Blind identification of overlapping communities from nodal observations

Wijnands, Ruben; Leus, Geert; Hunyadi, Borbála

DOI
10.23919/EUSIPCO63174.2024.10715311

Publication date
2024

Document Version
Final published version

Published in
32nd European Signal Processing Conference, EUSIPCO 2024 - Proceedings

Blind identification of overlapping communities from nodal observations

Ruben Wijnands, Geert Leus, Borbála Hunyadi
Department of Microelectronics, *Delft University of Technology*, Delft, The Netherlands
{R.Wijnands,G.J.T.Leus,B.Hunyadi}@tudelft.nl

***Abstract*—Identifying overlapping communities from data is crucial for grasping the complex structure and dynamics of networks, amongst others in fields such as computational neuroscience. Research using fMRI has demonstrated that brain regions can change their functional network membership over time using temporal independent component analysis (tICA). However, reproducibility of such overlapping communities remains a challenge. Recently, several alternative approaches have been proposed to identify such overlapping communities. While results are promising, less is known about the model and assumptions that underlie these approaches. This paper shows that the bilinear model, combined with the assumption of quasi-stationary and uncorrelated sources, underlies novel methods for identifying overlapping brain networks. Furthermore, we propose a new algorithm, and through simulations, we investigate the robustness of our algorithm and several existing methods to solve the problem in noisy conditions with few available data samples. We conclude that quasi-stationary blind source separation-based techniques can have a promising advantage over tICA in terms of identifiability of overlapping communities and thus have the potential to contribute towards greater reproducibility of results.**

***Index Terms*—Dynamic functional connectivity, overlapping communities, blind source separation, canonical polyadic decomposition, block-term decomposition**

I. Introduction

Blind source separation (BSS) techniques such as independent component analysis (ICA) have been frequently employed in brain network research. While the popularity of spatial ICA (sICA) has dominated temporal ICA (tICA) in functional magnetic resonance imaging (fMRI) studies, the work of [1] demonstrated that brain regions can be part of multiple networks that are temporally independent. However, the reproducibility of such overlapping communities remains challenging due to either a lack of temporal samples or strong non-Gaussianity assumptions on the source distributions. Nevertheless, [2] reached similar conclusions using a state-based analysis of stacked covariance matrices, illustrating that brain regions can indeed shift their network membership over time. While state-based methods thus allow for brain regions with changing network membership and capture valuable insights into the dynamics of brain region network memberships, it has been shown that the underlying assumption of having only one particular active network within a data window does not hold [3].

This work is part of the GraSPA project (project 19497 within the TTW OTP programme), which is financed by the Netherlands Organization for Scientific Research (NWO).

Following these developments, another line of research started investigating spatial patterns or communities that underlie the functional connectivity network matrices, including but not limited to the often-used covariance and Pearson correlation coefficient matrices. Several works [4]–[7] focus on stacked functional connectivity matrices that are obtained either by using a sliding window on subject-specific data or are gathered from multi-subject static functional connectivity matrices. Either way, the goal is to jointly identify communities in the brain that shape the observed connectivity patterns and infer the temporal activity information associated with an identified community. To name a few, [4] uses a tensor decomposition method to identify communities in the brain. Furthermore, sparse basis learning methods have been proposed in [5] that incorporate a rank-1 assumption on the communities underlying the functional connectivity matrices. Finally, instead of the last mentioned rank-1 assumption, [6], [7] impose a low-rank assumption on the underlying communities. We note that, in essence, the methods described above solve a BSS problem. However, up until now, the relation between the two-way matrix factorization model and the methods described above remains unclear. More importantly, the assumptions that underlie novel overlapping community identification methods need to be made more explicit.

In this paper, our contributions are as follows. First, we show the generative model with corresponding assumptions that form the base of state-of-the-art methods for overlapping community identification. Consequently, we provide insight into the physical interpretation of estimated parameters, thereby enhancing the interpretability of findings in the literature. Furthermore, we propose a new algorithm and conduct simulations to evaluate the performance of several methods that can solve the BSS problem. Through these simulations, we also investigate if these methods can be advantageous compared to tICA in noisy domains with fewer available data samples.

We adopt the following notation, based on matrix $\mathbf{X} = [\mathbf{x}_1, ..., \mathbf{x}_T]$. We use $\mathrm{vec}(\cdot)$ for the vectorization operator such that $\mathrm{vec}(\mathbf{X}) = [\mathbf{x}_1^\mathsf{T}, ..., \mathbf{x}_T^\mathsf{T}]^\mathsf{T}$; $\mathrm{vec}^{-1}(\cdot)$ denotes the inverse operation of $\mathrm{vec}(\cdot)$; $\mathrm{diag}(\cdot)$ takes the diagonal elements of a matrix such that $\mathrm{diag}(\mathbf{X}) = [x_{1,1}, ..., x_{T,T}]^\mathsf{T}$; $\mathrm{Diag}(\cdot)$ creates a diagonal matrix such that matrix $\mathrm{Diag}(\mathbf{x}_1)$ contains $\mathbf{x}_1$ on its diagonal; $\odot$ denotes the Khatri-Rao product; $\otimes$ denotes the Kronecker product; $\circ$ denotes the outer product; $\circledast$ denotes the Hadamard product.

II. Generative signal model

We start with the following noiseless bilinear data model

$$\mathbf{X} = \mathbf{AS}, \tag{1}$$

that decomposes the data $\mathbf{X} = [\mathbf{x}_1, ..., \mathbf{x}_T] \in \mathbb{R}^{N \times T}$ into low-rank factor matrices. Here, N and T are the number of nodes and time points, respectively. Furthermore, the community matrix $\mathbf{A} = [\mathbf{a}_1, ..., \mathbf{a}_R] \in \mathbb{R}^{N \times R}$ encodes the memberships of nodes to R functional networks or communities. Then, each row in the source matrix $\mathbf{S} = [\mathbf{s}_1, ..., \mathbf{s}_R]^\mathsf{T} \in \mathbb{R}^{R \times T}$ contains the temporal information of each column in $\mathbf{A}$. We absorb any noise or artifact sources into the signal model and aim to identify either the sources of interest or noise components because modeling noise sources present in neuroimaging data is hard, if not impossible. After blind identification of $\mathbf{A}$, we thus resort to ad-hoc methods for selecting sources of interest versus noise components. Also, we limit ourselves to the case of over-determined mixtures, i.e., $N > R$. In literature, $\mathbf{A}$ has been referred to as a mixing matrix, node-to-cluster or node-to-community assignment matrix, spatial maps [8], temporal functional modes (TFMs) [1], or basis vectors [5]. Note that $\mathbf{A}$ could also be viewed as a weighted hypergraph incidence matrix, where a source signal $\mathbf{s}_r$ describes the activity of nodes in a corresponding hyperedge $\mathbf{a}_r$.

Now, following the work of [9], we perform blind source separation based on the assumption that the sources are quasi-stationary. By quasi-stationary sources (QSS), we mean that the second-order statistics of the sources do not vary within a short time window of length L but vary from window to window. Under this assumption, we compute multiple sample covariance matrices in a sliding-window fashion:

$$\mathbf{R}_i = \frac{1}{L}\mathbf{X}_i\mathbf{X}_i^\mathsf{T} = \frac{1}{L}\mathbf{A}\mathbf{S}_i\mathbf{S}_i^\mathsf{T}\mathbf{A}^\mathsf{T}, \tag{2}$$

where windowed data $\mathbf{X}_i = [\mathbf{x}_{o(i-1)L+1}, ..., \mathbf{x}_{o(i-1)L+L}]$ is assumed zero mean, and $\mathbf{S}_i = [\mathbf{s}_{o(i-1)L+1}, ..., \mathbf{s}_{o(i-1)L+L}]$ represents the sources within a window. In this context, $i \in [1, ..., T_\mathrm{w}]$, where T_w represents the total number of time windows, and $o \in (0, 1]$ denotes an overlap factor. For a single window i, we vectorize the sample covariance matrix $\mathbf{R}_i$ and assume the sources are uncorrelated[1]:

$$\begin{aligned}
\mathbf{y}_i &= \mathrm{vec}(\mathbf{R}_i) \\
&= \mathrm{vec}(\frac{1}{L}\mathbf{A}\mathbf{S}_i\mathbf{S}_i^\mathsf{T}\mathbf{A}^\mathsf{T}) \\
&= (\mathbf{A} \otimes \mathbf{A})\mathrm{vec}(\frac{1}{L}\mathbf{S}_i\mathbf{S}_i^\mathsf{T}) \\
&= (\mathbf{A} \odot \mathbf{A})\mathrm{diag}(\frac{1}{L}\mathbf{S}_i\mathbf{S}_i^\mathsf{T}) \\
&= (\mathbf{A} \odot \mathbf{A})\mathbf{p}_i.
\end{aligned} \tag{3}$$

Note that each entry in $\mathbf{p}_i$ represents the estimated signal power of a source within the time window. Finally, we horizontally stack the vectorized sample covariance matrices:

$$\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_{T_\mathrm{w}}] = (\mathbf{A} \odot \mathbf{A})\mathbf{P}. \tag{4}$$

Now, from $\mathbf{Y}$, the goal is to identify (possibly overlapping) communities encoded in the columns of $\mathbf{A}$, assuming quasi-stationary and uncorrelated sources $\mathbf{s}_r$, without knowing $\mathbf{P} = [\mathbf{p}_1, ..., \mathbf{p}_{T_\mathrm{w}}] \in \mathbb{R}^{R \times T_\mathrm{w}}$.

A. Relation to tensor methods

As previously mentioned, several works implicitly work with the bilinear model and corresponding assumptions described above. We now show that (4) is equivalent to the canonical polyadic decomposition (CPD) of a tensor with stacked covariance matrices, used in [4], [5]. First, let us rewrite (4) as:

$$\mathbf{Y} = (\mathbf{A} \odot \mathbf{A})\mathbf{P} = \sum_{r=1}^{R}(\mathbf{a}_r \otimes \mathbf{a}_r) \circ \overline{\mathbf{p}}_r, \tag{5}$$

where $\overline{\mathbf{p}}_r$ denotes a row of $\mathbf{P}$, which we call a power time course that is associated with a community $\mathbf{a}_r$. If we apply the $\mathrm{vec}^{-1}(\cdot)$ operation on each vectorized sample covariance matrix:

$$\begin{aligned}
\mathbf{R}_i &= \mathrm{vec}^{-1}(\mathbf{y}_i) \\
&= \sum_{r=1}^{R}\mathrm{vec}^{-1}((\mathbf{a}_r \otimes \mathbf{a}_r)p_{r,i}) \\
&= \sum_{r=1}^{R}(\mathbf{a}_r \circ \mathbf{a}_r)p_{r,i},
\end{aligned} \tag{6}$$

we arrive at the rank-1 assumption on the components that shape the functional connectivity matrices made in [5]. Finally, by stacking the covariance matrices into a tensor $\mathcal{R} \in \mathbb{R}^{N \times N \times T_\mathrm{w}}$, we observe that the made assumptions on (1) result in a CPD of a tensor composed of stacked covariance matrices, used in [4]:

$$\mathcal{R} = \sum_{r=1}^{R}\mathbf{a}_r \circ \mathbf{a}_r \circ \overline{\mathbf{p}}_r. \tag{7}$$

Furthermore, the low-rank assumption in the works [6], [7] corresponds to the assumption that multiple uncorrelated sources have the same power time courses, represented by matrix $\mathbf{C} \in \{0, 1\}^{R \times R_\mathrm{c}}$, containing only a single 1 on each row, and $\mathbf{P} \in \mathbb{R}^{R_\mathrm{c} \times T_\mathrm{w}}$:

$$\mathbf{Y} = (\mathbf{A} \odot \mathbf{A})\mathbf{CP}. \tag{8}$$

By following the steps in (5)-(7), (8) can be proven equivalent to the block-term decomposition (BTD) with rank R_c.

[1]For simplicity, we assume that T_w is large enough to consider a diagonal source sample covariance matrix.

III. METHODS

In the remainder of this work, we aim to identify a suitable method for accurately identifying overlapping networks and investigate the necessity of data pre-whitening.

Several methods are listed in the literature that solve the overdetermined BSS-QSS problem, which can be grouped as joint diagonalization (JD) and CPD (a.k.a. PARAFAC or CANDECOMP) based approaches. We will consider the case of having a large number of nodes N, which is a realistic scenario when analyzing neuroimaging data on the voxel level or data segmented using detailed brain atlases. For such cases, pre-whitening of the data can be beneficial (and often necessary) in terms of memory requirements, decreasing from $\mathcal{O}(N^2 T_{\text{w}})$ to $\mathcal{O}(R^2 T_{\text{w}})$. Moreover, we specifically limit ourselves to methods that do not impose orthogonality assumptions on $\mathbf{A}$ or assume $\mathbf{R}_i$ to be positive-definite. We identified the pre-whitened alternating projections algorithm (PAPA) [9], [10], the tri-linear alternating least squares (TALS) algorithm [11], and the fast Frobenius diagonalization (FFDiag) algorithm [12] as potential candidate algorithms to solve the BSS-QSS problem.

We will begin by describing the pre-whitening procedure. Then, we present the different algorithms we will consider. We start with non-negative TALS (NN-TALS), which consists of TALS under a non-negativity constraint on $\mathbf{P}$. We also propose a pre-whitened version of this algorithm, labeled as pre-whitened non-negative TALS (P-NN-TALS). Then, we detail the PAPA, which is also based on the pre-whitened matrices. We label this algorithm as pre-whitened alternating projections (P-AP). The last algorithm we briefly outline is a pre-whitened version of FFDiag (P-FFDiag). Finally, we will argue for using covariances over Pearson correlation coefficients in [4], [6].

A. Pre-whitening

In [9], a pre-whitening matrix is found by computing the square-root factorization of the time-averaged covariance matrix:

$$\overline{\mathbf{R}} = \frac{1}{T_{\text{w}}} \sum_{i=1}^{T_{\text{w}}} \mathbf{A} \text{Diag}(\mathbf{p}_i) \mathbf{A}^{\mathsf{T}} = \mathbf{B}\mathbf{B}^{\mathsf{T}}, \tag{9}$$

where $\mathbf{B} \in \mathbb{R}^{N \times R}$ which can be found by computing the eigenvalue decomposition of $\overline{\mathbf{R}}$; more specifically, when $\overline{\mathbf{R}} = \mathbf{U}\mathbf{\Sigma}^2\mathbf{U}^{\mathsf{T}}$ we can set $\mathbf{B} = \mathbf{U}\mathbf{\Sigma}$. We can then (approximately) whiten the observations by a subsequent decorrelation and rescaling of the data: $\tilde{\mathbf{R}}_i = \mathbf{B}^{\dagger}\mathbf{R}_i(\mathbf{B}^{\dagger})^{\mathsf{T}} \in \mathbb{R}^{R \times R}$, with $\mathbf{B}^{\dagger} = (\mathbf{B}^{\mathsf{T}}\mathbf{B})^{-1}\mathbf{B}^{\mathsf{T}} = \mathbf{\Sigma}^{-1}\mathbf{U}^{\mathsf{T}}$. Note that we reduce the dimensionality of the problem by only using the first R eigenvectors, which is beneficial for efficiently storing all covariance matrices. After pre-whitening, we are left with estimating an unknown orthogonal rotation $\tilde{\mathbf{A}} \in \mathbb{R}^{R \times R}$ and the power time courses in $\mathbf{P}$. Stacking the vectorized pre-whitened covariances matrix, we obtain

$$\tilde{\mathbf{Y}} = (\tilde{\mathbf{A}} \odot \tilde{\mathbf{A}})\mathbf{P}. \tag{10}$$

After estimating $\tilde{\mathbf{A}}$, we can restore $\mathbf{A}$ easily by $\mathbf{A} = \mathbf{B}\tilde{\mathbf{A}}$.

B. Identification of communities

NN-TALS The first algorithm we consider is TALS, the classical workhorse for solving the CPD. Typically, the three factor matrices in (4), $\mathbf{A}$, $\mathbf{A}$, $\mathbf{P}$ are considered as separate entities, which we denote by $\mathbf{A}_1$, $\mathbf{A}_2$, $\mathbf{P}$, respectively. We implement TALS by solving two least-squares problems for $\mathbf{A}_1$ and $\mathbf{P}$ in an alternating fashion, where we set the duplicate factor matrix $\mathbf{A}_2$ equal to its updated version $\mathbf{A}_1$. Additionally, we note that $\mathbf{P}$ is non-negative by construction and thus include this constraint in the update.

P-NN-TALS To establish a second algorithm, note that (10) has a form similar to (4), which in turn can be rewritten to a similar form as the CPD in (7). Thus, it is also possible to perform NN-TALS on a tensor $\tilde{\mathcal{R}}$ of stacked pre-whitened covariance matrices, where the first two factor matrices should be constrained to be orthogonal. Writing $\tilde{\mathcal{R}}$ in its first-mode unfolding $\tilde{\mathcal{R}}_{(1)}$, and solving for $\tilde{\mathbf{A}}_1$ while fixing the other matrices, results in the following optimization problem:

$$\begin{aligned} \min_{\tilde{\mathbf{A}}_1} \quad & ||\tilde{\mathcal{R}}_{(1)} - \tilde{\mathbf{A}}_1(\mathbf{P}^{\mathsf{T}} \odot \tilde{\mathbf{A}}_2)^{\mathsf{T}}||_F \\ \text{s.t.} \quad & \tilde{\mathbf{A}}_1^{\mathsf{T}}\tilde{\mathbf{A}}_1 = \mathbf{I}, \end{aligned} \tag{11}$$

which is the orthogonal Procrustes problem. The solution of this problem can be found by computing the compact SVD of $\tilde{\mathcal{R}}_1(\mathbf{P}^{\mathsf{T}} \odot \tilde{\mathbf{A}}_2) = \mathbf{D}\mathbf{E}\mathbf{F}^{\mathsf{T}}$, and setting $\tilde{\mathbf{A}}_1 = \mathbf{D}\mathbf{F}^{\mathsf{T}}$, which is the closest orthogonal approximation. Then, we update the second duplicate factor matrix by $\tilde{\mathbf{A}}_2 = \tilde{\mathbf{A}}_1$ to ensure symmetry. Finally, we update $\mathbf{P}$ by solving the following non-negative least squares problem:

$$\begin{aligned} \min_{\mathbf{P}} \quad & ||\tilde{\mathcal{R}}_{(3)} - \mathbf{P}^{\mathsf{T}}(\tilde{\mathbf{A}}_1 \odot \tilde{\mathbf{A}}_2)^{\mathsf{T}}||_F \\ \text{s.t.} \quad & \mathbf{P} \geq 0, \end{aligned} \tag{12}$$

where $\geq$ denotes element-wise non-negativity, $\tilde{\mathcal{R}}_{(3)}$ is the third-mode unfolding of $\tilde{\mathcal{R}}$, and both $\tilde{\mathbf{A}}_1$ and $\tilde{\mathbf{A}}_2$ are fixed.

P-AP For the third algorithm, we use the algorithm proposed in [10]. Let us consider matrix $\tilde{\mathbf{Y}}$ in (10), which allows a compact singular value decomposition (SVD) $\tilde{\mathbf{Y}} = \mathbf{U}_s\mathbf{\Sigma}_s\mathbf{V}_s^{\mathsf{T}}$, where $\mathbf{U}_s \in \mathbb{R}^{N^2 \times R}$, $\mathbf{\Sigma}_s \in \mathbb{R}^{R \times R}$, and $\mathbf{V}_s^{\mathsf{T}} \in \mathbb{R}^{R \times T_{\text{w}}}$. Then, the method uses the fact that $\text{range}(\tilde{\mathbf{A}} \odot \tilde{\mathbf{A}}) = \text{range}(\mathbf{U}_s)$, and uses an alternating projection algorithm to find suitable orthonormal vectors $\tilde{\mathbf{a}}_r \; \forall \; r \in [1, ..., R]$ that lie in $\text{range}(\mathbf{U}_s)$ by deflation.

P-FFDiag The final algorithm we consider is a pre-whitened version of the FFDiag algorithm [12]. This algorithm is a well-known JD algorithm, which we now apply to the pre-whitened data. It basically minimizes the following cost function:

$$\begin{aligned} \min_{\tilde{\mathbf{A}}} \quad & \sum_{i=1}^{T_{\text{w}}} \sum_{j \neq k} ((\tilde{\mathbf{A}}^{\mathsf{T}}\tilde{\mathbf{R}}_i\tilde{\mathbf{A}})_{jk})^2 \\ \text{s.t.} \quad & \tilde{\mathbf{A}}^{\mathsf{T}}\tilde{\mathbf{A}} = \mathbf{I}. \end{aligned} \tag{13}$$

C. Pearson correlation coefficients versus covariance

Assuming the data $\mathbf{X}_i$ within window i is demeaned, the Pearson correlation coefficient matrix $\mathbf{R}_i^{\text{p}}$ is computed as follows:

$$\mathbf{R}_i^{\mathrm{p}} = \frac{1}{L}\mathbf{W}_i\mathbf{X}_i\mathbf{X}_i^{\mathsf{T}}\mathbf{W}_i^{\mathsf{T}} = \mathbf{W}_i\mathbf{R}_i\mathbf{W}_i^{\mathsf{T}}, \tag{14}$$

where matrix $\mathbf{W}_i = \mathrm{Diag}(\mathrm{diag}(\frac{1}{L}\mathbf{X}_i\mathbf{X}_i^{\mathsf{T}})^{-1/2})$ contains the inverse standard deviations on the diagonal and varies per window. Then, by substituting the bilinear model, we have

$$\mathbf{R}_i^{\mathrm{p}} = \mathbf{W}_i\mathbf{A}\mathrm{Diag}(\mathbf{p}_i)\mathbf{A}^{\mathsf{T}}\mathbf{W}_i^{\mathsf{T}}. \tag{15}$$

Crucially, the inferred networks in the columns of $\mathbf{A}$ will change by scaling their individual node memberships in each window. Consequently, the power time courses become dependent on the standard deviations of the data, as (the now distorted) node-membership matrix $\mathbf{A}$ is fixed. Furthermore, transforming covariances to Pearson correlation coefficients appears as a rather unconventional pre-whitening method, where one attempts to rescale the source covariances from the observations without first decorrelating the observations, as done in Sec. III-A. We believe that using Pearson correlation coefficients in single-subject data is suitable only when a single network is active in a window i, resembling state-based modeling. For the above reasons, we opt to utilize covariances when performing a CPD on single-subject data. Note that for the studies decomposing multi-subject static connectivity matrices [5], [7], using Pearson correlation coefficients may be inevitable due to the necessity for proper data standardization. Nevertheless, the fundamental assumptions outlined in Sec. II remain unchanged.

IV. Simulations

We perform simulations to analyze algorithm robustness to several noise intensities and different amounts of available temporal samples T. We include all algorithms described previously. Furthermore, we apply tICA using the FastICA algorithm described in [13].

Concerning synthetic data, we generate white noise sequences $\mathbf{S}_n \in \mathbb{R}^{R\times T}$ by sampling a normal distribution. Furthermore, we generate slowly time-varying power profiles $\mathbf{P}' = [\mathbf{p}'_1, ..., \mathbf{p}'_T]$ with $\mathbf{p}'_t \in \mathbb{R}^{R\times 1}$ following an AR-1 process, where each entry $\mathbf{p}'_t = \alpha_1\mathbf{p}'_{t-1} + \boldsymbol{\epsilon}_t$ with $\alpha_1 = 0.999$ and $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \sigma_\epsilon^2\mathbf{I})$ with $\sigma_\epsilon = 0.05$. Then, we create sources by point-wise multiplication: $\mathbf{S} = \mathbf{S}_n \circledast |\mathbf{P}'|$. This way, the sources are uncorrelated and have a quasi-stationary power profile. Finally, we generate noisy observations by $\mathbf{X} = \mathbf{AS} + \mathbf{N}$, where $\mathbf{n}_t \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$. Then, the signal-to-noise ratio is defined by $\mathrm{SNR} = \frac{1}{T}||\mathbf{AS}||_F^2/N\sigma^2$.

We use a window size $L = 20$ with 50% ($o = 0.5$) overlap to create the covariance matrices. All algorithms are initialized randomly. We allow the TALS-based methods, P-FFDiag, and P-AP to perform 300 iterations, which is sufficient for convergence in practice. For tICA, we allow for 1000 iterations.

To measure algorithm performance, we first undo the permutation ambiguity by matching the time courses in $\mathbf{P}$ and its estimate. In case a method does not explicitly solve for $\mathbf{P}$, we estimate $\mathbf{P}$ using (12). Then, we use the average absolute cosine similarity (AACS) as a performance metric:

$$\mathrm{AACS}(\mathbf{A}, \hat{\mathbf{A}}) = \frac{1}{R}\sum_{r=1}^{R}\left|\frac{\mathbf{a}_r^{\mathsf{T}}\hat{\mathbf{a}}_r}{||\mathbf{a}_r||_2||\hat{\mathbf{a}}_r||_2}\right|. \tag{16}$$

We start with a simple small-scale simulation, setting $R = 2$ and $N = 4$. We construct binary valued communities in $\mathbf{A}$ with 67% overlap. Figs. 1a-b show the performance of various algorithms under varying noise intensities for $T = 500$ and varying sample sizes for an SNR of 3 dB.

As a second simulation example, we slightly increase the complexity by setting $R = 5$ and $N = 10$. Now, we construct binary valued communities in $\mathbf{A}$ randomly, having an expected pairwise overlap of 25%. In Figs. 1c-d, the results for varying SNRs with $T = 500$ and sample sizes with an SNR of 3 dB are illustrated.

V. Discussion

From Fig. 1a, we observe that NN-TALS performs well under noisy conditions. Also, all the pre-whitening-based methods appear more or less on par, while tICA performs the worst in the range of the more challenging SNRs. It is clear that pre-whitening-based methods have limited performance due to their two-stage nature [14], [15]. In Fig. 1b, it is visible that tICA also suffers under small sample sizes. Several ICA implementations, such as FastICA [13], also incorporate a pre-whitening step applied to the matrix $\mathbf{X}$. Given that the slowly changing power time courses in $|\mathbf{P}'|$ do not significantly alter the cross-correlations among the rows of $\mathbf{S}_n$, both the pre-whitening procedures in the BSS-QSS methods and tICA are expected to unveil similar left eigenvectors $\mathbf{U}$, a fact we have confirmed. However, in the final stage, where the rotation matrix is determined, tICA encounters challenges due to the limited sample size and non-Gaussianity assumptions.

In Fig. 1c, we observe that when increasing N and R, surprisingly, NN-TALS suffers for higher SNRs (this result is also found in [9]). Note that very accurate solutions can still be obtained; however, the average performance is lower. We can see that increasing N makes the NN-TALS algorithm more prone to obtaining ill-conditioned factor matrices during the alternating updates. At this point, the necessity of pre-whitening becomes apparent, as pre-whitening-based methods outperform NN-TALS. For high-dimensional columns of $\mathbf{A}$, pre-whitening thus adds a layer of robustness by sacrificing performance. Moreover, our proposed method P-NN-TALS marginally outperforms P-FFDiag, as anticipated, due to incorporating a non-negativity constraint on $\mathbf{P}$. Also, it can be seen that P-AP has lower performance over a wide range of SNRs, which likely originates from its deflation-style nature, where accurate identification of the first $\tilde{\mathbf{a}}_r$ is key for the identification of others. Finally, we note that in our second simulation example, P-AP requires more samples to match with the other pre-whitening-based methods, as visible in Fig. 1d.

We leave further analysis on experimental neuroimaging data sets for future work, in which the potential improvement in reproducibility of overlapping communities between subjects has to be demonstrated.

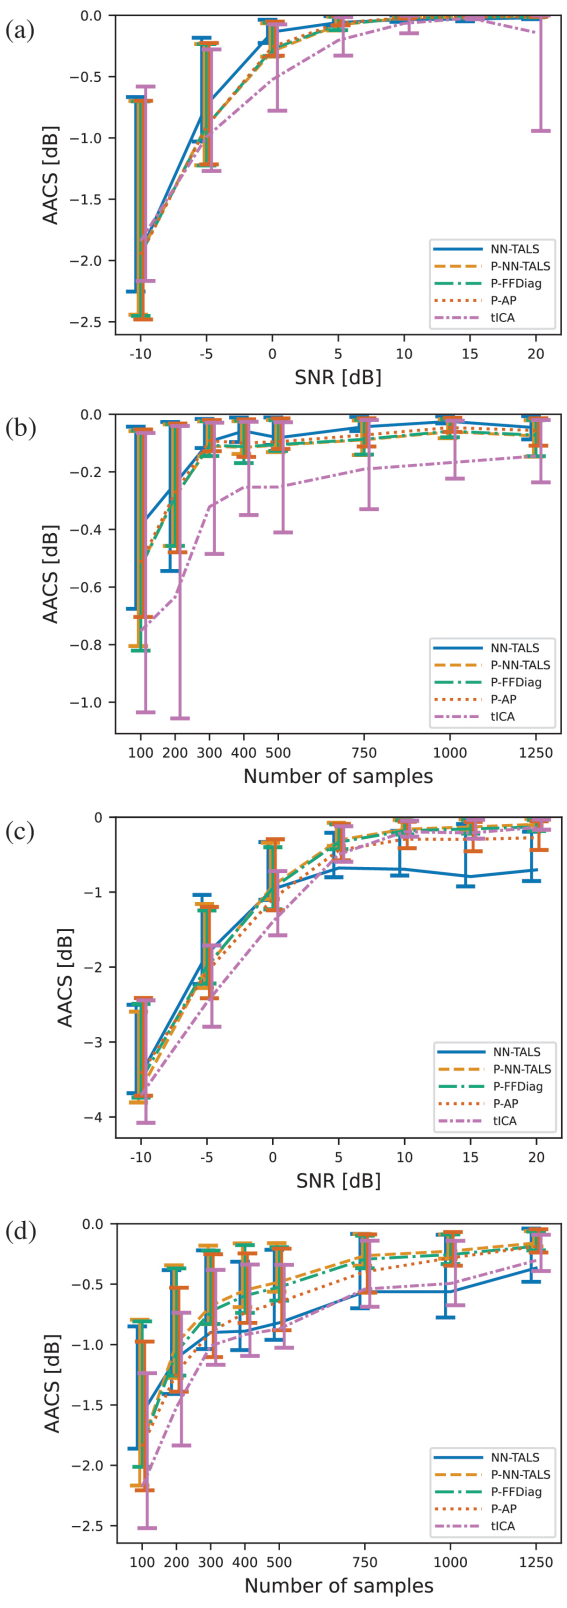

Fig. 1. Algorithm performance for **(a)** varying SNRs with fixed $T = 500$ and **(b)** varying number of samples T with fixed SNR $= 3$. Idem dito for **(c)** and **(d)**, respectively. For **(a,b)**, $R = 2$ and $N = 4$, while for **(c,d)**, $R = 5$ and $N = 10$. Error bars display the asymmetric 75% percentiles over 100 Monte Carlo runs, which are slightly offset for better visualization.

VI. Conclusion

We conclude that assumptions such as quasi-stationarity and uncorrelatedness of the hidden source signals form the basis of state-of-the-art methods that perform overlapping community detection. Furthermore, we briefly analyzed the robustness and performance of several methods that solve the BSS-QSS problem using simulations. Our simulations confirm that tICA suffers under a low number of samples and SNR. Moreover, the proposed algorithm marginally outperforms state-of-the-art methods, and is a viable and robust alternative for accurately identifying overlapping communities in scenarios with challenging SNRs and a low number of samples. Given that the assumptions mentioned earlier hold, the evaluated BSS-QSS-based methodology might improve the reproducibility of overlapping communities between experimental neuroimaging data sets through improved identification accuracy.

References

[1] S.M. Smith, K.L. Miller, S. Moeller, J. Xu, E.J. Auerbach, et al., "Temporally-independent functional modes of spontaneous brain activity," Proceedings of the National Academy of Sciences, vol. 109, pp. 3131–3136, 2012.

[2] E.A. Allen, E. Damaraju, S.M. Plis, E.B. Erhardt, T. Eichele, et al., "Tracking whole-brain connectivity dynamics in the resting state," Cerebral Cortex, vol. 24, pp. 663–676, 2014.

[3] X. Liu and J.H. Duyn, "Time-varying functional network information extracted from brief instances of spontaneous brain activity," PNAS, vol. 110, pp. 4392-4397, 2013.

[4] K. Glomb, A. Ponce-Alvarez, M. Gilson, P. Ritter, G. Deco, et al., "Resting state networks in empirical and simulated dynamic functional connectivity," NeuroImage, vol. 159, pp. 388–402, 2017.

[5] H. Eavani, T.D. Satterthwaite, R. Filipovych, R.E. Gur, R.C. Gur, et al., "Identifying sparse connectivity patterns in the brain using resting-state fMRI," NeuroImage, vol. 105, pp. 286–299, 2015.

[6] Y. Zhu, J. Liu, and F. Cong, "Dynamic community detection for brain functional networks during music listening with block component analysis," IEEE Transactions on Neural Systems and Rehabilitation Engineering, vol. 31, pp. 2438–2447, 2023.

[7] Y. Wang and Y. Guo, "LOCUS: A regularized blind source separation method with low-rank structure for investigating brain connectivity," The Annals of Applied Statistics, vol. 17, pp. 1307–1332, 2023.

[8] V.D. Calhoun, T. Adali, G.D. Pearlson, and J.J. Pekar, "Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms," Human Brain Mapping, vol. 13, pp. 43–53, 2001.

[9] K.-K. Lee, W.-K. Ma, X. Fu, T.-H. Chan, and C.-Y. Chi, "Blind identification of mixtures of quasi-stationary sources using a khatri-rao subspace approach," 2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR), pp. 2169-2173, 2011.

[10] K.-K. Lee, W.-K. Ma, X. Fu, T.-H. Chan, and C.-Y. Chi, "A Khatri–Rao subspace approach to blind identification of mixtures of quasi-stationary sources," Signal Processing, vol. 93, pp. 3515-3527, 2013.

[11] Y. Rong, S.A. Vorobyov, A.B. Gershman, and N.D. Sidiropoulos, "Blind spatial signature estimation via time-varying user power loading and parallel factor analysis," IEEE Transactions on Signal Processing, vol. 53, pp. 1697–1710, 2005.

[12] A. Ziehe, P. Laskov, and G. Nolte, "A fast algorithm for joint diagonalization with non-orthogonal transformations and its application to blind source separation," JMLR, vol. 5, pp. 777-800, 2004.

[13] A. Hyvärinen and E. Oja, "Independent component analysis: algorithms and applications," Neural Networks, vol. 13, pp. 411–430, 2000.

[14] J.-F. Cardoso, "On the performance of orthogonal source separation algorithms," Proc. EUSIPCO, pp. 776–779, 1994.

[15] L. De Lathauwer, B. De Moor, and J. Vandewalle, "A prewhitening-induced bound on the identification error in independent component analysis," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 52, pp. 546–554, 2005.