Flexible Data Fusion of EEG-fMRI: Revealing Neural-Hemodynamic Coupling Through Structured Matrix-Tensor Factorization

Simon Van Eyndhoven*†, Borbála Hunyadi*†, Lieven De Lathauwer*‡ and Sabine Van Huffel*†
*KU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Kasteelpark Arenberg 10, 3001 Leuven, Belgium
†imec, Leuven, Belgium ‡ KU Leuven Kulak, Group Science, Engineering and Technology, 8500 Kortrijk, Belgium
Email: simon.vaneyndhoven@esat.kuleuven.be

***Abstract*—Simultaneous recording of electroencephalographic (EEG) signals and functional magnetic resonance images (fMRI) has gained wide interest in brain research, thanks to the highly complementary spatiotemporal nature of both modalities. We propose a novel technique to extract sources of neural activity from the multimodal measurements, which relies on a structured form of coupled matrix-tensor factorization (CMTF). In a data-symmetric fashion, we characterize these underlying sources in the spatial, temporal and spectral domain, and estimate how the observations in EEG and fMRI are related through neurovascular coupling. That is, we explicitly account for the intrinsically variable nature of this coupling, allowing more accurate localization of the neural activity in time and space. By means of a simulation study, we illustrate the effectiveness of this approach, which is shown to be robust to noise. Hence, this might provide a conceptually simple, yet effective alternative to other data-driven analysis methods in event-related or resting-state EEG-fMRI studies.**

I. Introduction

In the field of neuroimaging and brain mapping, the combined use of EEG and fMRI has been widely recognized as a promising approach, since it exploits both the high temporal resolution from EEG and the excellent spatial resolution from fMRI. These two key factors enable researchers and clinicians to address a multitude of questions on cognitive or pathological neural activity, which underlies the data from both modalities [1].

The complementarity of the two modalities is at the same time its own pitfall: it is not straightforward to combine two such heterogeneous data streams - the fast electrical recording in EEG which is sparsely sampled over the scalp, and the sluggishly varying Blood Oxygen Level Dependent (BOLD) signal in fMRI which, however, can be localized accurately. Despite the mismatch in spatiotemporal resolution, we can assume that the two involved modalities measure the same neural processes in the brain, although in an indirect way and to a different extent. Thus, it is unknown to which degree the spatial and temporal observations in EEG and fMRI are related. This has led to the development of a wide variety of analysis tools, all approaching this data fusion problem from different angles.

One way to characterize these tools is by the way they try to fit a model to the data and pose restrictions to it. On the one hand, model-driven techniques exist, ranging from very complex [2] to simpler models such as the widely applied General Linear Model (GLM), which has been a standard workhorse to analyze fMRI datasets [3]. In the common way of applying the GLM, a single 'canonical' hemodynamic response function (HRF) is used, which describes the expected BOLD signal resulting from an impulse event in a reference stimulus time series. Hence, the HRF waveform is considered invariant over the different brain areas, over neural 'events', and even over different subjects in a group study. Although the GLM has its merit as a simple and attractive framework to make sense of (multimodal) brain data, this limitation obfuscates the interpretations that can be drawn through its use. Namely, it is well known that the hemodynamic response displays a significant variability over people, brain areas and over time [4]–[6], which might even be more severe for the diseased brain [7]. Hence, misspecification of the HRF can lead to biased estimates of activity and loss in statistical power [5], [6]. To cope with this variability, several methods have been proposed, such as adding the derivatives of the canonical HRF during GLM regression, or non-linear fitting methods [3], [8]. However, most of these methods still make an unfavorable trade-off between complexity versus flexibility: either they deal with only small variations, or they require the estimation of HRF parameters on a voxel-wise basis. On the other end of the spectrum, data-driven blind source separation (BSS) techniques such as independent component analysis can be used to decompose both EEG and fMRI data into a set of constituting sources, in which case neural sources are found by imposing sometimes arbitrary restrictions [9], [10]. Traditional BSS methods such as independent component analysis [9] or canonical correlation analysis [11] rely mostly

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Advanced Grant: BIOTENSORS (n° 339804). This paper reflects only the authors' views and the Union is not liable for any use that may be made of the contained information. The authors acknowledge funding from Research Council KU Leuven: C1 project c16/15/059-nD and CoE PFV/10/002 (OPTEC), and from F.W.O.: project G.0830.14N and G.0881.14N.

on matrix factorization techniques, which operate on data in a two-way format. Over the past years however, tensor-based techniques which factorize third or higher order arrays have been increasingly used [12], [13], also in EEG-fMRI data analysis [14]–[16]. They offer the advantage that the inherent multidimensional nature of datasets (such as those in brain analytics) is respected, and most importantly, that the factorizations of tensors can be unique under mild conditions, as opposed to matrix factorization [12]. However, when applied to recordings of brain activity, these data-driven techniques exploit (by definition) no model of the neurovascular coupling, and might thus neglect prior (approximate) knowledge on the joint 'behavior' of the datasets.

We aim to tackle this issue by taking a hybrid, semi-informed estimation approach, borrowing ideas from the GLM as well as from BSS methods. That is, we try to find sources of neural activity by fusing EEG and fMRI recordings in a data-driven fashion, while still accounting for a characteristic temporal relationship between the electrophysiological activity (as measured with EEG) and the BOLD signal (measured with fMRI), in the form of an a priori unknown HRF. Hence, we insert a prior on the neural-hemodynamic coupling into the problem formulation, which admits to co-estimate the true underlying HRF from the data itself. We achieve this combined goal of extracting neural-hemodynamic sources and their temporal coupling by expressing the problem as a coupled matrix-tensor factorization (CMTF) [16], [17], where the fMRI data is a matrix of time samples $\times$ voxels and the EEG spectrogram is represented as a 3rd order tensor of time samples $\times$ electrodes $\times$ frequencies as in [15]. The factorizations of both datasets are then constrained by a shared factor matrix in the temporal mode, which is transformed by a convolution with an unknown HRF in the fMRI factorization. This is a generic framework, in which different kinds of prior knowledge about the HRF can be inserted to reduce the search space when solving the problem.

In the following sections, we first explain the tensor-based framework that we use to tackle the problem. Then, we analyze the feasibility of the proposed approach by means of a simulation study. Finally, we present extensions to our approach that allow to cope with additional variability in the data.

II. Methods

A. Generative model

We assume that the recorded EEG and fMRI result from the superposition of several underlying neural sources, each with a distinct spatial origin in the brain, accompanied by a time course of their activity. The EEG recording $\mathbf{X}$ over channels (electrodes) and time is assumed to be transformed to a spectrogram representation, by means of e.g. the short-time Fourier transform, followed by squaring the absolute values of the complex coefficients. The resulting 3rd order tensor $\boldsymbol{\mathcal{X}}$ describes the variation over the spatial, temporal, and spectral mode. The fMRI data is a matrix $\mathbf{Y}$, describing the variation over the temporal and spatial mode.

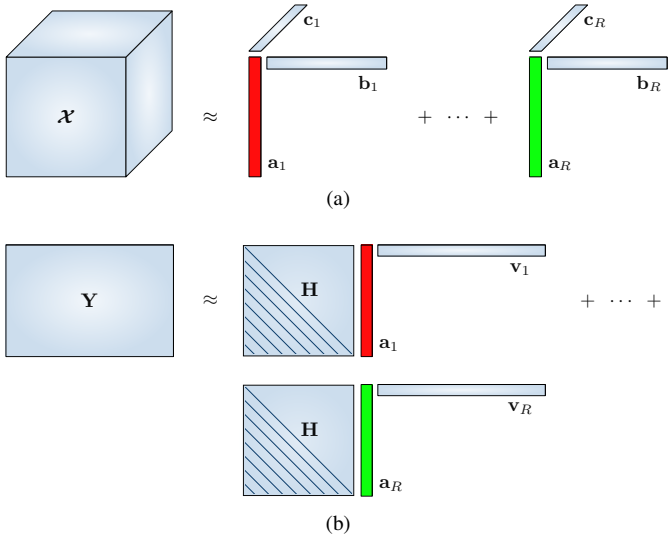

Fig. 1. Structured coupled matrix-tensor factorization of (a) EEG data and (b) fMRI data. The signatures $\mathbf{a}_r$ along the temporal mode are shared between the EEG factorization and the fMRI factorization, providing the coupling constraint.

Loosely following the notation in [15], the EEG and fMRI data can be written as the sum of R sources, which are reflected in both datasets as follows:

$$\begin{aligned} \boldsymbol{\mathcal{X}} &= \sum_{r=1}^{R} \mathbf{a}_r \otimes \mathbf{b}_r \otimes \mathbf{c}_r + \boldsymbol{\mathcal{E}}_x \\ \mathbf{Y} &= \sum_{r=1}^{R} \mathbf{u}_r \otimes \mathbf{v}_r + \mathbf{E}_y = \sum_{r=1}^{R} \mathbf{u}_r \mathbf{v}_r^T + \mathbf{E}_y \ , \end{aligned} \tag{1}$$

where $\otimes$ indicates the outer product. The column vectors $\mathbf{a}_r$, $\mathbf{b}_r$, and $\mathbf{c}_r$ describe the activity of source r over time, and its distribution over EEG channels and frequencies, respectively. Analogously, the vectors $\mathbf{u}_r$ and $\mathbf{v}_r$ describe the time course of source r and its distribution over voxels, respectively. The residuals are contained in $\boldsymbol{\mathcal{E}}_x$ and $\mathbf{E}_y$.

B. Structured matrix-tensor factorization

The electrophysiological time course $a_r(t)$ of the source in the EEG is assumed to predict the hemodynamic time course $u_r(t)$ in fMRI by means of a convolution with an unknown HRF $h(t)$, which can be rewritten as the matrix product of the temporal signature $\mathbf{a}_r$ with a Toeplitz matrix $\mathbf{H}$, holding the samples of $h(t)$ on its diagonals:

$$\mathbf{u}_r = \mathbf{H}\,\mathbf{a}_r = \text{Toeplitz}(h)\mathbf{a}_r \ . \tag{2}$$

To incorporate prior physiological knowledge, we constrain the HRF to a 'plausible' waveform. For this, we use a family of parametrized basis functions, i.e., we can model the HRF as the difference of two gamma functions (like is done for the often-used canonical HRF in the GLM). The HRF thus depends on a small number of unknown parameters as $h(t) = f(t,\theta) = \Gamma(\theta_1)^{-1} \cdot \theta_2^{\theta_1} t^{\theta_1 - 1} e^{-\theta_2 t} - \theta_3 \Gamma(\theta_4)^{-1} \cdot \theta_5^{\theta_4} t^{\theta_4 - 1} e^{-\theta_5 t}$.

Furthermore, note that the EEG tensor is non-negative, because it describes the spectral power on the different channels on different time instants. To obtain an interpretable

result, we impose that the signatures $\mathbf{a}_r$ and $\mathbf{c}_r$ are also non-negative, since they describe the time-varying power and spectra of the sources, respectively. Hence, we reformulate the data fusion problem as a structured matrix-tensor factorization [18], in which the factor matrix $\mathbf{A} = [\,\mathbf{a}_1\,\mathbf{a}_2 \ldots \mathbf{a}_R\,]$ is shared between the two datasets, while the factor matrices $\mathbf{B} = [\,\mathbf{b}_1\,\mathbf{b}_2 \ldots \mathbf{b}_R\,]$, $\mathbf{C} = [\,\mathbf{c}_1\,\mathbf{c}_2 \ldots \mathbf{c}_R\,]$, $\mathbf{V} = [\,\mathbf{v}_1\,\mathbf{v}_2 \ldots \mathbf{v}_R\,]$ are modality-specific and the parameter set θ describes the HRF factor matrix $\mathbf{H}$. The coupled factorization can then be solved by minimizing the following cost function:

$$J(\mathbf{A},\mathbf{B},\mathbf{C},\mathbf{V},\theta) = \|\boldsymbol{\mathcal{X}} - \sum_{r=1}^{R} \mathbf{a}_r \otimes \mathbf{b}_r \otimes \mathbf{c}_r\|^2 + \|\mathbf{Y} - \sum_{r=1}^{R} (\mathbf{H}(\theta)\mathbf{a}_r) \otimes \mathbf{v}_r\|^2 \;,\; \text{s.t. } \mathbf{A},\mathbf{C} \geq 0 \tag{3}$$

in which $\|\cdot\|$ indicates the Frobenius norm. Fig.1 shows the data model, including the dependencies (i.e. coupling) between EEG and fMRI, which is used to find a structured coupled matrix-tensor factorization. Under this model, the fMRI data can be written as $\mathbf{Y} = \mathbf{H}\mathbf{A}\mathbf{V}^T$. Contrary to other (standalone) matrix factorizations, rotational ambiguity is ruled out between any pair of the matrices $\mathbf{H}$, $\mathbf{A}$ and $\mathbf{V}^T$, since the middle factor $\mathbf{A}$ is constrained by the EEG factorization.

C. Simulation study

We conducted simulations in which the data was generated according to the model in (1) and (2). That is, we placed $R = 3$ neural sources at distinct positions in a template brain of $91 \times 109 \times 91$ voxels of $2 \times 2 \times 2\text{mm}^3$, and picked a HRF which differed from the canonical HRF, as shown in Fig.2a, 2b, 2c and Fig.4 respectively. We assigned (overlapping) temporal and spectral signatures to the sources (see Fig.3), and calculated the spatial distributions over 19 electrodes by means of a forward volume conduction model. The resulting signatures of the generated sources are depicted in Fig.2d, 2e, 2f. The temporal signatures span 60 seconds and are sampled at 4 Hz. After convolution with the chosen HRF, the obtained fMRI time courses are downsampled to 1 Hz, mimicking the discrepancy in temporal resolution between the two modalities. To deal with this mismatch, the downsampling operation is implicitly assumed in (3) and taken into account during optimization. The spectral signatures consist of 20 frequency bins.

We added white Gaussian noise of varying signal-to-noise ratio (SNR) to both EEG and fMRI data. In the remainder of the text, SNR is measured as the ratio of the variance of the noise-free time course, either in a voxel or in a electrode-frequency pair with the highest variance, and the variance of the noise. Before factorizing, we applied dimensionality reduction on the fMRI data, by computing the variance in every voxel, computing z-scores of all variances and rejecting voxels with a z-score below 1. To summarize, we jointly factorize a noisy EEG dataset $\boldsymbol{\mathcal{X}}$ of dimensions $240 \times 19 \times 20$ and a noisy fMRI dataset $\mathbf{Y}$ of dimensions $60 \times N$, where N indicates the remaining number of voxels (typically $\sim 10^5$).

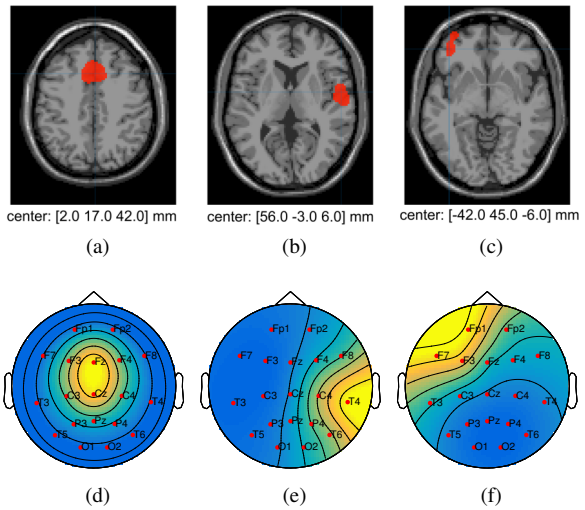

Fig. 2. Location and extent of the three neural sources in the fMRI domain (a, b, c) and topographic plots of the corresponding EEG spatial signatures, obtained through a forward volume conduction model, based on the anatomical MRI image and the sources' positions (d, e, f).

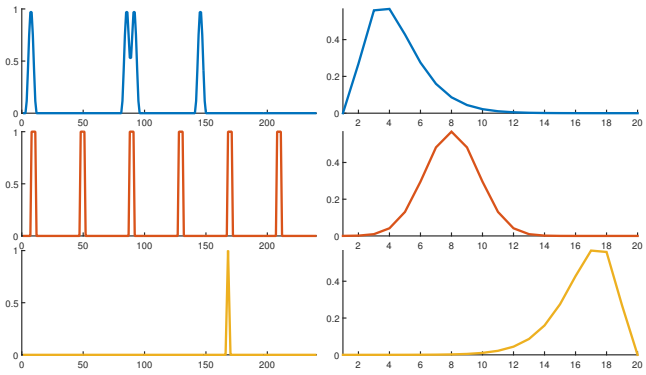

Fig. 3. Temporal (left column) and spectral (right column) EEG signatures, which belong to the neural sources in Fig.2.

Optimization of the cost function in (3) is done through the Structured Data Fusion framework in Tensorlab [18], [19]. In a realistic setting, the true number sources is unknown a priori. We therefore computed the factorization for R varying from 1 to 4, and picked the most likely number of sources according to the core consistency diagnostic [12], [16]. In all but a few cases, this heuristic correctly estimated the number of sources, which is why we only report results for the case where the number of extracted sources equals R. We aimed to address several subquestions, further described below.

How reliably can the true signatures be extracted?

Under SNR levels ranging from -8 dB to 12 dB, in steps of 4 dB, we performed the coupled EEG-fMRI factorization and checked the relative estimation errors. We exhaustively repeated this for data which was simulated using every possible combination of 2 sources, or using all 3 sources. For all sources, we computed relative estimation errors in the EEG temporal mode as $\delta_a = \frac{1}{R}\sum_{r=1}^{R} \frac{\|\hat{\mathbf{a}}_r - \mathbf{a}_r\|}{\|\mathbf{a}_r\|}$ after an optimal permutation and scaling of the estimated signatures $\hat{\mathbf{a}}_r$ (and analogously for the spectral and spatial modes). In the fMRI

domain, we quantified the spatial overlap of true and estimated sources. For a true underlying source s which is active in a set P consisting of N_s voxels, we took the set Q of N_s voxels with the highest value from the estimated signature $\hat{\mathbf{v}}_r$ (after applying the same permutation as for the EEG sources) and computed the Dice coefficient $d = \frac{2|P \cap Q|}{|P|+|Q|} = \frac{|P \cap Q|}{N_s}$. The EEG signatures were initialized by performing a factorization on the EEG dataset separately using parallel factor analysis (PARAFAC) [12], [16]. The HRF was initialized as the canonical HRF, and the initial voxel signatures were obtained from $\mathbf{Y}$ using the pseudo-inverse of the initialized factors $\mathbf{H}(\theta)$ and $\mathbf{A}$. We compared against reference results, which were obtained by fixing the parameter set θ in (3) to the parameters of the canonical HRF, similarly to the method in [15].

To which extent can the model cope with HRF variability?

We created different HRFs by varying the timing, width and relative amplitudes of the positive and negative lobes. The time of the positive peak varied from 2 s to 10 s, in steps of 2 s, and the negative peak came 5 s, 10 s, or 15 s later. The width (measured as the difference in time between the 2 points where the 2nd derivative is zero) of the first lobe was either 2, 5, or 8 s, and the width of the negative lobe was always 1.5 times larger. The 2nd peak had a relative amplitude of 0% (no peak), 20% or 60% w.r.t. the positive peak. We aimed to extract the signatures from the datasets generated with 3 sources and these different parameters, for SNR = 4 dB. We used the Pearson correlation coefficient to assess the similarity between the waveforms of the estimated and the true HRF.

What is the robustness versus non-white noise?

In the previous experiments, we used additive white Gaussian noise (WGN) for the simulations. To investigate the effect of noise which is colored (i.e. non-white) in the spatial or temporal domain, we conducted two additional tests. In the first test, referred to as 'SpCorr', we introduced spatial correlation in the fMRI noise by filtering it with a 3D Gaussian smoothing kernel with a standard deviation of 2 voxels, before adding it to the generated data. In the second test, referred to as 'TempCorr', we created the fMRI noise by means of an autoregressive model of order 1 (AR(1)), as $e_t = 0.7e_{t-1} + \xi_t$, where ξ_t is drawn from a white Gaussian noise distribution. The SpCorr experiment takes into account that (background) neural processes exhibit spatial covariance over neighbouring voxels, while the TempCorr experiment mimics AR(1)-structured residuals, which is an often-used noise model in fMRI [3]. We still added WGN to the EEG dataset in both cases. We evaluated this scenario for the combination with 2 active sources. Since only the noise for the fMRI data set was altered, we track performance in terms of the Dice coefficient.

III. Results

Reliability of extracted signatures

Fig.5a and Fig.5b show the mean relative errors δ for the EEG factorization in the case of 2 and 3 sources, respectively. It is clear that, on average, the signatures in the temporal,

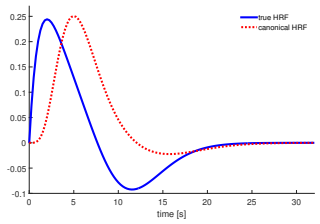

Fig. 4. True underlying hemodynamic response function (full blue line) in this study, and the canonical hemodynamic response function (dotted red line) as used in the SPM toolbox. The positive (resp. negative) peaks of both HRFs are 3 seconds (resp. 4 seconds) apart.

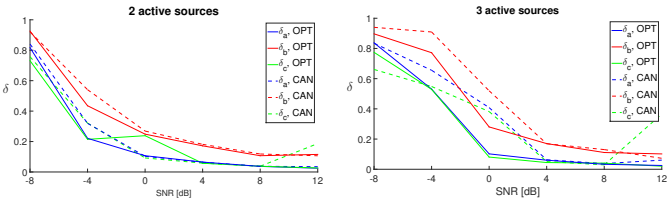

(a) (b)

Fig. 5. Mean relative errors δ_a, δ_b, and δ_c between true and estimated sources in the EEG's temporal, spatial, and spectral mode, respectively, in the presence of 2 sources (a) or 3 sources (b). Results are shown for the case where we find the HRF in the optimization procedure ('OPT') and for the case where we fix the HRF to the canonical HRF ('CAN').

spatial and spectral mode are more accurately estimated when we account for HRF variability ('OPT' condition). This is expected, since the EEG factorization is influenced by the fMRI factorization through coupling of the temporal signatures. A similar effect is observed when looking at the Dice coefficients: from Fig.6a we conclude that modeling HRF variability has a beneficial effect on the localization of neural sources in the brain, i.e. a higher overlap with the true sources is achieved.

Robustness versus HRF variability

In table I we show the correlation between estimated and true underlying HRF, for HRFs that vary according to 4 parameters: time to positive peak, separation between positive and negative peak, dispersion (i.e. the width of the positive lobe) and the relative amplitude of the negative peak compared to the positive peak. As a reference: the parameters of the canonical HRF are 5 s, 10.7 s, 4.5 s and 0.09, respectively. Values are

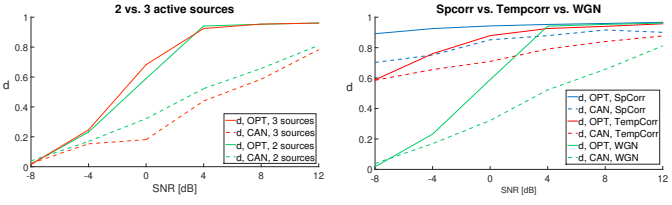

(a) (b)

Fig. 6. Mean dice coefficients d between true and estimated fMRI spatial signatures, in the presence of 2 or 3 active sources and WGN (a) or 2 active sources and spatially ('SpCorr') or temporally ('TempCorr') correlated noise (b). Results are shown for the case where we find the HRF in the optimization procedure ('OPT') and for the case where we fix the HRF to the canonical HRF ('CAN').

TABLE I
PEARSON CORRELATION BETWEEN ESTIMATED AND TRUE HRF.

time of 1st peak	peak separation	dispersion	relative amplitude
2 s 0.71±0.31	5 s 0.81±0.28	2 s 0.70±0.35	0 0.88±0.23
4 s 0.94±0.03	10 s 0.90±0.20	5 s 0.90±0.19	0.2 0.81±0.31
6 s 0.95±0.05	15 s 0.81±0.29	8 s 0.91±0.17	0.6 0.83±0.25
8 s 0.93±0.09			
10 s 0.66±0.41			

given as the mean ± one standard deviation of the pooled results, when only the indicated parameter is fixed. In general, the algorithm succeeds reasonably well in capturing the HRF waveform, admitting a decrease in performance when the true HRF differs substantially from the initial guess (which was the canonical HRF, in our case). We note that there was a significant dependence (correlation 0.84) between the average Dice coefficient and the correlation between true and estimated HRF. This confirms findings of previous research, which pointed out the potential bias resulting from HRF mismodeling [5], [6], and elicits the importance of HRF estimation.

Robustness versus spatially/temporally non-white noise

Fig.6b shows the mean Dice coefficients under the effect of spatially (SpCorr) or temporally (TempCorr) correlated noise. Interestingly, the spatial extent of the neural sources is estimated more accurately under AR(1) noise conditions, or under spatially correlated noise, compared to WGN conditions, especially at lower SNR. No significant difference in quality of the extracted EEG signatures was observed between conditions when the HRF was co-estimated, but using the canonical HRF resulted in an average increase in relative error of ±10-15% in all modes.

IV. DISCUSSION

We have demonstrated that accounting for HRF variability can have a beneficial effect on the estimation of spatial, temporal, and spectral aspects of neural sources from combined EEG-fMRI measurements. For this purpose, we have introduced a data-driven framework based on structured matrix-tensor factorization, which can be used to identify patterns of neural activity. It is generic in the sense that it describes neural-hemodynamic coupling solely by means of a convolution operation with an initially unknown HRF, and thus minimally makes assumptions on the neural sources of interest, adhering to the idea of blind source separation. Any prior model for the nature of the HRF can be plugged in, as long as it can be expressed as an explicit and differentiable function of some parameters [18]. Complete flexibility can be attained by not imposing any model for $h(t)$, which is conceptually the same as the 'FIR basis set' method [8], while the GLM can be mimicked by fixing $h(t)$ to the canonical HRF. In this paper, we have chosen a model consisting of two gamma functions, because of its simplicity and similarity to the often-used canonical HRF. While several methods have been proposed to find the HRF from the experimental data, they often require to re-estimate the HRF voxel by voxel [6]–[8], which is tedious and precludes easy interpretation. Instead, our proposed method seeks to find one HRF, which leads to the best fit for the whole dataset. Nonetheless, the framework allows to deal with variation of neurovascular coupling over brain regions, in a way that is analogous to well-known approaches which use a basis set of HRF functions [6], [8]. This is achieved by extending the fMRI factorization as $\mathbf{Y} = \mathbf{H}_1\mathbf{A}\mathbf{V}_1^T + \mathbf{H}_2\mathbf{A}\mathbf{V}_2^T + \ldots + \mathbf{H}_B\mathbf{A}\mathbf{V}_B^T$, in which B basis functions $h_1(t)$ to $h_B(t)$ are used. Hence, the BOLD response is described as a voxel-specific linear combination over basis functions - as well as over different neural sources, like in the simple case. Many other adaptations of this framework can be conceived, depending on the context of the data (e.g. event-related or resting state EEG-fMRI experiments), the desired flexibility, and the type of information which one tries to extract. This is the subject of future work.

REFERENCES

[1] P. Ritter and A. Villringer, "Simultaneous EEG–fMRI," *Neuroscience & Biobehavioral Reviews*, vol. 30, no. 6, pp. 823–838, 2006.

[2] P. A. Valdes-Sosa *et al.*, "Model driven EEG/fMRI fusion of brain oscillations," *Human brain mapping*, vol. 30, no. 9, pp. 2701–2721, 2009.

[3] M. M. Monti, "Statistical analysis of fMRI time-series: a critical review of the GLM approach," *Frontiers in human neuroscience*, vol. 5, no. 28, 2011.

[4] J.-R. Duann *et al.*, "Single-trial variability in event-related BOLD signals," *Neuroimage*, vol. 15, no. 4, pp. 823–835, 2002.

[5] D. A. Handwerker *et al.*, "Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses," *Neuroimage*, vol. 21, no. 4, pp. 1639–1651, 2004.

[6] M. A. Lindquist *et al.*, "Modeling the hemodynamic response function in fMRI: efficiency, bias and mis-modeling," *Neuroimage*, vol. 45, no. 1, pp. S187–S198, 2009.

[7] C.-G. Bénar *et al.*, "The BOLD response to interictal epileptiform discharges," *Neuroimage*, vol. 17, no. 3, pp. 1182–1192, 2002.

[8] F. Pedregosa *et al.*, "Data-driven HRF estimation for encoding and decoding models," *NeuroImage*, vol. 104, pp. 209–220, 2015.

[9] V. D. Calhoun *et al.*, "A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data," *Neuroimage*, vol. 45, no. 1, pp. S163–S172, 2009.

[10] D. Lahat *et al.*, "Multimodal data fusion: an overview of methods, challenges, and prospects," *Proceedings of the IEEE*, vol. 103, no. 9, pp. 1449–1477, 2015.

[11] Y.-O. Li *et al.*, "Joint blind source separation by multiset canonical correlation analysis," *IEEE Transactions on Signal Processing*, vol. 57, no. 10, pp. 3918–3929, 2009.

[12] A. Cichocki *et al.*, "Tensor decompositions for signal processing applications: From two-way to multiway component analysis," *IEEE Signal Processing Magazine*, vol. 32, no. 2, pp. 145–163, 2015.

[13] A. Cichocki, "Tensor decompositions: a new concept in brain data analysis?" *arXiv preprint arXiv:1305.0395*, 2013.

[14] B. Hunyadi *et al.*, "Tensor decompositions and data fusion in epileptic electroencephalography and functional magnetic resonance imaging data," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 7, no. 1, 2017.

[15] E. Martınez-Montes *et al.*, "Concurrent EEG/fMRI analysis by multiway partial least squares," *NeuroImage*, vol. 22, no. 3, pp. 1023–1034, 2004.

[16] E. Karahan *et al.*, "Tensor analysis and fusion of multimodal brain images," *Proceedings of the IEEE*, vol. 103, no. 9, pp. 1531–1559, 2015.

[17] E. Acar *et al.*, "Understanding data fusion within the framework of coupled matrix and tensor factorizations," *Chemometrics and Intelligent Laboratory Systems*, vol. 129, pp. 53–63, 2013.

[18] L. Sorber *et al.*, "Structured data fusion," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 4, pp. 586–600, 2015.

[19] N. Vervliet *et al.*, "Tensorlab 3.0," March 2016, Available online. URL: http://www.tensorlab.net.