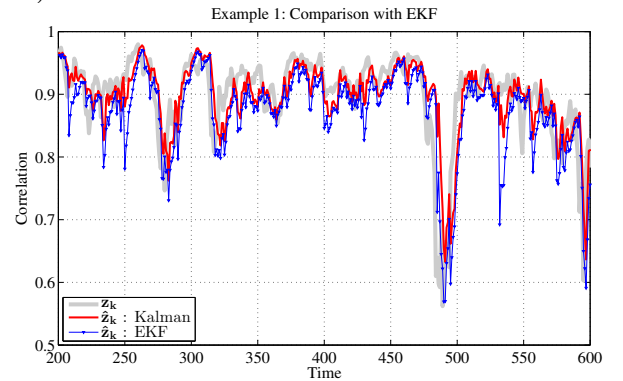


Fig. 1. **Example 1:** Estimation of the state in system (3) with $A = C = 1$, $Q_w = 0.1$ and $R_v = 0.05$ using optimal non-linear filter for bounded observation (red) and naive application of EKF (blue).




Fig. 2. Normalized MSE versus the ratio of process and measurement noise. Here, $r = \sigma_w^2/\sigma_v^2$ with $\mathbf{A} = \mathbf{C} = \mathbf{I}_4$, $\mathbf{Q}_w = \sigma_w^2 \mathbf{I}_4$ and $\mathbf{R}_v = \sigma_v^2 \mathbf{I}_4$.

Fig. 2 illustrates mean-square error (MSE) for the proposed optimal filtering compared to EKF for different noise parameter values. In this simulation, a 4-dimensional system with $\mathbf{A} = \mathbf{C} = \mathbf{I}_4$ is considered. The process and measurement noise are considered in the form of $\mathbf{Q}_w = \sigma_w^2 \mathbf{I}_4$, $\mathbf{R}_v = \sigma_v^2 \mathbf{I}_4$ and each time MSE is calculated for a specific ratio of $r = \sigma_w^2/\sigma_v^2$. It can be seen in this figure that by increasing r , MSE is decreased for both filters. Also, It can be seen that optimal filter always has better performance in term of MSE compare to EKF.

C. Covariance Estimation

The most significant obstacle in implementation of filters of the form (8) is selection of the covariance matrices \mathbf{Q}_w and \mathbf{R}_v . Here, we use the popular autoregressive least squares (ARLS)[12]–[14] method to compute, in a data-driven fashion, estimates of our filter covariances. In ARLS, a prior value is assigned for each of the covariances and the innovations of the subsequent filter is used to update the estimates of \mathbf{Q}_w and \mathbf{R}_v in the next iteration. More explicitly, assuming the the model (3) is time-invariant, we obtain an filtered estimate of the observations with non-exact covariance matrices. By computing the innovations and steady-state distributions, the problem of estimating \mathbf{Q}_w and \mathbf{R}_v then becomes a one-step least-squares optimization problem. In our case, the ALS solution is again facilitated by using the output linearized measurement (5).

Figure 3 shows the outcome of ALS when applied in a Monte Carlo ($n = 200$) fashion to the system described above. It is clear that the estimate covariances are in agree-




Fig. 3. Estimation of Q_w and R_v using auto-regressive least squares. Outcome of $n = 200$ Monte Carlo trials is shown, demonstrating concentration around the true value. System is as specified in the example of Section III-B.

ment of the true values.

D. Application to Multivariate Functional Magnetic Resonance Imaging (fMRI)

The Kalman-filtering approach outlined above is particularly useful under the assumption of edge independence, i.e., the matrix \mathbf{A} in (3) is the identity matrix. In this case, the n -dimensional update problem decomposes into n 1-dimensional problems, alleviating computational issues related to high-dimensionality. In brain imaging applications involving large numbers of recording sites, this yield a highly practical front-end correlation estimator.

To demonstrate the utility of the method, we apply it to an example of neural recordings obtained via fMRI in human volunteers¹. For this particular dataset, we concatenated trials to create a single long recording. The concatenation creates known time-points at which the temporal structure of the correlation would be expected to change. This particular dataset contains 160 brain regions, yielding 12720 correlation pairs. We bin the data over 5 datapoints (sampling rate of 0.45Hz), yielding an effective correlation window of around 10 seconds. We note, however, that the final result is largely insensitive to bin size.

Figure 4 illustrates how the method can yield particular insight into the temporal evolution of correlation structure. In Figure 4A, the average correlation map, calculated over the entire recording, is shown. This map highlights clear regions of interest that exhibit seemingly significant correlation structure. What this static map cannot discern is the extent to which these correlations may be changing over the course of the recording. By applying the ARLS-parametrized filter (8) to the data, we can obtain estimates of the correlation trajectories. These trajectories can be subsequently processed with any flavor of clustering or dimensionality reduction technique. In our case, we use the standard unsupervised clustering algorithm k -means [20]. A Bayesian information criterion (BIC) is used to determine the number of clusters returned by the k -means procedure [21]. A salient cluster of edges emerges that closely matches the high-correlation edges in the static map (Fig. 4B). Figure 4C plots the average trajectory of this cluster, where it is clearly seen that regions of ‘high’ correlation in the static map actually consist of epochs of very correlated activity, interspersed by epochs that are more weakly correlated. As noted above, in the case of this example, this temporal structure is expected as a consequence of the data concatenation described above. In this sense, the result demonstrates the capacity of the method to deliver meaningful spatiotemporal structure, even in the case of high dimensional data.

¹To demonstrate the utility of the method, we apply it to an example of fMRI data acquired from 15 human volunteers (Siemens 3T Trio, $4 \times 4 \times 4 \text{ mm}^3$ voxels, TR 2.2s, TE 27ms, flip angle = 90 degrees, 36 slices, 200 volumes/run). Two echo-planar BOLD imaging runs were acquired from each volunteer during quiet eyes-closed passivity (resting-state). Images underwent standard functional connectivity processing with regression of the mean whole brain signal [15], and motion censoring [16], [17]. BOLD signals were averaged within non-overlapping 5 mm radius spheres, using regions from prior studies [18], [19].

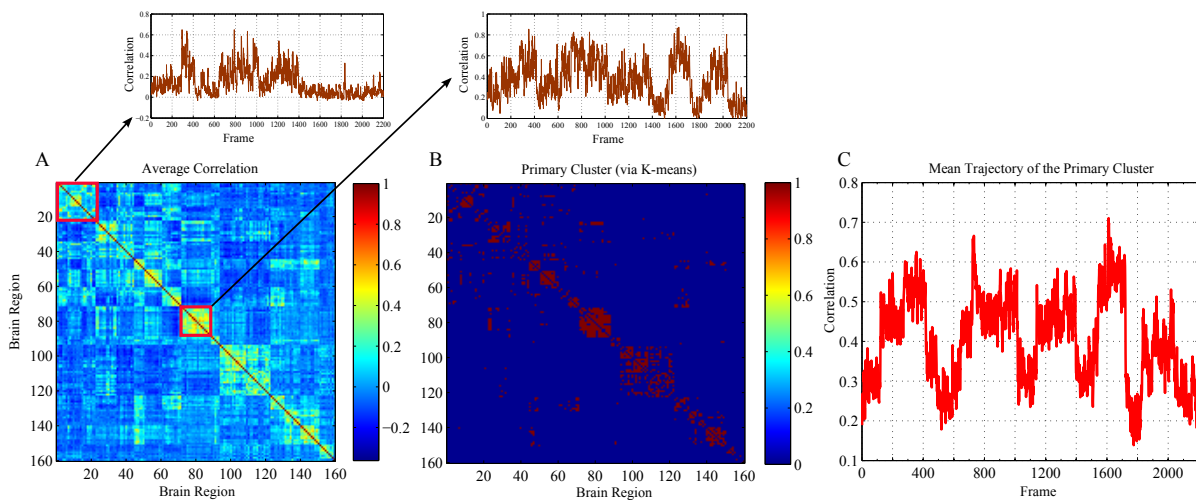


Fig. 4. **Disambiguating correlation dynamics.** (A) Static pairwise correlation for $n = 160$ channels of fMRI recordings. This ‘recording’ actually consisted of several concatenated datasets, thus creating artificial nonstationarity. Regions of interest are indicated by warmer colors. Application of the filtering procedure yields the temporal trajectory associated with each static correlation pair. The average of these trajectories over two regions of interest (indicated by red squares) are shown. (B) In total, we obtain 12720 pairwise correlation trajectories. Use of the standard clustering technique k -means yields a primary salient region of interest that closely matches the ‘warm’ regions in (A) (red indicates membership). (C) The mean of all trajectories represented in (B), illustrating the (expected) nonstationarity underlying the static characterization in (A).

IV. DISCUSSION & FUTURE WORK

This note presents a bounded-observation Kalman filter for estimating correlation dynamics in neural recordings. The metric of interest – the Pearson correlation coefficient – is constrained to $[-1, 1]$, necessitating the inclusion of a Fisher transform in the filter equations. Inversion of data to create a surrogate measurement facilitates straightforward derivation of the recursive filter equations and data-driven parametrization via the autoregressive least squares technique.

Under assumptions of independence, the consequent solution can be tractably obtained, even in the case of high dimensional data. We apply it to an example of fMRI data, showing that it can yield meaningful trajectories. We note that the results are relatively insensitive to the size of the sub-windows used to obtain the correlation time series. In future work, we will explore more principled ways of determining this window size. We note, unsurprisingly, that in the absence of filtering, the raw time series are exceedingly noisy and difficult to visualize.

The method outlined here can serve as an efficient front-end to downstream classification algorithms for extracting salient spatiotemporal dynamics and/or rejecting artifacts and noise. In future work, we plan to deploy the method in this capacity in a more systematic study of correlation dynamics in several neuroimaging recording modalities.

REFERENCES

- [1] O. Sporns, G. Tononi, and G. M. Edelman, “Connectivity and complexity: the relationship between neuroanatomy and brain dynamics.” *Neural Networks*, vol. 13, no. 8-9, pp. 909–922, 2000.
- [2] J. S. Damoiseaux and M. D. Greicius, “Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity,” *Brain Struct. Funct.*, vol. 213, no. 6, pp. 525–533, 2009.
- [3] M. Kamiński, M. Ding, W. A. Truccolo, and S. L. Bressler, “Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance,” *Biological Cybernetics*, vol. 85, no. 2, pp. 145–157, 2001.
- [4] K. J. Friston, L. Harrison, and W. Penny, “Dynamic causal modelling,” *NeuroImage*, vol. 19, no. 4, pp. 1273–1302, 2003.
- [5] K. R. Van Dijk, M. R. Sabuncu, and R. L. Buckner, “The influence of head motion on intrinsic functional connectivity MRI,” *NeuroImage*, vol. 59, no. 1, pp. 431–438, 2012.
- [6] J. Kang, L. Wang, C. Yan, J. Wang, X. Liang, and Y. He, “Characterizing dynamic functional connectivity in the resting brain using variable parameter regression and kalman filtering approaches,” *Neuroimage*, vol. 56, no. 3, pp. 1222–1234, 2011.
- [7] H. Hotelling, “New light on the correlation coefficient and its transforms,” *J. R. Statist. Soc. B*, vol. 15, no. 2, pp. pp. 193–232, 1953.
- [8] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.
- [9] —, “Contributions to the theory of optimal control,” *Bol. Soc. Mat. Mexicana*, vol. 5, no. 2, pp. 102–119, 1960.
- [10] Z. Chen, “Bayesian filtering: From kalman filters to particle filters, and beyond,” *Statistics*, vol. 182, no. 1, pp. 1–69, 2003.
- [11] G. A. Einicke, “Smoothing, filtering and prediction-estimating the past, present and future,” *New York: InTech*, 2012.
- [12] B. J. Odelson, M. R. Rajamani, and J. B. Rawlings, “A new autocovariance least-squares method for estimating noise covariances,” *Automatica*, vol. 42, no. 2, pp. 303–308, 2006.
- [13] B. M. Åkesson, J. B. Jørgensen, N. K. Poulsen, and S. B. Jørgensen, “A tool for kalman filter tuning,” *Comput.-Aided Chem. Eng.*, vol. 24, pp. 859–864, 2007.
- [14] —, “A generalized autocovariance least-squares method for kalman filter tuning,” *J. Process Control*, vol. 18, no. 7, pp. 769–779, 2008.
- [15] M. D. Fox, D. Zhang, A. Z. Snyder, and M. E. Raichle, “The global signal and observed anticorrelated resting state brain networks,” *J. Neurophysiol.*, vol. 101, no. 6, pp. 3270–3283, 2009.
- [16] J. D. Power, A. Mitra, T. O. Laumann, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, “Methods to detect, characterize, and remove motion artifact in resting state fMRI,” *NeuroImage*, vol. 84, pp. 320–341, 2014.
- [17] C. D. Smyser, A. Z. Snyder, and J. J. Neil, “Functional connectivity MRI in infants: exploration of the functional organization of the developing brain,” *NeuroImage*, vol. 56, no. 3, pp. 1437–1452, 2011.
- [18] J. D. Power, A. L. Cohen, S. M. Nelson, G. S. Wig, K. A. Barnes, J. A. Church, A. C. Vogel, T. O. Laumann, F. M. Miezin, B. L. Schlaggar *et al.*, “Functional network organization of the human brain,” *Neuron*, vol. 72, no. 4, pp. 665–678, 2011.
- [19] M. E. Raichle, “The restless brain,” *Brain Connectivity*, vol. 1, no. 1, pp. 3–12, 2011.
- [20] J. Hartigan, *Clustering Algorithms*, ser. Wiley Series in Probability and Mathematical Statistics. New York, 1975.
- [21] C. Goutte, L. K. Hansen, M. G. Liptrot, and E. Rostrup, “Feature-space clustering for fMRI meta-analysis,” *Human brain mapping*, vol. 13, no. 3, pp. 165–183, 2001.