

Augmenting interictal mapping with neurovascular coupling biomarkers by structured factorization of epileptic EEG and fMRI data

Van Eyndhoven, Simon ; Dupont, Patrick; Tousseyn, Simon ; Vervliet, Nico; Van Paesschen , Wim; Van Huffel, Sabine; Hunyadi, Borbala

DOI

[10.1016/j.neuroimage.2020.117652](https://doi.org/10.1016/j.neuroimage.2020.117652)

Publication date

2021

Document Version

Final published version

Published in

NeuroImage

Citation (APA)

Van Eyndhoven, S., Dupont, P., Tousseyn, S., Vervliet, N., Van Paesschen , W., Van Huffel, S., & Hunyadi, B. (2021). Augmenting interictal mapping with neurovascular coupling biomarkers by structured factorization of epileptic EEG and fMRI data. *NeuroImage*, 228, 1-20. Article 117652. <https://doi.org/10.1016/j.neuroimage.2020.117652>

Important note

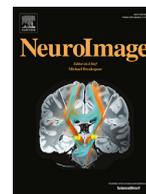
To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Augmenting interictal mapping with neurovascular coupling biomarkers by structured factorization of epileptic EEG and fMRI data

Simon Van Eynhoven^{a,*}, Patrick Dupont^{b,c}, Simon Tousseyn^d, Nico Vervliet^a, Wim Van Paesschen^{e,f}, Sabine Van Huffel^a, Borbála Hunyadi^g

^a Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, KU Leuven, Belgium

^b Laboratory for Cognitive Neurology, Department of Neurosciences, KU Leuven, Leuven, Belgium

^c Leuven Brain Institute, Leuven, Belgium

^d Academic Center for Epileptology, Kempenhaeghe and Maastricht UMC+, Heeze, the Netherlands

^e Laboratory for Epilepsy Research, KU Leuven, Leuven, Belgium

^f Department of Neurology, University Hospitals Leuven, Leuven, Belgium

^g Circuits and Systems Group (CAS), Department of Microelectronics, Delft University of Technology, Delft, the Netherlands

ARTICLE INFO

Keywords:

EEG-fMRI
Blind source separation
Tensor factorization
Interictal epileptic discharge
Neurovascular coupling
Hemodynamic response function

ABSTRACT

EEG-correlated fMRI analysis is widely used to detect regional BOLD fluctuations that are synchronized to interictal epileptic discharges, which can provide evidence for localizing the ictal onset zone. However, the typical, asymmetrical and mass-univariate approach cannot capture the inherent, higher order structure in the EEG data, nor multivariate relations in the fMRI data, and it is nontrivial to accurately handle varying neurovascular coupling over patients and brain regions. We aim to overcome these drawbacks in a data-driven manner by means of a novel structured matrix-tensor factorization: the single-subject EEG data (represented as a third-order spectrogram tensor) and fMRI data (represented as a spatiotemporal BOLD signal matrix) are jointly decomposed into a superposition of several sources, characterized by space-time-frequency profiles. In the shared temporal mode, Toeplitz-structured factors account for a spatially specific, neurovascular ‘bridge’ between the EEG and fMRI temporal fluctuations, capturing the hemodynamic response’s variability over brain regions. By analyzing interictal data from twelve patients, we show that the extracted source signatures provide a sensitive localization of the ictal onset zone (10/12). Moreover, complementary parts of the IOZ can be uncovered by inspecting those regions with the most deviant neurovascular coupling, as quantified by two entropy-like metrics of the hemodynamic response function waveforms (9/12). Hence, this multivariate, multimodal factorization provides two useful sets of EEG-fMRI biomarkers, which can assist the presurgical evaluation of epilepsy. We make all code required to perform the computations available at <https://github.com/svaneyn/structured-cmf>.

1. Introduction

Refractory epilepsy is a neurological disorder suffered by 30% of approximately 50 million epilepsy patients worldwide (World Health Organization, 2019), in which seizures cannot adequately be controlled by anti-epileptic medication. In the preparation of treatment via resective surgery, interictal epileptic discharges (IEDs) can be localized in the brain with simultaneous EEG-fMRI, which provides a good surrogate for mapping the seizure onset zone (An et al., 2013; Grouiller et al., 2011; van Houdt et al., 2013; Khoo et al., 2017; Lemieux et al., 2001; Thornton et al., 2010; Vulliemoz et al., 2009; Zijlmans et al., 2007). This is often conducted via EEG-correlated fMRI analysis, wherein a reference temporal representation of the IEDs is used to interrogate all brain regions’ blood oxygen level dependent (BOLD) signals for significant cor-

relations; voxels for which a statistical threshold is exceeded can then be considered part of the epileptic brain network, along which epileptic seizures are generated and propagated (Gotman, 2008; Lemieux et al., 2001; Salek-Haddadi et al., 2003; Thornton et al., 2010; Zijlmans et al., 2007).

Since its inception, the workhorse for conducting EEG-correlated fMRI analysis has been the general linear model (GLM) framework (Friston et al., 1994; Poline and Brett, 2012; Salek-Haddadi et al., 2006). Over the past years, it has become clear that using the GLM comes with several hurdles, related to the many modeling assumptions, that may reduce its sensitivity or specificity (increasing Type I errors) when violated (Lindquist et al., 2009; Monti, 2011; Poline and Brett, 2012). Remedies for several of these issues are not yet widely applied, or are not yet available.

* Corresponding author.

E-mail addresses: simon.vaneynhoven@gmail.com, simon.vaneynhoven@gmail.com (S. Van Eynhoven).

<https://doi.org/10.1016/j.neuroimage.2020.117652>

Received 15 April 2020; Received in revised form 28 November 2020; Accepted 4 December 2020

Available online 24 December 2020

1053-8119/© 2020 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

First of all, the adoption of a relevant representation of IED occurrences to construct a regressor for the design matrix has proven vital to the sensitivity. This aspect has been investigated in Rosa et al. (2010), Murta et al. (2015), Abreu et al. (2018), Van Eyndhoven et al. (2019a). In previous work (Van Eyndhoven et al., 2019a), we addressed this issue by pre-enhancing the EEG signals using a spatiotemporal filter that is tuned to maximize the signal-to-noise ratio (SNR) of IEDs with respect to the background EEG. We have shown that taking the time-varying power of the filtered EEG leads to a robust regressor, which is more performant than many other types of regressors, including those based on stick functions (Lemieux et al., 2001; Salek-Haddadi et al., 2006), ICA (Abreu et al., 2016; Formaggio et al., 2011) or EEG synchronization (Abreu et al., 2018).

Model mismatch may occur due to the unknown neurovascular coupling from electrophysiological phenomena measured on the EEG to hemodynamic variations captured by the BOLD signals. In many papers on EEG-correlated fMRI, a canonical hemodynamic response function (HRF) based on two gamma density functions is used to translate IED-related temporal dynamics to BOLD fluctuations (Friston et al., 1998). However, there is insurmountable evidence that the HRF is not fixed, but varies substantially over subjects (Aguirre et al., 1998), over brain regions (Handwerker et al., 2004), with age (Jacobs et al., 2008), or even with stress level (Elbau et al., 2018). For the diseased brain, this issue may be even greater: i.e., additional variation, e.g. in brain areas involved in the epileptic network, has been observed compared to healthy controls (Bénar et al., 2002; Grouiller et al., 2010; Van Houdt et al., 2013; Jacobs et al., 2009; Lemieux et al., 2008). Plenty of previous research has shown that failing to account for this variability may lead to substantial bias and increased variance of the estimated activation, which in turn inflates Type I and/or Type II error rates (Calhoun et al., 2004; Lindquist et al., 2009; Lindquist and Wager, 2007; Monti, 2011).

Several methods have been devised to deal with this variability. A widely used approach is to model the HRF as a linear combination of several basis functions. Some popular choices for these bases, which are also supported by open source toolboxes like SPM are the ‘informed basis set’ (Friston et al., 1998), consisting of the HRF plus its derivative w.r.t. time and its derivative w.r.t. the dispersion parameter (leading to a Taylor-like extension which can capture slight changes in peak onset and width), and the finite impulse response (FIR) basis set, in which every basis function fits exactly one sample of the HRF in every voxel (Aguirre et al., 1998; Glover, 1999). Other researchers have aimed to find a basis set by computing a low-dimensional subspace of a large set of ‘reasonable’ HRFs (Woolrich et al., 2004) or by fitting nonlinear functions to given fMRI data (Lindquist and Wager, 2007; Van Eyndhoven et al., 2017). Alternatively, multiple copies of a standard HRF, which differ only in their peak latencies, can be used (Bagshaw et al., 2004). Finally, approaches exist that aim to be immune to differences in neurovascular coupling, such as those based on mutual information (MI), which does not rely on any predefined model or even linearity of the HRF (Caballero-Gaudes et al., 2013; Ostwald and Bagshaw, 2011). Perhaps surprisingly, the authors of Caballero-Gaudes et al. (2013) found that the results based on MI were often very similar to those based on the informed basis set, leading to the conclusion that the assumption of a linear time-invariant system, as described by the convolution with an appropriate HRF, is sufficiently accurate. Instead, it may be useful to not make abstraction of the variable neurovascular coupling, but rather consider it as an additional biomarker to localize epileptogenic zones (van Houdt et al., 2013). Indeed, in several studies, HRFs that deviate from the canonical model were found in regions of the epileptic network (Bénar et al., 2002; Hawco et al., 2007; van Houdt et al., 2013; Jacobs et al., 2009; Lemieux et al., 2008; Moeller et al., 2008; Pittau et al., 2011). Several hypotheses have been postulated to explain this variability, including altered autoregulation due to higher metabolic demand following (inter)ictal events (Schwartz, 2007), vascular reorganization near the epileptogenic region (Rigau et al., 2007), or the existence of pre-spike changes in neuro-electrical activity which are not visible on

EEG and which culminate in the IED (Jacobs et al., 2009). It is thus an opportunity to map not only regions with statistically significant BOLD changes in response to IEDs, but also the spatial modulation of the HRF waveform itself, in order to discover regions where an affected HRF shape may provide additional evidence towards the epileptic onset.

The previous considerations indicate that it is difficult to meet all assumptions in the general linear model, which may compromise inference power (Handwerker et al., 2004; Lindquist et al., 2009; Monti, 2011). Data-driven alternatives may relieve this burden, since they adapt to the complexity of the data more easily compared to model-based approaches, and are especially suited for exploratory analyses (Mantini et al., 2007; Mareček et al., 2016). Blind Source Separation (BSS) techniques consider EEG and/or fMRI data to be a superposition of several ‘sources’ of physiological activity and nonphysiological influences. Based on the observed data alone, BSS techniques are used to estimate both the sources and the mixing system, by means of a factorization of the data into two (or more) factor matrices, holding sources or mixing profiles along the columns. They naturally allow a symmetrical treatment of EEG and fMRI data, enabling true fusion of both modalities (Calhoun et al., 2009; Lahat et al., 2015; Valdes-Sosa et al., 2009), which is in contrast to EEG-correlated fMRI, where EEG-derived IEDs inform the fMRI analysis. While the information-theoretic approach in Caballero-Gaudes et al. (2013) also shares this symmetry feature, it purposely avoids the estimation of HRFs, which is our goal here. Furthermore, BSS techniques naturally accommodate higher-order representations of the data in the form of tensors or multiway arrays, which can capture the rich structure in the data. Indeed, measurements of brain activity inherently vary along several modes (subjects, EEG channels, frequency, time,...), which cannot be represented using matrix-based techniques like ICA without loss of structure or information (Acar et al., 2007; Lahat et al., 2015; Sidiropoulos et al., 2017). Tensor-based BSS techniques have been used to mine unimodal EEG data by decomposing third-order spectrograms (channels \times time points \times wavelet scales) into several ‘atoms’ (also coined ‘components’ or ‘sources’), each with a distinct spatial, temporal and spectral profile/signature (Mareček et al., 2016; Miwakeichi et al., 2004; Mørup et al., 2006), with successful application in seizure EEG analysis (Acar et al., 2007; De Vos et al., 2007). While a tensor extension of ICA for group fMRI data (in the form of subjects \times time points \times voxels) exists (Beckmann and Smith, 2005), matrix representations of fMRI remain dominant for single-subject analyses. Moreover, such a tensor-based extension implicitly assumes that sources have the same time course over subjects, which is not an adequate model for IED occurrences, nor resting-state fluctuations. Coupled BSS techniques can estimate components which are shared between both modalities, providing a characterization in both domains (Hunyadi et al., 2017). For example, in Acar et al. (2017), Acar et al. (2019), Hunyadi et al. (2016), Chatzichristos et al. (2018), multi-subject EEG and fMRI data have been analyzed using coupled matrix-tensor factorization (CMTF), wherein the ‘subjects’ factor is shared between the EEG trilinear tensor decomposition and the fMRI matrix decomposition. In Hunyadi et al. (2016), the resulting factor signatures revealed onset and propagation zones of an interictal epileptic network that was common over patients, as well as the modulation of the default-mode network (DMN) activity. Also single-subject data can be decomposed into distinct components, using a shared temporal factor for EEG and fMRI. This requires the use of a model of the neurovascular coupling, to ensure temporal alignment of EEG and BOLD dynamics. In Martínez-Montes et al. (2004), a fixed canonical HRF was used, followed by multiway partial least squares to extract components with spatial, temporal, and spectral signatures. In previous work, we proposed an extension to this technique, where a subject-specific HRF is co-estimated from the available data, along with the components (Van Eyndhoven et al., 2017).

In this paper, we extend this latter technique in order to account not only for subject-wise variation of the HRF, but also capture variations over brain regions. This results in a highly structured CMTF (sCMTF) of

Table 1
Clinical patient data.

patient	gender	ictal onset zone	etiology	surgery	ILAE outcome	follow-up time (y)	# IEDs	# TRs (# sessions)	# EEG channels	IED loc.
p01	F	L temporal	HS	temporal lobe resection	3	5	15	540 (1)	29	F7-T1
p02	F	L parietal	FCD	partial lesionectomy	4	5	663	1620 (3)	29	Pz
p03	F	R parieto-occipito-temporal	Sturge-Weber				105	1080 (4)	21	F8
p04	M	R temporal	unknown				825	1620 (3)	21	F8-T4
p05	F	L anterior temporal	HS	temporal lobe resection	1	8	117	1080 (3)	29	F7-T1
p06	F	R frontal	FCD	partial lesionectomy	5	2	640	1080 (3)	29	Cz-C4
p07	F	L anterior temporal	DNET	temporal lobe resection	1	4	126	1080 (4)	29	F7-T1
p08	M	L temporo-parietal	unknown	overlap eloquent cx			11	1080 (4)	21	T5
p09	F	L temporo-occipital	FCD	overlap eloquent cx			1815	1620 (3)	29	T3-T5
p10	F	R temporal	HS	refused			226	540 (1)	29	F8-T2
p11	M	L anterior temporal	HS	temporal lobe resection	1	6	6	1080 (2)	29	F7-T1
p12	F	R temporal	CNS infection	refused			966	1350 (5)	27	T4

Abbreviations: F = female, M = male, L = left, R = right, CNS = central nervous system, DNET = dysembryoplastic neuroepithelial tumor, FCD = focal cortical dysplasia, HS = hippocampal sclerosis, cx = cortex, IED = interictal epileptic discharge, TR = repetition time, IED loc. = localization of the IED on EEG.

the interictal multimodal data, in which HRF basis functions and spatial weighting coefficients are estimated along with spatial, spectral and temporal signatures of components. By preprocessing the EEG using the data-driven filters from Van Eyndhoven et al. (2019a), we aim to maximize the sensitivity in mapping the interictal discharges. We analyze whether the estimated spatial modulation of the HRF waveform is a viable biomarker when localizing the ictal onset zone, besides the BOLD spatial signatures themselves.

2. Methods and materials

2.1. Patient group

We use data of twelve patients, whom we previously studied in Tousseyn et al. (2014a), Tousseyn et al. (2014b), Tousseyn et al. (2015), Hunyadi et al. (2015), Van Eyndhoven et al. (2019a). These patients were selected based on the following criteria: (1) they were adults which underwent presurgical evaluation for refractory focal epilepsy using EEG-fMRI, and for which there was concordance of all the available clinical evidence regarding the epileptic focus; (2) subtraction ictal single-photon emission tomography (SPECT) coregistered to MRI (SISCOM) images were available for all patients, as well as post-surgery MRI scans when patients were seizure-free (international league against epilepsy (ILAE) outcome classification 1-3 (1, completely seizure-free; 2, only auras; 3, one to three seizure days per year \pm auras; 4, four seizure days per year to 50% reduction of baseline seizure days \pm auras; 5, <50% reduction of baseline seizure days to 100% increase of baseline seizure days \pm auras; 6, more than 100% increase of baseline seizure days \pm auras)); (3) IEDs were recorded during the EEG-fMRI recording session.

This study was carried out in accordance with the recommendations of the International Conference on Harmonization guidelines on Good Clinical Practice with written informed consent from all subjects. All subjects gave written informed consent in accordance with the Declaration of Helsinki, for their data to be used in this study, but not to be made publicly available. The protocol was approved by the Medical Ethics Committee of the University Hospitals KU Leuven. For the patients' complete clinical data, we refer to Table 1.

2.2. Data acquisition and preprocessing

Functional MRI data were acquired on one of two 3T MR scanners (Achieva TX with a 32-channel head coil and Intera Achieva with an eight-channel head coil, Philips Medical Systems, Best, The Netherlands) with an echo time (TE) of 33 ms, a repetition time (TR) of either 2.2 or 2.5 s, and a voxel size of $2.6 \times 3 \times 2.6 \text{ mm}^3$. EEG data were recorded according to the international 10–20 system using MR-compatible caps, sampled at 5 kHz, with Cz reference. The EEG signals were band-pass filtered offline between 1–50 Hz, gradient artifacts were removed using the Bergen plug-in (Bergen fMRI Group, Bergen, Norway) for EEGLAB (Moosmann et al., 2009) and pulse artifacts were subtracted with the Brain Vision Analyzer software (Brain Products, Munich, Germany) (Allen et al., 1998). The signal of every channel was divided by its standard deviation. Two neurologists subsequently inspected and annotated the EEG signals for IEDs.

The fMRI images were realigned, slice-time corrected and normalized to MNI space, resampled to a voxel size of $2 \times 2 \times 2 \text{ mm}^3$, and smoothed using a Gaussian kernel of 6 mm full width at half maximum (FWHM). These processing steps were carried out using SPM8 (Functional Imaging Laboratory, Wellcome Center for Human Neuroimaging, University College London, UK) (Friston et al., 1994). We refer the reader to Tousseyn et al. (2014a) for a detailed description of these preprocessing steps.

We regress out covariates of no interest from the fMRI data. These include: the six motion-correction parameters (plus their squares and derivatives); boxcar regressors at moments of substantial scan-to-scan head movement (larger than 1 mm based on the translation parameters); the first five principal components of the BOLD time series within the cerebrospinal fluid and white matter regions (Behzadi et al., 2007). Subsequently, the BOLD time series are filtered between 0.008–0.20 Hz using the CONN toolbox (Whitfield-Gabrieli and Nieto-Castanon, 2012). For an analysis of the effect of the ordering of these preprocessing steps, we refer to the supplementary material.

The dimensionality of the fMRI data is reduced by means of an anatomical parcellation of the brain. The initial $79 \times 95 \times 68$ images are segmented into regions-of-interest (ROIs) according to the Brainnetome

atlas, which consists of 246 parcels in the grey matter (Fan et al., 2016). For every ROI, one BOLD time series is constructed as the average of the time series of all voxels within the ROI. If multiple acquisition runs (within the same recording session) had been done, the EEG and fMRI data of the different runs are temporally concatenated. Further customized preprocessing steps are treated in Sections 2.3 and 2.4.

2.3. Multi-channel Wiener filtering for spatio-temporal EEG enhancement

In previous work (Van Eyndhoven et al., 2019a), we have shown that pre-enhancing the EEG signals using a data-driven, spatiotemporal filter that is tuned to maximize the signal-to-noise ratio (SNR) of IEDs with respect to the background EEG and artifacts, leads to a BOLD predictor that is more performant than many other predictors, including those based on simple stick functions (Lemieux et al., 2001; Salek-Haddadi et al., 2006), ICA (Abreu et al., 2016; Formaggio et al., 2011) or EEG synchronization (Abreu et al., 2018). This pre-enhancement strategy based on multi-channel Wiener filters (MWF) has error-correcting capabilities and produces an IED representation that improves the localization sensitivity of EEG-correlated fMRI (Van Eyndhoven et al., 2019a).

In brief, the MWF is estimated by first performing time-delay embedding of the multi-channel EEG signals $\mathbf{x}[t] \in \mathbb{R}^{I_m}$, leading to an extended multi-channel, multi-lag signal $\tilde{\mathbf{x}}[t] \in \mathbb{R}^{2I_m\tau+I_m}$ as

$$\tilde{\mathbf{x}} = \begin{bmatrix} \mathbf{x}[t - \tau] \\ \vdots \\ \mathbf{x}[t] \\ \vdots \\ \mathbf{x}[t + \tau] \end{bmatrix} \quad (1)$$

and subsequently computing the filter coefficients as

$$\hat{\mathbf{W}} = \mathbf{R}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}}^{-1}(\mathbf{R}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}} - \mathbf{R}_{nn}), \quad (2)$$

where $\mathbf{R}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}} = \mathbb{E}\{\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T | H = 1\}$ is the covariance matrix of the EEG observed during annotated IED segments ($H = 1$), and $\mathbf{R}_{nn} = \mathbb{E}\{\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T | H = 0\}$ is the covariance matrix of the EEG outside of IED segments ($H = 0$). For the full derivation, we refer the reader to (Somers et al., 2018; Van Eyndhoven et al., 2019a). The EEG signals are then filtered as $\hat{\mathbf{W}}^T \tilde{\mathbf{x}}$. Due to the extension with lagged copies of the signals, channel-specific finite impulse response filters are found. Hence, $\hat{\mathbf{W}}^T \tilde{\mathbf{x}}$ is a set of spatiotemporally filtered output signals, in which IED-like waveforms are preserved while other waveforms, which are not specific to epilepsy, are suppressed¹.

We train the MWF for each patient individually, after embedding the EEG signals using $\tau = 4$ positive and negative lags², and compute the final filter using the generalized eigenvalue decomposition, which ensures the positive definiteness property of the subtracted covariance matrix in (2) (Somers et al., 2018).

2.4. Higher-order data representation

To preserve the intrinsic multiway nature of the data, we represent the preprocessed EEG and fMRI as a tensor and matrix respectively, which are subsequently factorized jointly. This approach differs from the mass-univariate treatment in the traditional GLM, where each voxel is treated individually, and only ‘flattened’ EEG time courses can be entered as regressors. Since epilepsy is manifested with considerable variability between patients, we handle the multimodal data of each patient separately.

¹ To retrieve filtered versions of the original set of channels only (and not of their time-delay embedded copies, which would be redundant), only the middle I_m columns of $\hat{\mathbf{W}}$ are used (cf. (1)).

² We observed in Van Eyndhoven et al. (2019a) that for most patients, the output SNR saturates around this value, corresponding to an interval of -16 ms to +16 ms.

2.4.1. Spatio-temporal-spectral tensor representation of EEG

We adopt a tensorization strategy based on time-frequency transformation of the EEG data to third-order spectrograms (time points \times frequencies \times channels). After the pre-enhancement step described in Section 2.3, we create a spectrogram using the Thomson multitaper method, applied on nonoverlapping EEG segments with a length equal to one repetition time (TR) of the fMRI acquisition. The squared Fourier magnitudes are averaged into 1 Hz bins, from 1 Hz to 40 Hz. Hence, for every EEG channel, we obtain a spectrogram which is synchronized to the fMRI time series. The time points \times frequencies \times channels spectrogram $\mathcal{X} \in \mathbb{R}^{I_s \times I_f \times I_m}$ is further normalized as described in A.1, to equalize the influence of each channel and each frequency, and to focus on relative signal increases or decreases (Mareček et al., 2017; 2016)

2.4.2. Spatio-temporal matrix representation of fMRI

The average BOLD time series are stacked in a time points \times ROIs matrix $\mathbf{Y} \in \mathbb{R}^{I_s \times I_v}$, where $I_v = 246$ ROIs. We normalize each ROI’s time series by subtracting its mean and dividing by its standard deviation.

2.4.3. Neurovascular coupling in the temporal mode

EEG and fMRI data are acquired simultaneously per subject, and are thus naturally coregistered along the ‘time’ mode. This is captured in a temporal factor matrix that is common between the EEG factorization and the fMRI factorization. However, the electrophysiological changes that are picked up by EEG vary on a much more rapid time scale than the sluggish BOLD fluctuations that (indirectly) correspond to the same neural process. The neurovascular coupling that describes the relation between these two complementary signals can be described by a convolution with an HRF³.

In previous work, we developed a CMTF model in which the HRF itself is parametrically estimated from the data (Van Eyndhoven et al., 2017), and a matrix multiplication with Toeplitz structure implements the HRF convolution, as proposed in Valdes-Sosa et al. (2009). In the same paper, we hinted towards an extension based on multiple basis functions to account for the variability of the HRF over brain regions. In the following, we assume that the time course of each EEG source is convolved with an a priori unknown, ROI-specific HRF, which is a superposition of K parametrized basis functions, which leads to a modelled contribution of this source to the ROI’s BOLD signal. Hence, in every ROI i_v , the modeled (unscaled) BOLD time course $\mathbf{z}_{i_v}^{(r)}$ of the r -th neural source is

$$\mathbf{z}_{i_v}^{(r)} = \mathbf{H}_{i_v} \mathbf{s}_r \quad (r = 1 \dots R) \quad (3)$$

$$= \sum_{k=1}^K b_{k,i_v} \mathbf{H}_k \mathbf{s}_r \quad (4)$$

$$= \sum_{k=1}^K b_{k,i_v} \mathcal{T}(\mathbf{h}_k) \mathbf{s}_r \quad (5)$$

$$= \sum_{k=1}^K b_{k,i_v} \mathcal{T}(\mathcal{H}(\theta^k)) \mathbf{s}_r. \quad (6)$$

Here, \mathbf{s}_r is a factor vector holding the time course of the r -th EEG source; \mathcal{H} is an operator that transforms a set of parameters $\theta^{(k)}$ into a full HRF, represented as a vector \mathbf{h}_k ; \mathcal{T} is an operator that transforms \mathbf{h}_k into a Toeplitz matrix \mathbf{H}_k by populating the main and lower diagonals with the HRF samples (see also A.2); b_{k,i_v} is the weight for the k -th HRF basis function in the i_v -th ROI; \mathbf{H}_{i_v} is the Toeplitz matrix holding the total HRF in the i_v -th ROI. This operation is clarified in Fig. 1b.

³ In this paper, we use the term ‘neurovascular coupling’ to describe the coupling characteristic between EEG and fMRI temporal dynamics, and make the silent assumption that this characteristic is a proxy/surrogate for ‘neurovascular coupling’ as it is understood in neuroscience: the model that describes BOLD changes in response to electrical neural ‘events’, which take the form of local field potentials at the synapses.




Fig. 1. Structured coupled matrix-tensor factorization (sCMTF) of EEG and fMRI data can reveal neural sources that are encoded in both modalities, as well as capture the varying neurovascular coupling between the electrophysiological and BOLD changes. (a) The EEG signals vary over time points \times frequencies \times electrodes. The resulting third-order spectrogram tensor \mathcal{X} is factorized according to (8) into R rank-1 components, which each consist of a temporal signature s_r , a spectral signature g_r and a spatial signature m_r . (b) The fMRI data consist of the average BOLD signal in different brain parcels or regions of interest (ROIs), represented in a time points \times ROI matrix \mathbf{Y} , and are factorized according to (11). Neurovascular coupling is modeled as a convolution of the EEG temporal dynamics with a ROI-specific hemodynamic response function (HRF), as in (11)–(13). In this example, each local HRF is represented as a linear combination (encoded by coefficients b_k) of $K = 3$ optimized basis functions, each populating a Toeplitz matrix \mathbf{H}_k which implements a convolution through matrix multiplication with the temporal signatures s_r . Afterwards, each smoothed component r is spatially weighted by a signature v_r . This is accomplished by the elementwise product $b_k * v_r$ of the HRF basis function-specific coefficients b_k and the component-specific amplitudes v_r .

This time course $z_{i_v}^{(r)}$ is conceptually equivalent to a regressor in the GLM's design matrix. We treat the HRF parameter sets $\theta^{(k)}$, $k = 1 \dots K$ as unknown variables, which need to be fitted to the data at hand (Lindquist and Wager, 2007). By parametrizing each basis function, we embed protection against nonsensical HRF shapes, and against overfitting, since the number of parameters to be estimated is greatly reduced compared to the FIR basis in Glover (1999), Aguirre et al. (1998). We employ a double-gamma HRF, i.e., each HRF basis function k is described by five parameters as $h_k(t) = f(t; \theta) = \Gamma(\theta_1)^{-1} \cdot \theta_2^{\theta_1} t^{\theta_1 - 1} e^{-\theta_2 t} - \theta_5 \Gamma(\theta_3)^{-1} \cdot \theta_4^{\theta_3} t^{\theta_3 - 1} e^{-\theta_4 t}$, where we omit the superscript (k) from the parameters θ to not overload the notation.

2.5. Coupled matrix-tensor factorization of EEG and fMRI

After tensorization, we jointly decompose the EEG tensor \mathcal{X} and the fMRI matrix \mathbf{Y} into a set of distinct sources.

The third-order EEG spectrogram is approximated by a sum of R rank-1 terms according to the trilinear canonical polyadic decomposition (CPD) (also referred to as Parallel Factor Analysis (PARAFAC)) as in Miwakeichi et al. (2004), Mareček et al. (2016), Martínez-Montes et al. (2004), Van Eyndhoven et al. (2017). Each rank-1 term $s_r \circ g_r \circ m_r$ describes a source (also called 'component') in terms of an outer product (\circ) of a temporal, spectral, and spatial signature, respectively. Unlike matrix decompositions, the decomposition of a higher-order tensor into a set of sources is unique, up to scaling and permutation ambiguities, without imposing constraints (under mild conditions).

The fMRI matrix is similarly approximated as a sum of rank-1 terms. Coupling arises from the temporal signatures s_r , which are shared between the EEG and fMRI factorization. After processing through a hemodynamic system (as described in Section 2.4.3), each source's BOLD temporal signature is weighted with a spatial signature v_r .

To accommodate additional structured variation in the fMRI data, that is not related to electrophysiological dynamics, we allow a low-rank term to the fMRI factorization which is not coupled with the EEG factorization. We have empirically found that such a low-rank term can capture structured noise, preventing it from biasing the estimation of the parameters which are coupled with the EEG factorization.

The full sCMTF model is then described as:

$$\mathcal{X} = \hat{\mathcal{X}} + \mathcal{E}_x \quad (7)$$

$$= \sum_{r=1}^R s_r \circ g_r \circ m_r + \mathcal{E}_x \quad (8)$$

$$= [\mathbf{S}, \mathbf{G}, \mathbf{M}] + \mathcal{E}_x \quad (9)$$

$$\mathbf{Y} = \hat{\mathbf{Y}} + \mathbf{E}_y \quad (10)$$

$$= \sum_{r=1}^R \sum_{k=1}^K (\mathbf{H}_k s_r) \circ (b_k * v_r) + \sum_{q=1}^Q n_q \circ p_q + \mathbf{E}_y \quad (11)$$

$$= \sum_{k=1}^K (\mathbf{H}_k \mathbf{S}) (b_k^T \circ \mathbf{V}^T) + \mathbf{N} \mathbf{P}^T + \mathbf{E}_y \quad (12)$$

$$= [\mathbf{H}_1 \mathbf{S} \dots \mathbf{H}_K \mathbf{S}] \cdot [\mathbf{B}^T \odot \mathbf{V}^T] + [\mathbf{N}, \mathbf{P}] + \mathbf{E}_y, \quad (13)$$

where $\hat{\mathcal{X}}$ and $\hat{\mathcal{Y}}$ are the low-rank approximations; \mathcal{E}_x and \mathbf{E}_y hold the residuals of both factorizations; $[\mathbf{S}, \mathbf{G}, \mathbf{M}]$ describes the CPD model composed of factor matrices $\mathbf{S} \in \mathbb{R}^{I_s \times R}$, $\mathbf{G} \in \mathbb{R}^{I_g \times R}$, $\mathbf{M} \in \mathbb{R}^{I_m \times R}$, which hold the temporal, spectral and EEG spatial signatures in the columns; the HRF matrices \mathbf{H}_k are constructed as in (3)–(6); $\mathbf{V} \in \mathbb{R}^{I_v \times R}$ is the fMRI spatial factor matrix; $\mathbf{B} \in \mathbb{R}^{I_v \times K}$ is the HRF basis coefficient matrix; $[\mathbf{N}, \mathbf{P}]$ is a rank- Q term to capture fMRI-only structured nuisance; $*$ denotes the Hadamard or elementwise product; \odot denotes the Khatri–Rao product.

Note that the coupled part of \mathbf{Y} is described by RK nonindependent rank-1 terms, or equivalently, by K rank- R block terms. Namely, each rank-1 term $(\mathbf{H}_k \mathbf{s}_r) \odot (\mathbf{b}_k * \mathbf{v}_r)$ describes the convolution of the r -th source's temporal signature with the k -th basis function, after which a spatial loading with vector $(\mathbf{b}_k * \mathbf{v}_r)$ is performed; in all ROIs, there is one source-nonspecific relative weight for each basis function k (captured in \mathbf{b}_k), and source-specific amplitudes (captured in \mathbf{v}_r) (Calhoun et al., 2004).

It is not our aim to estimate HRF variability over sources, but rather, for the sake of easier interpretation, to estimate only variability over patients and ROIs. Hence, to limit the degrees of freedom, the HRF in every ROI does not depend on r , but is shared between all sources, as in Makni et al. (2008), Vincent et al. (2010), Pedregosa et al. (2015). This is expressed by the the Khatri–Rao product in (12)–(13), which forms a constraint that has earlier been used to robustify GLM parameter estimation (Pedregosa et al., 2015). I.e., there are not $RK I_v$ spatial coefficients, but $(R + K) I_v$, i.e., K basis function weights and R source amplitudes in all I_v ROIs. In this way, the Khatri–Rao structure also breaks the curse of dimensionality in the fMRI decomposition if either the number of sources R or the number of basis functions K is high (or both).

The model is depicted in Fig. 1, omitting $[\mathbf{N}, \mathbf{P}]$ to not overload the diagram.

We estimate all parameters of the model in (8) and (11) by iteratively minimizing the cost function J , composed of two data fitting terms and two regularization terms as in (Acar et al., 2014):

$$J(\mathbf{S}, \mathbf{G}, \mathbf{M}, \mathbf{B}, \mathbf{V}, \theta) = \beta_x \|\mathcal{A} - \hat{\mathcal{X}}\|_F^2 + \beta_y \|\mathbf{Y} - \hat{\mathcal{Y}}\|_F^2 + \gamma_x \|\lambda_x\|_1 + \gamma_y \|\lambda_y\|_1 \quad (14)$$

$$\begin{aligned} \text{s.t. } \mathbf{H}_k &= \mathcal{T}(\mathbf{h}_k) = \mathcal{T}(\mathcal{H}(\theta^{(k)})) \\ \lambda_x &= [\lambda_{x,1} \dots \lambda_{x,R}]^T \\ \lambda_{x,r} &= \|\mathbf{s}_r\|_2 \cdot \|\mathbf{g}_r\|_2 \cdot \|\mathbf{m}_r\|_2 \\ \lambda_y &= [\lambda_{y,1} \dots \lambda_{y,R}]^T \\ \lambda_{y,r} &= \sum_{k=1}^K \|\mathbf{b}_k * \mathbf{v}_r\|_2, \end{aligned} \quad (15)$$

where the squared Frobenius norm $\|\mathcal{A}\|_F^2$ of a tensor \mathcal{A} is the sum of its squared elements; $\|\mathbf{a}\|_2$ and $\|\mathbf{a}\|_1$ denote the Euclidean or ℓ_2 -norm and the ℓ_1 -norm or sum of the elements' absolute values of a vector \mathbf{a} , respectively; β_x , β_y , γ_x and γ_y are positive weights; λ_x and λ_y are vectors which hold the amplitudes with which each source is expressed in the EEG and fMRI data, respectively. The squared Frobenius norms of the residuals promote a good fit of the low-rank approximations to the data, while the ℓ_1 -regularization terms penalize excessive source amplitudes and promote a parsimonious⁴ model, similar to the group-LASSO method (Acar et al., 2014; Yuan and Lin, 2006). At the same time, the latter penalty also tends to prevent the occurrence of degenerate terms (Bro, 1997). We minimize (14) using the Structured Data Fusion

⁴ The sparsity-promoting properties of the LASSO penalty are most useful in the context of an underdetermined system, with more coefficients than observations, e.g. in dictionary learning. Here, the problem is heavily overdetermined, and we do not expect that the amplitudes λ_x and λ_y go exactly to zero. However, the ℓ_1 -penalty is still a useful heuristic to avoid degenerate components in the EEG's CP decomposition.

framework in Tensorlab (Sorber et al., 2015; Vervliet et al., 2016), using a quasi-Newton method based on a limited-memory BFGS algorithm, for 50 independent initializations (see Appendix A for details regarding the optimization procedure and parameters). After convergence, each set of estimated factors needs to be calibrated to remove certain ambiguities, and model selection must be performed to pick the best solution, with an appropriate R (see Appendix B for details).

2.6. Statistical inference

We create statistical nonparametric maps (SnPMs) of the obtained spatial signatures \mathbf{v}_r to determine which ROIs sources are significantly (de)activated in relation to the found sources (Nichols and Holmes, 2002; Waites et al., 2005). Namely, under the null hypothesis of no significant BOLD effect related to the EEG dynamics, the fMRI data may be temporally reshuffled without a significant loss of fit to the EEG dynamics in s_r . To this end, we use permutation-based inference, in which the spatial signatures \mathbf{v}_r are compared against their empirically derived distributions, which are obtained via resampling of the fMRI data while freezing the other estimated sCMTF factors. To account for serial correlations in the fMRI time series, we use the robust wavelet-based resampling approach in Bullmore et al. (2001) to ensure exchangeability and to preserve spatiotemporal correlation structure of the original data in the produced surrogate datasets. For each fMRI dataset and every sCMTF solution, we generate $L = 250$ surrogate fMRI $\hat{\mathbf{Y}}^{(l)}$ datasets using the procedure in Bullmore et al. (2001). We resample only the adjusted data $\mathbf{Y} - \mathbf{N}\mathbf{P}^T$, i.e., after removing the components which model variation specific to the fMRI data. We perform inference on a pseudo t-statistic, which we compute for every ROI and for every source as follows:

1. construct a local 'design matrix' with all estimated temporal signatures as in (3): $\mathbf{D}_{i_v} = [\mathbf{z}_{i_v}^{(1)} \dots \mathbf{z}_{i_v}^{(R)}]$,
2. find the new 'betas' by solving $\beta_{i_v}^{(l)} = \mathbf{D}_{i_v}^* \hat{\mathbf{y}}_{i_v}^{(l)} \forall l$,
3. convert the betas to a t-statistic per source by dividing them by their estimated standard deviation (see Friston et al., 1994; Poline and Brett, 2012).

Through this procedure, we obtain L -point empirical null distributions for every source and every ROI. We set the significance threshold as to control the familywise error (FWE) rate at $\alpha = 0.05$, according to the maximum statistic procedure outlined in Nichols and Hayasaka (2003). That is, for every source r , we form the empirical distribution of the maximal t-statistic over all I_v ROIs, and determine source-specific thresholds $T_{(1-\alpha)}^{(r)}$ as the 95%-percentile (to test for activation) and $T_{(\alpha)}^{(r)}$ as the 5%-percentile (to test for deactivation). Finally, we obtain statistical maps for all sources r by applying these thresholds to the original spatial signatures \mathbf{v}_r , which can be considered as the betas of the unshuffled data.

Furthermore, we create a map of the HRF variability over ROIs. For every ROI, we assess how 'unusual' the local HRF is, by measuring its calibrated distance in HRF space to all other ROIs' HRFs. We use two metrics to quantify this (see Appendix C for details on the computation).

1. *Extremity* is computed as one minus the average of the absolute values of the correlations between a HRF waveform and all other HRFs' waveforms.
2. *Entropy* of the HRF waveform is computed as the negative logarithm of the conditional probability of the HRF.

Both for the pseudo t-maps as for the HRF extremity and entropy maps, we furthermore limit the inspection to the 20 ROIs with the highest values, if applicable.

An end-to-end overview of our pipeline, from data preprocessing up until statistical inference, is depicted in Fig. 2.

2.7. Model performance

We use several metrics to quantify the quality of the obtained sCMTF solutions.




Fig. 2. Interictal EEG and fMRI data can be analyzed via structured coupled matrix-tensor factorizations (sCMTF), which reveals both spatial localization of interictal discharges (spikes), and also localized deviations in neurovascular coupling between electrical and BOLD fluctuations. (a) fMRI and EEG data are first separately preprocessed (yellow block). The fMRI data (top row) are structured as a time points \times regions of interest (ROIs) matrix, after BOLD time courses are averaged within predefined or data-driven parcels. The EEG data (bottom row) are structured as a channels \times time points \times frequencies tensor, after the signals are enhanced via a multi-channel Wiener filter (MWF) which is calibrated based on spike annotations, and subsequently undergo a time-frequency transform. (b) The sCMTF of the EEG and fMRI data (blue block) reveals temporally, spatially and spectrally resolved components, and captures spatially varying hemodynamic response functions (HRFs) (cfr. Fig. 1). We show the EEG temporal, spatial and spectral signatures in Figs. 4a and 6a, and the HRFs in Figs. 4b and 6b, for two selected patients. To initialize the sCMTF factors, first a canonical polyadic decomposition (CPD) of the EEG tensor is computed, from which the remaining fMRI factors are initialized. The full sCMTF model is then computed N times, from these N different initializations, and the stability of the resulting factors over runs is assessed. (c) Statistical images are created for the patient's data and the corresponding sCMTF factors (green block). From the sCMTF factors, the spike-related component is picked as the one with the highest temporal correlation to the filtered EEG signals's broadband power envelope. A statistical nonparametric map (SnPM) of this interictal spike-related component is created, revealing co-activated ROIs in a pseudo-t-map (red). For every ROI, the entropy (and also the extremity) of the HRF is computed by assessing its likelihood under the distribution of all other ROIs' HRFs, and a map of this metric is constructed (blue) to reveal localized HRF abnormalities. Both maps can be used to form a hypothesis on the location of the epileptogenic zone, as we show in in Figs. 5 and 7 for the two selected patients. In this paper, we validate our technique on a set of patients for which the outcome is known.

We compare the statistical maps with a ground truth delineation of the ictal onset zone (IOZ) to assess their concordance. This ground truth is the manually delineated resection zone for patients that had undergone surgical treatment and that were seizure-free afterwards (An et al., 2013; Grouiller et al., 2011; van Houdt et al., 2013; Thornton et al., 2010; Zijlmans et al., 2007), or otherwise the hypothetical resection zone, based on concordant evidence from multiple modalities other than EEG-fMRI (cfr. Section 2.1), for patients that were ineligible for or refused surgery (Tousseyn et al., 2014a). The sensitivity for detecting the IOZ is then computed as the fraction of 'true positive' cases, which are determined by the presence or absence of significant activation clusters which overlap the IOZ in the spatial signatures v_r . Following the reasoning in Tousseyn et al. (2014a), we do not consider significantly active voxels or regions outside of the delineated IOZ as false positives. Acknowledging epilepsy as a network disorder, such active regions might reflect seizure or IED propagation, despite not being involved in their generation.

Furthermore, we hypothesize that the spatial variation of the HRF over the brain might reveal additional localizing information regarding the IOZ, i.e., based on considerations explained in Section 1, we

assume that the HRF in or near the IOZ might be distorted compared to nonepileptic brain regions. We test this hypothesis by assessing whether those regions correspond to high values in the HRF entropy and HRF extremity maps (cfr. Section 2.6).

Additionally, we inspect the spectral, spatial and temporal EEG signatures of the extracted sources, and we measure whether the spatial fMRI signatures bear any similarity to known networks of resting-state human brain activity (Shirer et al., 2012).

3. Experiments

3.1. Patient-specific model selection

Table B.3 compiles the results of the model selection described in Appendix B. For each patient, we select the set of sCMTF factors of rank \hat{R} , which best fulfill the criteria. In all cases, we found at least one such a solution, including an IED-related component within that solution. Note that sometimes models with different R might score well on different (subsets of the) criteria, so the selection of the rank is inevitably ambiguous. In the next section, we analyze the individual set of results for




Fig. 3. Goodness of fit of each patient's EEG tensor \mathcal{X} and fMRI matrix \mathbf{Y} , for varying choices of the rank R in the sCMTF. Naturally, the EEG approximation error decreases monotonically for increasing rank (intra-patient). For the fMRI data, the fit already plateaus for very low R . This is due to the presence of additional, uncoupled components $\mathbf{n}_q \circ \mathbf{p}_q$ in the fMRI factorization, which can absorb some of the variance when the number of coupled components is low, but which lose their relevance at higher ranks.

each patient, based on the selected rank, and we analyze the sensitivity of the results to the choice of R .

We show the goodness of fit of the estimated factors for the EEG tensor \mathcal{X} and the fMRI matrix \mathbf{Y} in Fig. 3. Due to the normalization steps which have been applied to the data (cfr. Section 2.2), the sCMTF operates in a regime of moderately high relative approximation errors.

3.2. Spatio-temporo-spectral profiles of interictal discharges

We analyze for each patient the sources which have been estimated via the sCMTF model. We discuss the results of two patients in detail in the next subsections, and include complete results for all other patients in the supplementary material.

Every time, we show (1) the thresholded pseudo t-maps of the IED-related source in the fMRI domain, both for significant activation as for significant deactivation; (2) maps highlighting the ROIs of high HRF entropy and extremity; (3) the temporal profile (time-varying power) s_r , spatial profile (topography) \mathbf{m}_r , and spectral profile \mathbf{g}_r , of each source in the EEG domain; (4) the HRF waveforms in the different ROIs, and the HRF basis functions at convergence of the algorithm.

We plot maximally 800 s of the temporal signatures, to ensure readability. For ease of comparison, we always overlay the broadband MWF envelope (with an arbitrary vertical offset for visualization only), which is the reference time course s_{ref} for selecting the IED-related component (cfr. B.3). For considerations of space, we generally only show the maps of the fMRI spatial signature \mathbf{v}_r for the IED-related components, but discuss the maps of other components when relevant. We show five axial slices of each map: in each case, we show two slices near the highest and lowest voxels of the IOZ or significant regions of the fMRI spatial signature (whichever lies furthest); if applicable, the middle slice is the cross-section with most overlap between IOZ and spatial signature, and the two remaining slices lie halfway between this slice and the extremal slices; otherwise all three bulk slices are chosen with equal spacing between the extremal slices. We cross-validate the maps against known resting state networks (RSNs) of human brain activity from the Stanford atlas (Shirer et al., 2012).

We stress again at this point that a subset of the results is prone to errors due to imperfect sign normalization (cfr. B.1). While it is relatively straightforward to unambiguously determine the 'right' sign of the EEG signatures, this is more challenging for fMRI. That is, frequently, the polarity of the HRF waveform is ambiguous, and making the 'wrong' choice in a voxel i_v (i.e., the HRF has the opposite effect of the true physical cerebral blood flow change) immediately leads to the wrong sign of the spatial coefficients in \mathbf{v}_r in the respective voxel, and their pseudo t-values, for all sources r . To track the occurrence of this fore-

seen failure mode, we also investigate the significant deactivations of the sources⁵.

Note that we designed the HRF variability metrics so that they are *immune* to the polarity of the HRFs. Hence, any high score of the HRF metrics can be reliably interpreted. For each case, we separate the twenty waveforms with the highest entropy scores, and report how many of those are found in ROIs that overlap with the IOZ, along with the probability (in the form of a p-value from a binomial distribution) that this would occur by randomly sampling as many ROIs (under a given fraction of brain that is covered by the IOZ). Hence, this metric is analogous to one minus the false discovery rate (FDR).


3.2.1. Patient 3

We analyze the solution with $\hat{R} = 2$ sources, and show the results in Figs. 4 and 5. Besides one clear IED-related source, there is one other source that is substantially correlated to the reference time course, but with a homogeneous distribution over the head and an unclear spectrum. This may signify that the IEDs do not follow exactly a rank-1 structure in the spectrogram, and that they may be nonstationary in time or space (cfr. the argument made for nonstationary seizures in Hunyadi et al., 2014). The second source's pseudo t-map had significantly active areas symmetrically in the left and right parietal lobe, much more focalized than the EEG topography. In the EEG time courses, we found indeed IED-like waveforms at the times of the peaks in the temporal signature. Hence, we suspect that both sources may reflect the onset and propagation of the IEDs to other areas, respectively. Five out of the twenty ROIs with high-entropy HRFs overlapped with the IOZ, and a significant finding is that several of them are highly noncausal, i.e., with a positive peak before zero seconds. Fig. 5 confirms this, and also shows that the IED-related source is significantly active in different ROIs of the IOZ.


3.2.2. Patient 10

We analyze the solution with $\hat{R} = 5$ sources, and show the results in Figs. 6 and 7. There is a clear IED-related source, and also an artifactual source at ± 33 Hz, which is also present in other patients. Due to its relatively consistent occurrence, we hypothesize that this artifact is due to the MR acquisition. For example, it may be a remnant of a gradient

⁵ Alternatively, it is possible to use a pseudo F-statistic, e.g. the squared pseudo t-value, to bypass the sign correction altogether. The downside of such an approach is that it is then impossible to distinguish activation and deactivation, which may be meaningful.



(a) Temporal (\mathbf{s}_r , top), spatial (\mathbf{m}_r , middle), and spectral (\mathbf{g}_r , bottom) profiles of the 2 sources in the EEG domain, and reference IED time course (\mathbf{s}_{ref} , in green).



(b) Estimated HRF waveforms in all ROIs, split in HRFs with low (grey) and high (blue) entropy

Fig. 4. (a) In the selected solution for patient 3 ($\hat{R} = 2$), both sources have a temporal signature that correlated strongly to the reference IED time course. The first source modeled the main onset of IEDs and was low-frequency and topographically focal, while the second source was spatially and spectrally diffuse and captured the propagation of IEDs to remote areas (cfr. Fig. 8). (b) Five out of the twenty most deviant HRFs were found inside the ictal onset zone (bold lines, $p < 10^{-4}$). These HRFs had main peaks before 0 s, i.e., they led to BOLD changes before the corresponding EEG correlate of the IED was seen.

artifact which is not adequately removed from the data of some channels, cfr. the observation made in Mareček et al. (2016). Surprisingly, this source is significantly active in an extended area in the occipital lobe, overlapping with the visual network. Both HRF metrics reached extreme values at some (distinct) ROIs within the IOZ. The pseudo-t map of the IED-related source shows significantly active ROIs that are concordant with the IOZ, and deactivation of a large part of the default mode network. Furthermore, the IED-related source's EEG topography is very consistent with the clinical diagnosis. The fourth source is active in the default mode network, predominantly in the α band (cfr. Fig. 9).

The fifth source had an unclear spectrum, but its temporal signature corresponds to the occurrence of high-amplitude IEDs. Its pseudo t-map shows widespread activations over the brain, which did not include the IOZ. We expect that this component captures the propagation of IEDs, after onset near the IOZ, similarly to patient 3.

3.2.3. Summary of all patient's results

We provide an overview of the results w.r.t. IOZ detection in Table 2. All results taken together, the sCMTF results allow a correct detection of the IOZ based on the significant IED activation (10/12 cases) and




Fig. 5. The statistical nonparametric maps of the IED-related component (top two rows) and HRF entropy/extremity maps (bottom two rows) of patient 3 show concordance with the ictal onset zone (IOZ). Especially the regions of significant IED activation were accurate, but also five out of the twenty regions with the most deviant (highest entropy) HRFs were found in the IOZ (cfr. Fig. 4b). The ground truth ictal onset zone is highlighted in dark gray with a white contour. ROIs with high values for both HRF variability metrics are colored in orange.

Table 2

The sCMTF leads to three types of spatial information, which can be cross-validated against the ground truth IOZ, as defined in Section 2.7 and summarized for all patients in Table 1: (1) the EEG topography m_{IED} of the IED-related component; (2) the significantly activated and deactivated ROIs in the fMRI spatial signature v_{IED} ; (3) the ROIs with strongly deviating HRF waveforms, as measured via entropy and extremity. Since the EEG topography has a very low spatial resolution, and depends on the attenuation properties of the tissue as well as the orientation of the neural sources in the cortex, we only expect partial similarity to the IOZ's spatial focus; hence, we use the term 'consistent' rather than 'concordant'. The patients who had a good outcome after surgery (patients 5, 7 and 11) had a higher concordance between the three types of spatial clues than patients with a poorer outcome (patients 1, 2 and 6).


patient ID	selected solution \hat{R}	EEG topography consistent with IOZ?	spatial signature v_{IED} concordant with IOZ?		HRF variability metrics complementary to v_{IED} ?		20 highest-entropy ROIs	
			activation	deactivation	entropy	extremity	# in IOZ	(p-value)
p01	6				✓	✓	1	(0.34)
p02	3		✓	✓	✓		1	(0.59)
p03	2		✓		✓	✓	5	(< 10 ⁻⁴)
p04	4	✓	✓		✓	✓	2	(0.32)
p05	5	✓	✓	✓	✓	✓	6	(< 10 ⁻³)
p06	2		✓				0	/
p07	4	✓	✓				1	(0.57)
p08	2		✓	✓			0	/
p09	2	✓	✓	✓		✓	0	/
p10	5	✓	✓		✓	✓	3	(0.02)
p11	2	✓	✓	✓	✓	✓	4	(0.01)
p12	2			✓	✓	✓	7	(< 10 ⁻³)

significant IED deactivation (6/12 cases). For many patients, some of the ROIs with the highest HRF entropy (9/12 cases, of which 8 were complementary to the SnPM) and highest HRF extremity (8/12 cases) also overlapped with the IOZ, which was shown to be (very) unlikely due to chance. All cases are covered by at least one of the metrics, and all patients besides patient 6 had at least two metrics providing correct and complementary localizing info on the IOZ. For nearly all cases, the IED-related component's time course was highly correlated to a reference IED time course, and its spectrum was plausible. In many, but not all cases, this component's EEG topography was also consistent with the location of the IOZ, though this notion is slightly fuzzy because of the very different spatial domains of EEG and (f)MRI—hence we do not use the term 'concordant'. Analysis of the spatial, spectral, and temporal signatures, in combination with inspection of the filtered EEG signals, reveals the identity of RSN oscillations and/or artifacts in the majority of cases. For several patients, we found sources that are active in a narrow spectral band near 33 Hz. While these likely reflect a technical artifact as the result of the MR acquisition, we found no concomitant changes at this frequency in the EEG. This may be the result of the normalization procedure which we applied prior to the decomposition: since every frequency bin was given equal importance, even unnotice-


able but structured fluctuations at higher frequencies may be captured in a component.

3.2.4. Sensitivity to model selection

For many patients, selecting \hat{R} is ambiguous, since more than one solution (with different R) score well on some of the criteria (cfr. Table B.3). Therefore, we analyze the impact of the choice of R on the sCMTF results. For each patient, we select the solution with the rank which is next in line, i.e., which would be a second best (or equally good) choice, based on the same criteria. This is the solution with $R = 1$ for patient 12, $R = 2$ for patients 1, 2, 5 and 7, $R = 3$ for patients 3, 4, 6, 8, 9 and 10, and $R = 4$ for patient 11. For patients 1, 6 and 8, the results deteriorate drastically, as no metric correctly localizes the IOZ. For patient 11, no ROI within the IOZ is significantly activated due to IEDs anymore, but the HRF metrics are still informative. The results for patients 9 and 12 improve, since all metrics are now sensitive to the IOZ. For the other patients, the situation stays more or less the same, i.e., the same metrics are valuable for IOZ localization. However, the maximum value under the different metrics is generally attained at different ROIs compared to the initially selected model.



(a) Temporal (s_r , top), spatial (m_r , middle), and spectral (g_r , bottom) profiles of the 5 sources in the EEG domain, and reference IED time course (s_{ref} , in green).



(b) Estimated HRF waveforms in all ROIs, split in HRFs with low (grey) and high (blue) entropy

Fig. 6. (a) The sCMTF solution with $\hat{R} = 5$ sources was selected for patient 10. One source's temporal signature is highly correlated with the reference IED time course and is identified as the IED-related source, which has a characteristic low-frequency behaviour and with a frontotemporal topography, consistent with the IOZ location. The second source, which has very narrow-band power around ± 33 Hz, likely captured an artifact of the MR acquisition. The fourth source captured α activity in the default mode network (cfr. also Fig. 9). (b) Three out of the twenty most deviant HRFs were found inside the ictal onset zone (bold lines, $p = 0.02$).

4. Discussion

A novel EEG-fMRI data fusion framework

We have proposed an integrated and structured coupled matrix-tensor factorization (sCMTF) framework, which can be used to make inferences on the localization of the ictal onset zone in refractory focal epilepsy based on simultaneous EEG and fMRI recordings. Our approach aims to perform blind source separation of the neural activity related to interictal epileptic discharges (IEDs), and to characterize it in the spatial, temporal, and spectral domain. To this end, we developed a pipeline consisting of (1) semi-automated EEG enhancement based on annota-

tions of the IEDs; (2) modality-specific preprocessing and tensorization steps, which lead to a third-order EEG spectrogram tensor varying over electrodes, time points, and frequencies, and an fMRI matrix with BOLD time courses for a predefined set of regions of interest or parcels; (3) coupled matrix-tensor factorization of the EEG tensor and fMRI matrix along the shared temporal mode, while accounting for variations in the local neurovascular coupling; (4) automated selection of a robust, and relevant IED-related component, and nonparametric testing to infer its spatial distribution in the brain.

We have stressed the importance of and accounted for the variability of the hemodynamic response function (HRF) over different patients




Fig. 7. The statistical nonparametric maps of the IED-related component and HRF entropy/extremity maps of patient 10 show concordance with the ictal onset zone (IOZ). IED occurrences were associated with significantly active (red) regions in and near the IOZ, and at the same time with a deactivation (blue) in a part of the default mode network. Three out of the twenty regions with the most deviant (highest entropy) HRFs were found in the IOZ (cfr. Fig. 6b). The ground truth ictal onset zone is highlighted in dark gray with a white contour. ROIs with high values for both HRF variability metrics are colored in orange.




Fig. 8. The second component in patient 3 likely captured the propagation of IEDs from the irritative zone, given its relatively large correlation to the MWF envelope (cfr. Fig. 4a). The ground truth ictal onset zone is highlighted in dark gray with a white contour.




Fig. 9. The fourth component in patient 10 seemed to pick up activity in the Default Mode Network (DMN), predominantly in the α band (cfr. Fig. 6a).

and brain regions, by equipping the CMTF with the required expressive power via a set of adaptive basis functions. Moreover, after estimating the EEG and fMRI factor signatures, as well as the HRF parameters, we have computed different summary metrics (entropy and extremity) that measure the local deviance of a ROI's HRF compared to other HRFs in the same brain, and have cross-validated the spatial map of these metrics against the ground truth localization of the ictal onset zone.

The sCMTF pipeline proved to be sensitive in detecting the IOZ in all twelve patients in this study. The statistical nonparametric map (SnPM) of the spatial signature of the IED-related component's activation, obtained with the sCMTF, is the best biomarker, which is in line with the traditional EEG-correlated fMRI approach (Lemieux et al., 2001). In the large majority of the patients, several of these significantly active ROIs overlapped with the IOZ. When used in conjunction with the IED-related SnPM, also the HRF entropy and extremity, as measures of how unlikely an HRF is within a specific set of other HRFs, are promising biomarkers, which identified regions of the IOZ that were complementary to those already found by tracking significant IED activation for

nine out of twelve patients in this study. In roughly half of all cases, we also found regions within the IOZ that significantly deactivated in association to IEDs. The patients who had a good outcome after surgery (patients 5, 7 and 11) had a higher concordance between the three types of spatial clues (EEG topography, fMRI (de)activation, HRF variability) than patients with a poorer outcome (patients 1, 2 and 6). While the number of patients is too low to draw conclusions, this observation supports the hypothesis that the degree of such concordance might help to predict surgical outcome. In its current form, the pipeline still predicted too many ROIs that did not overlap with the IOZ. This might in part be due to IED propagation throughout the epileptic network, as postulated in Tousseyn et al. (2014a), but is likely also a result of inherent limitations of the model. Hence, the output of this analysis is only to be used in conjunction with other modalities (e.g. SISCOM) during presurgical evaluation, in order to assess cross-modality agreement, as is already common for EEG-fMRI.

We previously studied the same patient cohort using classical EEG-correlated fMRI analysis, using different types of EEG-derived regres-

sors, in Van Eyndhoven et al. (2019a). There, we concluded that the time-varying power of MWF-enhanced EEG was preferable over regressors based on ICA, EEG synchronization, or simple stick functions, allowing a sensitivity of the IED activation map for IOZ localization of eleven out of twelve cases. Compared to this prior study, the current pipeline yielded an additional correct detection for patient 11, but failed to correctly predict ROIs that overlap with the IOZ for patients 1 and 12. However, for patient 12, the HRF entropy and extremity metrics were very informative, in that they predicted many IOZ ROIs correctly. Hence, the additional flexibility of this pipeline was probably key to the IOZ detection for patient 11. Yet, the two ‘misses’ for the other two patients indicate that the robustness of the current design still needs improvement. For comparison, we include in the supplementary material the statistical maps which we obtained via the analysis in Van Eyndhoven et al. (2019a). It is worth noting that both patient 1 and patient 11 had only few IEDs (15 and 6, respectively), which makes mining of their EEG–fMRI data for BOLD activation difficult a priori (Salek-Haddadi et al., 2006; 2003; Zijlmans et al., 2007). Hence, also for our more flexible method, the probable yield scales with the number of IEDs during recording. We inspected the 20 HRFs and ROIs with the highest extremity and entropy. Hence, it is inevitable that some or most of these ROIs are not within the IOZ, or the IOZ might even not have a deviant HRF. Standalone HRF metrics would hence have a high false discovery rate, even though for several patients, the high proportion of IOZ-covering ROIs among the 20 selected ROIs was very unlikely due to chance (as measured with p-values). Still, it is not always the case that the HRF inside the IOZ is measurably different than in the rest of the brain, in which case these metrics would not be very sensitive.

However, the ROIs that were highlighted by the HRF metrics were often distinct from the ROIs identified as significantly activated to the IEDs. Hence, the SnPMs of the IEDs and the entropy metrics provide very complementary information, and when analyzed jointly, they may infer the location of the IOZ with more certainty, by selecting brain areas where both IED-related and HRF-related metrics have a high value.

We envision that this approach, with minor modifications, may also be used to analyze resting-state EEG–fMRI activity, since also the estimation of extended networks of correlated spontaneous BOLD fluctuations can be biased by spatially varying HRFs. Since in such a case, no IEDs are present, EEG-enhancement like we have done in this study would no longer take place. However, an MWF may still be used to clean up the EEG, e.g. by annotating artifactual periods, which can be removed from the data by the MWF in its dual form (or another tool) (Somers et al., 2018). Similar to the method in Martínez-Montes et al. (2004), the current method could then allow to extract RSNs which are reflected in both EEG and fMRI data, given that sufficiently many components R are extracted.

HRFs vary strongly over subjects and brain regions

There were substantial differences in (estimated) neurovascular coupling over patients and brain regions, as expected. Since we used ‘regularized’ basis functions, which are parametrized as smooth gamma density functions, the resulting HRFs generally had a plausible shape. However, in some cases we found nonsensical shapes, in which, e.g., the waveform had the same polarity over the whole time course, potentially with a bimodal shape (cfr. patient 4). This serves as a humble reminder to not blindly trust the outputted HRFs (or other factor signatures, for that matter). While we have empirically verified that the optimization algorithm converges properly to the true factor signatures and HRFs for synthetic data under mild conditions, there is no guarantee that this holds true for real-life data, which are orders of magnitude more complex, so that a linear generative model like the sCMTF may not be sufficient to describe the interplay between EEG and fMRI. Moreover, the proper behavior of the sCMTF estimation depends on careful preprocessing, and on a proper selection of hyperparameters (in casu: a good value for the number of sources \hat{R}). Hence, manual inspection of the data quality and the solution is still required. Even if the estimated

HRFs or factor signatures may not fully reflect the ‘correct’ underlying physical phenomena, we have demonstrated that they offer actionable information. Not in the least, via summarizing metrics such as HRF entropy and extremity, our algorithm manages to be reasonably robust to subtle changes in the waveform—which is less of interest here than spatial cues towards the IOZ. We reckon that other parametrizations for HRF than the one we have used, might be better suited for the task. These basis functions could even be picked a priori, e.g., from a set of sensible generating parameters (Woolrich et al., 2004). This would even simplify the optimization problem, since the parameter vectors θ_k no longer need to be estimated, at the expense of being less data-driven.

The algorithm used its modeling freedom to fit ‘noncausal’ HRFs, which are ahead of the EEG by as much as 10 s. Generally, we indeed found that many of the estimated HRFs had significant positive or negative amplitudes already before the neural correlate visible on the EEG. This is in line with recurrent findings on BOLD changes that precede the IEDs which were observed in the EEG (Hawco et al., 2007; Jacobs et al., 2009; Moeller et al., 2008; Pittau et al., 2011). We stress that this noncausality may only be in the observation, and not in the underlying physical chain of events: here, it strictly means that we *observed* BOLD changes in the fMRI data that occur before the corresponding *observed* neural correlate on the EEG. Despite the fact that many of the HRFs differed substantially from the canonical HRF, which is causal and peaks approximately 6 s *after* its neural input, we obtained good results as well with the latter HRF as a nonadaptive model for neurovascular coupling (Van Eyndhoven et al., 2019a). The same conclusion was reached by Caballero-Gaudes et al. (2013) for the comparison of the canonical HRF and an information-theoretic approach for BOLD mapping. The reason for this agreement between these different models—which differ substantially in terms of flexibility—is likely that the canonical HRF is positively correlated to the true HRFs which are found inside the IOZ, and as such the resulting activation maps may still be sufficiently informative. In our data and sCMTF results this is indeed the case for many patients.

Prior EEG signal enhancement aids analysis

Importantly, our pipeline heavily relies on a prior enhancement of the interictal spikes in the EEG data, which would otherwise have a too low SNR for the sCMTF algorithm to pick up IED-related sources. We employ multi-channel Wiener filters, which solely rely on the annotation of a sufficient amount of IEDs in the data itself, or in related data (e.g., data from the same patient, recorded outside the MR scanner). While this task still frequently relies on the skill of human EEG readers and neurologists, advanced automated solutions for interictal spike detection are available (Scheuer et al., 2017; Wilson et al., 1999). Within each solution of a specific rank, we picked the IED component as the one with the highest correlation with a reference time course directly derived from the enhanced EEG. Some of the presented results make clear that this reference time course is not completely free from artifacts, hence caution is warranted when many high-amplitude artifacts are still present in the reference. In this study, however, we have not encountered any issues that seemed to be the direct results of a noisy reference during IED component selection. For the fMRI data, we have carried out a relative strict, but unsupervised ‘enhancement’, by regressing out multiple potential confounds. Hence, it would be worthwhile to perform the fMRI cleanup according to a task-based or supervised criterion as well, e.g., using ICA combined with noise component identification (Salimi-Khorshidi et al., 2014).

The interpretation of components

Overall, the sCMTF pipeline succeeded in extracting meaningful IED-related components, alongside components that modeled resting-state neural fluctuations and physiological and technical artifacts. The fact that the sCMTF can estimate signatures and statistical maps for multiple components is a powerful advantage over classical EEG-correlated analysis. As we demonstrated in the experiments, artifactual influences may

be isolated in separate components, which could reduce their impact on IED mapping in the brain. Additionally, we encountered cases where two components were correlated to the IED occurrences: the component with the highest temporal correlation to a reference IED time course then correctly revealed the localization of the IOZ, while the other component presumably modeled the propagation of IEDs to remote brain regions. This observation is analogous to the finding in Hunyadi et al. (2016), where a different type of CMTF was applied to average EEG waveforms of IEDs and statistical BOLD maps, which revealed a dissociation between the early IED spike and the subsequent wave, which were related to the onset and spread of the IEDs, respectively. Since we transformed the data with a time-frequency transform that used windows of length TR, our algorithm is unable to unravel different phases within one IED, since they occur in a shorter time frame. However, we identified these different IED-related sources by their significantly correlated temporal signatures, and their distinct spatial and spectral profiles. While we did not impose nonnegativity constraints, many estimated EEG spectral and spatial signatures were approximately nonnegative. This need not be the case, however, since the EEG data are normalized in a way such that the resulting signatures would reveal relative increase/decrease, rather than absolute time-varying spectral power (Mareček et al., 2016). For example, if a certain component is associated with a power increase in one spectral band, and a simultaneous power decrease in another band, this would be reflected in a spectrum with both a positive and a negative peak.

Practical considerations

The end-to-end sCMTF pipeline can provide a richer set of results compared to classical EEG-correlated fMRI analysis. In this respect, it is a more powerful data exploration tool. The tradeoff to be made is that significant computation time goes into the sCMTF and subsequent inference—if one wants to apply it as rigorously as we have done in the current experiments. We seem to be doing a lot of unnecessary work, by computing the sCMTF factors for several numbers of sources, and by repeating the optimization several times for a fixed number of sources. Unfortunately, both ways of repetition seem required to obtain robust results, as we have argued in Appendix A and Appendix B. However, the EEG-only CP decomposition, which lies at the heart of our initialization strategy, seemed very robust: we found highly similar EEG signatures for almost all random initializations. Probably, this is thanks to the use of the powerful Gauss–Newton-type optimization. Hence, fewer repetitions of the sCMTF may be already sufficient to arrive at the same robust results. Despite the very reproducible EEG signatures in the initial CP decompositions, we still performed 50 repetitions of the sCMTF, each time slightly varying the initial HRF parameters. As such, we believe our findings are reasonably robust to poor initialization of the HRFs. Performing the sCMTF for many choices of R may still be required, as the quality of the result depends on the extraction of an appropriate number of sources. In our study, no prohibitive computations were needed, since the heuristic selection procedure preferred low ranks, which was still sufficient to model the IED-related dynamics. To also capture more resting-state activity, the candidate ranks and the model selection process could be chosen differently. Furthermore, we have demonstrated in our experiments that the summary metrics (sensitivity for localizing the IOZ based on different statistical scores) are fairly robust to the choice of R , although the estimated signatures themselves differ.

For many patients, the available data were split across multiple runs (i.e., with a few minutes break in between), and we opted to temporally concatenate data over runs, as explained in Section 2.2. While this violates the coupling model based on HRFs for time samples near the boundaries, we consider the effect minimal, given that the number of those ‘affected’ time samples represents a very tiny fraction of the whole time series. However, a more rigorous approach would be to ‘inverse-impute’ these samples and consider them as missing values: as such, they are ignored during the sCMTF optimization and will not affect the results.

Strategies to alleviate the computational demand

Due to the repeated decomposition and the nature of the nonparametric inference, the computations are highly parallelizable. For a typical dataset with available IED annotations, and with the parameters we have used for this study, the end-to-end computation for one patient took at most five hours on a machine with twelve cores. To alleviate the computational burden, we have parcellated the fMRI data into 246 regions, based on the Brainnetome atlas (Fan et al., 2016). This is clearly suboptimal, as the atlas is not patient-specific, and is mostly designed to study healthy brains. There is a serious risk for partial volume effects, in which the IOZ is scattered over several ROIs. As such, the IED-related BOLD changes in the part of the IOZ that falls within a certain ROI may get swamped by the remaining BOLD fluctuations within the ROI delineation. Hence, we hope to be able overcome this problem, either by algorithmic improvements, including a speed up of the optimization, or by the use of a patient-specific parcellation or PCA-like compression of the fMRI data. As of yet, it is hard to say whether the fixed atlas had an adverse effect on the results, and it is not so straightforward to compare the statistical maps from this study to maps which are voxel-based. We are currently pursuing experiments in which we employ a hierarchical parcellation: in a first step, the BOLD time series are grouped (but not yet averaged) according to the Brainnetome atlas; subsequently, we use spectral clustering to further refine each Brainnetome parcel based on the correlation matrix of its BOLD time series. As such, this hybrid approach combines a fixed, coarse-grained atlas with a further data-driven subdivision, which can mitigate partial volume effects, while still providing a significant data compression. For patients with lesions, a customized parcellation can be used, in which the lesion itself coincides with one parcel or with the union of several parcels. Alternatively, it is possible to achieve a data reduction while still preserving voxelwise BOLD signals, by limiting the scope of the sCMTF to an a priori defined ROI (e.g., based on a clinical hypothesis stemming from other modalities).

Summary

In summary, we have developed and empirically validated a fully integrated framework for EEG–fMRI data fusion, which yielded a rich characterization of the interictal activity in time, space, and frequency, and which accounts for variations in neurovascular coupling over the brain. Such spatial variation can be exploited to obtain complementary information for IOZ localization. The ability to separate local (de)activation of IEDs from local deviations in the HRF makes the sCMTF a powerful tool for exploratory analysis of interictal EEG and fMRI data. This approach may also be used for RSN analysis, a field where estimation bias due to HRF variation has so far largely been ignored.

Our complete MATLAB code to execute the pipeline is available at <https://github.com/svaneynd/structured-cmtf>.

Acknowledgment

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 (no. 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. This research received funding from the Flemish Government (AI Research Program). SVE and SVH are affiliated to Leuven.AI - KU Leuven institute for AI, B-3000, Leuven, Belgium. This research furthermore received funding from the Bijzonder Onderzoeksfonds KU Leuven (BOF) under the project numbers C24/15/036 and C24/18/097; from the Agentschap Innoveren en Ondernemen (VLAIO) under the project number 150466; from the EU for Horizon 2020 projects 766456, 813120 and 813483; and from EIT for the project SeizeIT (no. 19263).

We thank the reviewers for their in-depth, insightful feedback, which has helped to improve the manuscript’s quality.

Declaration of Competing Interest

Declarations of interest: none

Appendix A. Nonlinear fitting of the sCMTF model

A1. Normalizing the EEG-fMRI data

After construction of the EEG spectrogram tensors and BOLD signal matrix, we perform a normalization procedure, to ensure that each frequency bin, each channel, and each ROI contribute the same variance to the data (cfr. Sections 2.4.1 and 2.4.2). From all mode-1 fibers in the EEG tensor and fMRI matrix (i.e., in the temporal mode), we subtract the mean. Subsequently, for the fMRI data, we divide each ROI's time series by its standard deviation. For the EEG data, we carry out this latter step alternatingly over the spectral and spatial mode, until convergence, since normalizing over one mode generally does not preserve normalization in the other modes (Bro, 1997).

This normalization is less stringent than z-scoring each mode-1 fiber, since contrary to the latter, the former does not enforce equal variance for each time course, but only equal variance for whole spatiotemporal or spectrotemporal slabs of the tensor. After computing the structured factorization, the signatures in \mathbf{G} and \mathbf{M} can be converted back to the original scale of the data by premultiplying with Σ_g^{-1} and Σ_m^{-1} , if desired (which we have not done).

A2. Accommodating noncausal HRFs

In Section 2.4.3, we derived the implementation of the convolution with an HRF as a left multiplication of the temporal signatures \mathbf{S} with a Toeplitz matrix, whose diagonals hold the HRF samples. For a causal convolution, in which the BOLD signal strictly lags its neural correlate, $\mathbf{H}_k(i, j) = h_k(i - j)$ if $i - j \geq 0$, $\mathbf{H}_k(i, j) = 0$ otherwise, hence the matrix is lower triangular. This is the situation depicted also in Fig. 1.

However, a recurring observation is that BOLD changes can be observed that precede the IEDs themselves (Hawco et al., 2007; Jacobs et al., 2009; Moeller et al., 2008; Pittau et al., 2011). Hence, we allowed noncausal HRFs that start at most 4 samples before the EEG, which allows for BOLD responses preceding the IEDs by up to 10 s at a typical TR of 2.5 s.

A3. Initialization and optimization

Since the cost function J in (14) is nonconvex, any optimization procedure can only guarantee to converge to a local optimum, hence selecting a good starting point is crucial to obtain a reliable solution.

Firstly, we decomposed the EEG data \mathcal{X} individually according to the CP or PARAFAC model (Miwakeichi et al., 2004; Mørup et al., 2006), to obtain a good initialization for the factors $[\mathbf{S}, \mathbf{G}, \mathbf{M}]$ in the sCMTF model. To this end, we used a Gauss-Newton algorithm (cpd_nls with 2000 iterations, 400 conjugate gradient iterations for the step computation, and tolerance on the relative cost function update of 10^{-8} , in Tensorlab 3.0 (Vervliet et al., 2016)), which we ran 50 times, from randomly drawn initial factors. We observed that the resulting factors lied very often close together over runs, indicating the algorithm had found a robust solution.

We always employed $K = 3$ HRFs, which we manually initialized. To assess whether the eventual sCMTF solution was also robust to the initialization of the HRF basis functions, we used a slightly different set of HRF-generating parameters $\theta^{(k)}$ in each repetition of the optimization. Fig. A.10 shows some typical HRF waveforms, which are used to generate the Toeplitz blocks in Fig. 1b.

From there, we initialized also the fMRI factors in the sCMTF model in (11)–(13). We constructed a flattened ‘design matrix’ $\mathbf{D} = [\mathbf{H}_1\mathbf{S} \dots \mathbf{H}_K\mathbf{S}]$ in (13) and obtained a rough estimate for $\mathbf{B}^T \odot \mathbf{V}^T$ as $\mathbf{U}^T = \mathbf{D}^+ \mathbf{Y}$ via regression—albeit this does not yet disentangle \mathbf{B} and \mathbf{V} .




Fig. A1. Different initializations of the HRF basis functions are used in every repetition of the sCMTF fitting procedure. In each repetition, one early-peaking, one mid-peaking and one late-peaking HRF were sampled from probability distributions of the HRF's generating parameters ($\theta^{(1)}$, $\theta^{(2)}$ and $\theta^{(3)}$, respectively) that was created by applying some multiplicative noise to baseline parameters (shown as the bold waveforms). The support extends to four negative samples, which gives the model the freedom to fit noncausal HRFs (relatively to the synchronized EEG).

Input : unnormalized EEG tensor $\mathcal{X} \in \mathbb{R}^{I_s \times I_g \times I_m}$
Output: normalized tensor $\mathcal{X} \in \mathbb{R}^{I_s \times I_g \times I_m}$ and total scaling matrices Σ_g and Σ_m

Initialization
 $\Sigma_g \leftarrow \mathbf{I}_{I_g}$, $\Sigma_m \leftarrow \mathbf{I}_{I_m}$ // identity matrices
while not converged do
 for $i_g = 1$ **to** I_g **do**
 $\sigma_{i_g} \leftarrow \sqrt{\sum_{i_s=1}^{I_s} \sum_{i_m=1}^{I_m} X(i_s, i_g, i_m)^2}$
 $\Sigma_g(i_g, i_g) \leftarrow \sigma_{i_g} \Sigma_g(i_g, i_g)$
 $X(:, i_g, :) \leftarrow \sigma_{i_g}^{-1} X(:, i_g, :)$
 end for
 for $i_m = 1$ **to** I_m **do**
 $\sigma_{i_m} \leftarrow \sqrt{\sum_{i_s=1}^{I_s} \sum_{i_g=1}^{I_g} X(i_s, i_g, i_m)^2}$
 $\Sigma_m(i_m, i_m) \leftarrow \sigma_{i_m} \Sigma_m(i_m, i_m)$
 $X(:, :, i_m) \leftarrow \sigma_{i_m}^{-1} X(:, :, i_m)$
 end for
end while
return X , Σ_g , Σ_m

Algorithm 1. Tensor scale normalization

To obtain initializations for the individual spatial factors, we exploit the fact that in every ROI i_v , the Khatri-Rao product of the i_v -th columns of \mathbf{B}^T and \mathbf{V}^T corresponds to a rank-1 constraint when folded into a $K \times R$ matrix (Beckmann and Smith, 2005; Boussé et al., 2018); hence, a rank-1 truncated singular value decomposition of the folded i_v -th column of \mathbf{U}^T leads to the desired vectors (Beckmann and Smith, 2005), which are further refined via a constrained Gauss-Newton algorithm (Boussé et al., 2018). We approximated the residual of the fMRI data under the initialized (coupled) factors using a rank- Q truncated SVD to capture fMRI nuisances. The parameter Q was chosen as twice the number of acquisition runs that had been done for a subject.

After each of the 50 runs of the initialization procedure, we iteratively optimized (14) with a quasi-Newton algorithm (sdf_minf with 1000 iterations, and tolerance on the relative cost function update of 10^{-8} , in Tensorlab 3.0 (Vervliet et al., 2016)). Both \mathcal{X} and \mathbf{Y} were divided by their Frobenius norm, such that afterwards $\|\mathcal{X}\|_F = \|\mathbf{Y}\|_F = 1$, and

we chose $\beta_x = \beta_y = 1$, such that their fit had equal contribution to the cost function. For the regularization penalties, we pick $\gamma_x = \gamma_y = 10^{-3}$, as in (Acar et al., 2014).

Appendix B. Model selection

We fitted the sCMTF model to each patient’s data for varying number of sources (rank), i.e., $R = 1 \dots 6$. For each choice of R , we ran the optimization procedure 50 times, as explained in A.3. Afterwards, the results need to be aggregated, such that clear conclusions on the sources of interest can be drawn. This involves several steps, which we explain below in chronological order.

B1. Sign and scale standardization of factor estimates

Some of the factors that are estimated via minimization of (14) are subject to sign and scale ambiguities, which are inevitable in many BSS contexts (Kolda and Bader, 2009; Sidiropoulos et al., 2017). In the EEG factor model in (8), the factor vectors \mathbf{s}_r , \mathbf{g}_r , and \mathbf{m}_r belonging to the same component r may be multiplied with arbitrary scaling factors whose product is one, without altering the goodness of fit. Similarly, in the fMRI factor model in (11), sign and scale are exchangeable between corresponding columns of \mathbf{S} and \mathbf{V} , and between rows of \mathbf{V} and rows of \mathbf{B} . However, to conduct proper statistical inference on each source r ’s spatial amplitudes, the elements in \mathbf{v}_r must be calibrated, in the sense that it must be possible to compare them across ROIs, and determining the sign is crucial to distinguish local activation from deactivation. Hence, we sequentially fix these ambiguities as follows:

1. For every component r :
 - (a) \mathbf{s}_r and \mathbf{g}_r are rescaled to unit ℓ_2 -norm, and \mathbf{m}_r is counterscaled
 - (b) for both \mathbf{s}_r and \mathbf{g}_r , the sign is flipped if the sum of squares of the negative elements exceeds that of the positive elements; the sign of \mathbf{m}_r is adjusted to preserve the global sign of the EEG rank-1 term
 - (c) if the sign of \mathbf{s}_r was flipped, also the sign of \mathbf{v}_r is flipped to preserve the global sign of the r -th fMRI block term
2. For every ROI i_v (cfr. also (Calhoun et al., 2004)):
 - (a) the local HRF \mathbf{h}_{i_v} is reconstructed using (6)
 - (b) the i_v -th row of \mathbf{B} is rescaled and sign-corrected as to make \mathbf{h}_k unit ℓ_1 -norm and as to ensure that the HRF’s largest overshoot precedes the largest undershoot; the i_v -th row of \mathbf{V} is counterscaled

B2. Stability analysis

To assess the reproducibility of the factors, we use the graph-structured clustering algorithm that we proposed in Van Eyndhoven et al. (2019b) and briefly summarize here. We represent the factor sets for all 50 repetitions of the fitting procedure as $[[\mathbf{S}, \mathbf{G}, \mathbf{M}, \mathbf{V}]]$, and use a threshold of 0.85 to construct a binary link matrix that encodes similarities between components from different runs of the optimization (empirically, we found that this threshold led to acceptable cluster definitions in this context). Via low-rank matrix approximation of this link matrix, we then obtained clusters of components that were encountered in varying numbers of repetitions. High cardinality of a cluster is then a sign that the involved component is very reproducible or ‘stable’, since it is part of the factor set upon convergence, in many repetitions. We suggest to assign higher trust in such components, as opposed to components in small clusters, which are likely specific to one (potentially poor) local minimum. For the further steps, we condensed each cluster to one of its components, i.e., its centroid. In each cluster, the centroid

component is defined as the component which has the largest accumulated similarity with all other components in the cluster. By extension, the centroid repetition is defined as the repetition to which the centroid component belongs.

B3. IED component selection

Out of the centroids of the clustered components, we identified the (most) IED-related component as the one whose temporal signature s_r was most correlated to a reference time course, which is constructed as the average over channels of the MWF’s output signal’s time-varying (broadband) power. This reference is the BOLD predictor we have proposed for EEG-correlated fMRI analysis in (Van Eyndhoven et al., 2019a), and which provides a good baseline for identifying temporal dynamics that are timelocked to the IEDs.

B4. Choice of the rank R of the factorization

After the previous steps have been carried out for each setting of R , we are left to select the rank $R \in \{1, 2, 3, 4, 5, 6\}$ whose set of results we proceed with. We heuristically determined an appropriate value \hat{R} by selecting the model which fulfills several criteria:

1. *high core consistency of the EEG decomposition*
We compute the core consistency diagnostic (CorConDia) (Bro and Kiers, 2003) for the EEG tensor in combination with its estimated factor set $[[\mathbf{S}, \mathbf{G}, \mathbf{M}]]$ from the centroid repetition. The consistency describes how suitable a rank- R CPD is for the given tensorial data and given factors, and is expressed as a percentage (100% being a very adequate model, and percentages below 70–80% indicating that the model is not appropriate)⁶.
2. *reproducible IED-related component*
We count the number of repetitions in the cluster that was most related to the IED (cfr. B.2 and B.3), and used this as a measure of reproducibility. We rejected clusters whose cardinality was lower than 10.
3. *similarity to a reference IED time course*
We track the correlation between the IED-related component’s temporal signature s_{IED} and the reference temporal signature s_{ref} , as explained in B.3. We expect higher correlations to signify a more suitable model, since s_{ref} generally led to good results in our previous study (Van Eyndhoven et al., 2019a).
4. *high significance in the IED-related spatial map*
We track the highest pseudo t-value in the SnPM that was created based on the IED-related component’s spatial signature \mathbf{v}_{IED} (cfr. Section 2.6). A high statistical score indicates a good model fit for the IED-related component (Abreu et al., 2018).

⁶ CorConDia is a popular and robust model selection tool for tensor decompositions (Acar et al., 2007; Liu et al., 2016; Miwakeichi et al., 2004; Mørup and Hansen, 2009; Papalexakis, 2016). To compute it, first the core tensor which is most appropriate (in minimum mean squared error sense) for the given data and CPD-derived factors is estimated. Subsequently, CorConDia is computed as the fraction of the core tensor’s sum of squares which is due to off-superdiagonal elements. When, for a given set of factor matrices, the CP structure is indeed ideal, the core tensor is superdiagonal and CorConDia equals 100%. Note that for a rank-1 model, this notion is meaningless, since the core tensor is a scalar, and CorConDia would trivially be 100% always.

Table B1

For every patient (ID 1–12), the sCMTF model can be fitted for a varying number of sources or rank R . We select a ‘good’ value for R post hoc, based on four criteria which are checked intra-patient: 1) the core consistency of the EEG tensor decomposition should be high ($> 70\%$); 2) the IED-related source should be found in sufficiently many (≥ 10) of the 50 repetitions of the estimation procedure; 3) the correlation of the IED-related source’s temporal signature with the reference time course, namely the MWF’s broadband envelope, is preferably high; 4) the maximal pseudo t-statistic for the IED-related source’s spatial signature is preferably high.

ID	$R =$	core consistency diagnostic (%)						reproducibility (# repetitions in IED cluster)						correlation of s_{IED} with reference MWF envelope s_{ref}						maximal t-statistic of v_{IED}						selected rank
		1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	
p01	-	100.0	99.6	97.1	90.0	76.8	25	19	14	7	16	20	0.35	0.91	0.38	0.94	0.97	0.97	14.2	20.0	21.5	16.0	7.7	20.1	$\hat{R} = 6$	
p02	-	100.0	94.9	58.0	19.6	7.2	14	23	27	26	17	18	0.10	0.93	0.95	0.83	0.96	0.88	8.7	17.0	23.4	22.6	17.7	16.9	$\hat{R} = 3$	
p03	-	100.0	95.6	70.7	29.0	28.4	29	23	24	21	23	19	0.29	0.93	0.92	0.93	0.93	0.93	15.2	16.5	14.9	14.3	13.9	15.0	$\hat{R} = 2$	
p04	-	100.0	87.0	74.4	37.3	12.3	25	22	15	27	15	20	0.09	0.20	0.61	0.71	0.68	0.81	9.7	6.7	15.4	14.3	11.7	23.4	$\hat{R} = 4$	
p05	-	100.0	98.7	94.2	93.2	89.7	3	14	14	21	13	17	0.28	0.85	0.00	0.01	0.92	0.01	19.6	5.5	6.5	4.4	9.0	4.6	$\hat{R} = 5$	
p06	-	100.0	94.6	80.9	33.8	-409	15	21	14	14	11	13	0.55	0.92	0.87	0.80	0.80	0.22	9.9	17.1	9.9	9.6	16.4	36.9	$\hat{R} = 2$	
p07	-	100.0	97.7	96.7	-76.0	-0.6	18	17	12	15	11	12	0.09	0.74	0.80	0.80	0.80	0.22	7.4	10.0	9.4	8.6	11.2	9.9	$\hat{R} = 4$	
p08	-	100.0	99.7	30.6	-67.2	-178	12	21	21	11	14	23	0.61	0.95	0.44	0.95	0.94	0.94	13.8	16.0	10.1	20.3	15.5	18.0	$\hat{R} = 2$	
p09	-	100.0	98.8	95.9	90.1	23.6	21	27	21	15	19	19	0.19	0.66	0.67	0.13	0.48	0.49	14.7	22.5	21.9	10.1	8.0	19.2	$\hat{R} = 2$	
p10	-	100.0	98.0	95.0	91.7	80.3	23	34	13	24	25	14	0.47	0.29	0.96	0.96	0.96	0.19	11.9	12.9	10.0	8.6	9.9	12.8	$\hat{R} = 5$	
p11	-	100.0	97.9	78.3	-307	49.1	21	18	12	22	13	12	0.65	0.91	0.50	0.74	0.66	0.67	18.8	12.8	12.0	28.9	18.6	12.6	$\hat{R} = 2$	
p12	-	100.0	97.3	89.0	59.3	69.4	23	14	14	12	15	12	0.78	0.79	0.60	0.07	0.47	0.50	15.8	15.0	10.9	7.9	10.2	5.5	$\hat{R} = 2$	

Appendix C. Computing HRF deviation metrics

C1. HRF extremity

The extremity of a specific ROI's HRF is computed as one minus the average of the absolute values of the Pearson correlation between the HRF waveform and all other ROIs' HRFs waveforms. I.e., for the j_v -th ROI, the extremity is computed as

$$\text{extremity}(j_v) = 1 - \frac{1}{I_v - 1} \sum_{i_v \neq j_v} |\text{corr}(h_{i_v}, h_{j_v})| \quad (\text{C.1})$$

Only the first twenty samples (~ 50 s) are considered. Note that the extremity does not change if the (global, not samplewise) sign (polarity) of one or more HRFs changes.

C2. HRF entropy

The entropy of a specific ROI's HRF is computed as the negative logarithm of the probability of this HRF, conditional on all other ROIs' HRFs. For example, we first estimated a probability density in HRF space based on all other ROIs' HRFs, and then evaluated this density at the HRF of the ROI under inspection. From every HRF, we considered the first twenty samples, and then estimated a nonparametric multivariate kernel density in 20-dimensional space, by placing a multivariate Gaussian probability kernel at the location of each HRF except one. We made this entropy metric insensitive to the signs of the HRFs, by extending the set of HRF waveforms by their flipped counterparts, and computing the nonparametric density using the resulting $2(I_v - 1)$ HRFs in a leave-one-ROI-out fashion.

$$\text{entropy}(j_v) = -\log\left(\frac{1}{2(I_v-1)} \sum_{i_v \neq j_v} \left(K(h_{i_v}, h_{j_v}; \Sigma) + K(h_{i_v}, -h_{j_v}; \Sigma) \right)\right), \quad (\text{C.2})$$

in which $K(h_{i_v}, h_{j_v}; \Sigma)$ is a Gaussian kernel distance, which is proportional to

$$\exp\left(-\frac{1}{2}(\mathbf{h}_{i_v} - \mathbf{h}_{j_v})^T \Sigma^{-1}(\mathbf{h}_{i_v} - \mathbf{h}_{j_v})\right), \quad (\text{C.3})$$

in which \mathbf{h}_{i_v} and \mathbf{h}_{j_v} are column vectors that store the twenty first samples of the HRFs $h_{i_v}(t)$ and $h_{j_v}(t)$, and Σ is a diagonal covariance or bandwidth matrix. We used Silverman's heuristic to set the kernel bandwidths for each individual dimension, corresponding to one HRF time sample (Silverman, 1986). I.e., the n -th bandwidth σ_{nn}^2 , which corresponds to the HRF amplitudes at sample n , is given by

$$\sigma_{nn}^2 = \left(\frac{4}{20+2}\right)^{\frac{2}{20+4}} (2(I_v - 1))^{-\frac{2}{20+4}} s_n^2, \quad (\text{C.4})$$

in which s_n^2 is the observed variance (over ROIs) of the HRFs' amplitudes at the n -th sample.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.neuroimage.2020.117652](https://doi.org/10.1016/j.neuroimage.2020.117652)

Credit authorship contribution statement

Simon Van Eyndhoven: Conceptualization, Methodology, Software, Formal analysis, Writing - original draft. **Patrick Dupont:** Conceptualization, Investigation, Writing - original draft. **Simon Tousseyn:** Data curation, Investigation, Writing - original draft. **Nico Vervliet:** Software, Writing - original draft. **Wim Van Paesschen:** Methodology, Supervision, Writing - original draft. **Sabine Van Huffel:** Conceptualization, Supervision, Writing - original draft. **Borbála Hunyadi:** Conceptualization, Supervision, Writing - original draft.

References

- Abreu, R., Leal, A., da Silva, F.L., Figueiredo, P., 2018. EEG synchronization measures predict epilepsy-related BOLD-fMRI fluctuations better than commonly used univariate metrics. *Clin. Neurophys.* 129 (3), 618–635.
- Abreu, R., Leite, M., Leal, A., Figueiredo, P., 2016. Objective selection of epilepsy-related independent components from EEG data. *J. Neurosci. Methods* 258, 67–78.
- Acar, E., Aykut-Bingol, C., Bingol, H., Bro, R., Yener, B., 2007. Multiway analysis of epilepsy tensors. *Bioinformatics* 23 (13), i10–i18.
- Acar, E., Levin-Schwartz, Y., Calhoun, V.D., Adali, T., 2017. Acmtf for fusion of multi-modal neuroimaging data and identification of biomarkers. In: *Proceedings of the 25th European Signal Processing Conference (EUSIPCO)*. IEEE, pp. 643–647.
- Acar, E., Papalexakis, E.E., Gürdeniz, G., Rasmussen, M.A., Lawaetz, A.J., Nilsson, M., Bro, R., 2014. Structure-revealing data fusion. *BMC Bioinform.* 15 (1), 239.
- Acar, E., Schenker, C., Levin-Schwartz, Y., Calhoun, V.D., Adali, T., 2019. Unraveling diagnostic biomarkers of schizophrenia through structure-revealing fusion of multi-modal neuroimaging data. *Front. Neurosci.* 13, 416.
- Aguirre, G.K., Zarahn, E., D'esposito, M., 1998. The variability of human, bold hemodynamic responses. *Neuroimage* 8 (4), 360–369.
- Allen, P.J., Polizzi, G., Krakow, K., Fish, D.R., Lemieux, L., 1998. Identification of EEG events in the MR scanner: the problem of pulse artifact and a method for its subtraction. *Neuroimage* 8 (3), 229–239.
- An, D., Fahoum, F., Hall, J., Olivier, A., Gotman, J., Dubeau, F., 2013. Electroencephalography/functional magnetic resonance imaging responses help predict surgical outcome in focal epilepsy. *Epilepsia* 54 (12), 2184–2194.
- Bagshaw, A.P., Aghakhani, Y., Bénar, C.-G., Kobayashi, E., Hawco, C., Dubeau, F., Pike, G.B., Gotman, J., 2004. EEG-fMRI of focal epileptic spikes: analysis with multiple haemodynamic functions and comparison with gadolinium-enhanced mr angiograms. *Hum. Brain Mapp.* 22 (3), 179–192.
- Beckmann, C.F., Smith, S.M., 2005. Tensorial extensions of independent component analysis for multisubject fMRI analysis. *Neuroimage* 25 (1), 294–311.
- Behzadi, Y., Restom, K., Liau, J., Liu, T.T., 2007. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage* 37 (1), 90–101.
- Bénar, C.-G., Gross, D.W., Wang, Y., Petre, V., Pike, B., Dubeau, F., Gotman, J., 2002. The BOLD response to interictal epileptiform discharges. *Neuroimage* 17 (3), 1182–1192.
- Boussé, M., Vervliet, N., Domanov, I., Debals, O., De Lathauwer, L., 2018. Linear systems with a canonical polyadic decomposition constrained solution: Algorithms and applications. *Numer. Linear Algebra Appl.* 25 (6), e2190.
- Bro, R., 1997. PARAFAC. Tutorial and applications. *Chemometr. Intell. Lab. Syst.* 38 (2), 149–171.
- Bro, R., Kiers, H.A., 2003. A new efficient method for determining the number of components in PARAFAC models. *J. Chemometr.: J. Chemometr. Soc.* 17 (5), 274–286.
- Bullmore, E., Long, C., Suckling, J., Fadili, J., Calvert, G., Zelaya, F., Carpenter, T.A., Brammer, M., 2001. Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains. *Hum. Brain Mapp.* 12 (2), 61–78.
- Caballero-Gaudes, C., Van de Ville, D., Grouiller, F., Thornton, R., Lemieux, L., Seeck, M., Lazeyras, F., Vulliemoz, S., 2013. Mapping interictal epileptic discharges using mutual information between concurrent EEG and fMRI. *Neuroimage* 68, 248–262.
- Calhoun, V.D., Liu, J., Adali, T., 2009. A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. *Neuroimage* 45 (1), S163–S172.
- Calhoun, V.D., Stevens, M., Pearlson, G., Kiehl, K., 2004. fMRI analysis with the general linear model: removal of latency-induced amplitude bias by incorporation of hemodynamic derivative terms. *Neuroimage* 22 (1), 252–257.
- Chatzichristos, C., Davies, M., Escudero, J., Kofidis, E., Theodoridis, S., 2018. Fusion of EEG and fMRI via soft coupled tensor decompositions. In: *Proceedings of the 26th European Signal Processing Conference (EUSIPCO)*. IEEE, pp. 56–60.
- De Vos, M., Vergult, A., De Lathauwer, L., De Clercq, W., Van Huffel, S., Dupont, P., Palmieri, A., Van Paesschen, W., 2007. Canonical decomposition of ictal scalp EEG reliably detects the seizure onset zone. *NeuroImage* 37 (3), 844–854.
- Elbau, I.G., Brückmeier, B., Uhr, M., Arloth, J., Czamara, D., Spormaker, V.I., Czisch, M., Stephan, K.E., Binder, E.B., Sämann, P.G., 2018. The brains hemodynamic response function rapidly changes under acute psychosocial stress in association with genetic and endocrine stress response markers. *Proc. Natl. Acad. Sci.* 115 (43), E10206–E10215.
- Fan, L., Li, H., Zhuo, J., Zhang, Y., Wang, J., Chen, L., Yang, Z., Chu, C., Xie, S., Laird, A.R., et al., 2016. The human brainnetome atlas: a new brain atlas based on connectome architecture. *Cerebral Cortex* 26 (8), 3508–3526.
- Formaggio, E., Storti, S.F., Bertoldo, A., Manganotti, P., Fiaschi, A., Toffolo, G.M., 2011. Integrating EEG and fMRI in epilepsy. *Neuroimage* 54 (4), 2719–2731.
- Friston, K.J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M., Turner, R., 1998. Event-related fMRI: characterizing differential responses. *Neuroimage* 7 (1), 30–40.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.-P., Frith, C.D., Frackowiak, R.S., 1994. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2 (4), 189–210.
- Glover, G.H., 1999. Deconvolution of impulse response in event-related BOLD fMRI. *Neuroimage* 9 (4), 416–429.
- Gotman, J., 2008. Epileptic networks studied with EEG-fMRI. *Epilepsia* 49, 42–51.
- Grouiller, F., Thornton, R.C., Groening, K., Spinelli, L., Duncan, J.S., Schaller, K., Sinatchkin, M., Lemieux, L., Seeck, M., Michel, C.M., et al., 2011. With or without spikes: localization of focal epileptic activity by simultaneous electroencephalography and functional magnetic resonance imaging. *Brain* 134 (10), 2867–2886.
- Grouiller, F., Vercueil, L., Krainik, A., Segebarth, C., Kahane, P., David, O., 2010. Characterization of the hemodynamic modes associated with interictal epileptic activity using a deformable model-based analysis of combined EEG and functional MRI recordings. *Hum. Brain Mapp.* 31 (8), 1157–1173.

- Handwerker, D.A., Ollinger, J.M., D'Esposito, M., 2004. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage* 21 (4), 1639–1651.
- Hawco, C.S., Bagshaw, A.P., Lu, Y., Dubeau, F., Gotman, J., 2007. BOLD changes occur prior to epileptic spikes seen on scalp EEG. *Neuroimage* 35 (4), 1450–1458.
- van Houdt, P.J., de Munck, J.C., Leijten, F.S., Huiskamp, G.J., Colon, A.J., Boon, P.A., Ossenblok, P.P., 2013. EEG-fMRI correlation patterns in the presurgical evaluation of focal epilepsy: a comparison with electrocorticographic data and surgical outcome measures. *Neuroimage* 75, 238–248.
- Hunyadi, B., Camps, D., Sorber, L., Van Paesschen, W., De Vos, M., Van Huffel, S., De Lathauwer, L., 2014. Block term decomposition for modelling epileptic seizures. *EURASIP J. Adv. Signal Proc.* 2014 (1), 139.
- Hunyadi, B., Dupont, P., Van Paesschen, W., Van Huffel, S., 2017. Tensor decompositions and data fusion in epileptic electroencephalography and functional magnetic resonance imaging data. *Wiley Interdisc. Rev.: Data Min. Knowl. Discov.* 7 (1), e1197.
- Hunyadi, B., Tousseyn, S., Dupont, P., Van Huffel, S., De Vos, M., Van Paesschen, W., 2015. A prospective fMRI-based technique for localising the epileptogenic zone in presurgical evaluation of epilepsy. *Neuroimage* 113, 329–339.
- Hunyadi, B., Van Paesschen, W., De Vos, M., Van Huffel, S., 2016. Fusion of electroencephalography and functional magnetic resonance imaging to explore epileptic network activity. In: *Proceedings of the 24th European Signal Processing Conference (EUSIPCO)*. IEEE, pp. 240–244.
- Jacobs, J., Hawco, C., Kobayashi, E., Boor, R., LeVan, P., Stephani, U., Siniatchkin, M., Gotman, J., 2008. Variability of the hemodynamic response as a function of age and frequency of epileptic discharge in children with epilepsy. *Neuroimage* 40 (2), 601–614.
- Jacobs, J., LeVan, P., Moeller, F., Boor, R., Stephani, U., Gotman, J., Siniatchkin, M., 2009. Hemodynamic changes preceding the interictal EEG spike in patients with focal epilepsy investigated using simultaneous EEG-fMRI. *Neuroimage* 45 (4), 1220–1231.
- Khoo, H.M., Hao, Y., von Ellenrieder, N., Zazubovits, N., Hall, J., Olivier, A., Dubeau, F., Gotman, J., 2017. The hemodynamic response to interictal epileptic discharges localizes the seizure-onset zone. *Epilepsia* 58 (5), 811–823.
- Kolda, T.G., Bader, B.W., 2009. Tensor decompositions and applications. *SIAM Rev.* 51 (3), 455–500.
- Lahat, D., Adali, T., Jutten, C., 2015. Multimodal data fusion: an overview of methods, challenges, and prospects. *Proc. IEEE* 103 (9), 1449–1477.
- Lemieux, L., Laufs, H., Carmichael, D., Paul, J.S., Walker, M.C., Duncan, J.S., 2008. Non-canonical spike-related BOLD responses in focal epilepsy. *Hum. Brain Map.* 29 (3), 329–345.
- Lemieux, L., Salek-Haddadi, A., Josephs, O., Allen, P., Toms, N., Scott, C., Krakow, K., Turner, R., Fish, D.R., 2001. Event-related fMRI with simultaneous and continuous EEG: description of the method and initial case report. *Neuroimage* 14 (3), 780–787.
- Lindquist, M.A., Loh, J.M., Atlas, L.Y., Wager, T.D., 2009. Modeling the hemodynamic response function in fMRI: efficiency, bias and mis-modeling. *Neuroimage* 45 (1), S187–S198.
- Lindquist, M.A., Wager, T.D., 2007. Validity and power in hemodynamic response modeling: a comparison study and a new approach. *Hum. Brain Map.* 28 (8), 764–784.
- Liu, K., Da Costa, J.P.C., So, H.C., Huang, L., Ye, J., 2016. Detection of number of components in CANDECOMP/PARAFAC models via minimum description length. *Digit. Signal Proc.* 51, 110–123.
- Makni, S., Beckmann, C., Smith, S., Woolrich, M., 2008. Bayesian deconvolution fMRI data using bilinear dynamical systems. *Neuroimage* 42 (4), 1381–1396.
- Mantini, D., Puerucci, M.G., Del Gratta, C., Romani, G.L., Corbetta, M., 2007. Electrophysiological signatures of resting state networks in the human brain. *Proc. Natl. Acad. Sci.* 104 (32), 13170–13175.
- Mareček, R., Lamoš, M., Labounek, R., Bartoň, M., Slavíček, T., Mikl, M., Rektor, I., Brázdil, M., 2017. Multiway array decomposition of EEG spectrum: implications of its stability for the exploration of large-scale brain networks. *Neural Comput.* 29 (4), 968–989.
- Mareček, R., Lamoš, M., Mikl, M., Bartoň, M., Fajkus, J., Rektor, I., Brázdil, M., 2016. What can be found in scalp EEG spectrum beyond common frequency bands. *EEG-fMRI study*. *J. Neural Eng.* 13 (4), 046026.
- Martínez-Montes, E., Valdés-Sosa, P.A., Miwakeichi, F., Goldman, R.I., Cohen, M.S., 2004. Concurrent EEG/fMRI analysis by multiway partial least squares. *Neuroimage* 22 (3), 1023–1034.
- Miwakeichi, F., Martínez-Montes, E., Valdés-Sosa, P.A., Nishiyama, N., Mizuhara, H., Yamaguchi, Y., 2004. Decomposing EEG data into space-time-frequency components using parallel factor analysis. *Neuroimage* 22 (3), 1035–1045.
- Moeller, F., Siebner, H.R., Wolff, S., Muhle, H., Boor, R., Granert, O., Jansen, O., Stephani, U., Siniatchkin, M., 2008. Changes in activity of striato-thalamo-cortical network precede generalized spike wave discharges. *Neuroimage* 39 (4), 1839–1849.
- Monti, M.M., 2011. Statistical analysis of fMRI time-series: a critical review of the GLM approach. *Front. Hum. Neurosci.* 5 (28).
- Moosmann, M., Schönfelder, V.H., Specht, K., Scheeringa, R., Nordby, H., Hugdahl, K., 2009. Realignment parameter-informed artefact correction for simultaneous EEG-fMRI recordings. *Neuroimage* 45 (4), 1144–1150.
- Mørup, M., Hansen, L.K., 2009. Automatic relevance determination for multi-way models. *J. Chemometr.: J. Chemometr. Soc.* 23 (7-8), 352–363.
- Mørup, M., Hansen, L.K., Herrmann, C.S., Parnas, J., Arnfred, S.M., 2006. Parallel factor analysis as an exploratory tool for wavelet transformed event-related EEG. *Neuroimage* 29 (3), 938–947.
- Murta, T., Leite, M., Carmichael, D.W., Figueiredo, P., Lemieux, L., 2015. Electrophysiological correlates of the BOLD signal for EEG-informed fMRI. *Hum. Brain Map.* 36 (1), 391–414.
- Nichols, T., Hayasaka, S., 2003. Controlling the familywise error rate in functional neuroimaging: a comparative review. *Stat. Methods Med. Res.* 12 (5), 419–446.
- Nichols, T.E., Holmes, A.P., 2002. Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum. Brain Map.* 15 (1), 1–25.
- Ostwald, D., Bagshaw, A.P., 2011. Information theoretic approaches to functional neuroimaging. *Magn. Reson. Imaging* 29 (10), 1417–1428.
- Papalexakis, E.E., 2016. Automatic unsupervised tensor mining with quality assessment. In: *Proceedings of the SIAM International Conference on Data Mining*. SIAM, pp. 711–719.
- Pedregosa, F., Eickenberg, M., Ciuciu, P., Thirion, B., Gramfort, A., 2015. Data-driven HRF estimation for encoding and decoding models. *Neuroimage* 104, 209–220.
- Pittau, F., LeVan, P., Moeller, F., Gholipour, T., Haegelen, C., Zermann, R., Dubeau, F., Gotman, J., 2011. Changes preceding interictal epileptic eeg abnormalities: comparison between EEG/fMRI and intracerebral EEG. *Epilepsia* 52 (6), 1120–1129.
- Poline, J.-B., Brett, M., 2012. The general linear model and fMRI: does love last forever? *Neuroimage* 62 (2), 871–880.
- Rigau, V., Morin, M., Rousset, M.-C., de Bock, F., Lebrun, A., Coubes, P., Picot, M.-C., Baldy-Moulinier, M., Bockaert, J., Crespel, A., et al., 2007. Angiogenesis is associated with blood-brain barrier permeability in temporal lobe epilepsy. *Brain* 130 (7), 1942–1956.
- Rosa, M.J., Kilner, J., Blankenburg, F., Josephs, O., Penny, W., 2010. Estimating the transfer function from neuronal activity to BOLD using simultaneous EEG-fMRI. *Neuroimage* 49 (2), 1496–1509.
- Salek-Haddadi, A., Diehl, B., Hamandi, K., Merschhemke, M., Liston, A., Friston, K., Duncan, J.S., Fish, D.R., Lemieux, L., 2006. Hemodynamic correlates of epileptiform discharges: an EEG-fMRI study of 63 patients with focal epilepsy. *Brain Res.* 1088 (1), 148–166.
- Salek-Haddadi, A., Friston, K., Lemieux, L., Fish, D., 2003. Studying spontaneous EEG activity with fMRI. *Brain Res. Rev.* 43 (1), 110–133.
- Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L., Smith, S.M., 2014. Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* 90, 449–468.
- Scheuer, M.L., Bagic, A., Wilson, S.B., 2017. Spike detection: Inter-reader agreement and a statistical Turing test on a large data set. *Clin. Neurophysiol.* 128 (1), 243–250.
- Schwartz, T.H., 2007. Neurovascular coupling and epilepsy: hemodynamic markers for localizing and predicting seizure onset. *Epilepsy Currents* 7 (4), 91–94.
- Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D., 2012. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex* 22 (1), 158–165.
- Sidiropoulos, N.D., De Lathauwer, L., Fu, X., Huang, K., Papalexakis, E.E., Faloutsos, C., 2017. Tensor decomposition for signal processing and machine learning. *IEEE Trans. Signal Process.* 65 (13), 3551–3582.
- Silverman, B., 1986. Density estimation for statistics and data analysis. *Monographs on Statistics and Applied Probability*. Chapman & Hall.
- Somers, B., Francart, T., Bertrand, A., 2018. A generic EEG artifact removal algorithm based on the multi-channel Wiener filter. *J. Neural Eng.* 15 (3), 036007.
- Sorber, L., Van Barel, M., De Lathauwer, L., 2015. Structured data fusion. *IEEE J. Sel. Top. Signal Process.* 9 (4), 586–600.
- Thornton, R., Laufs, H., Rodionov, R., Cannadathu, S., Carmichael, D.W., Vulliemoz, S., Salek-Haddadi, A., McEvoy, A.W., Smith, S.M., Lhatoo, S., et al., 2010. EEG correlated functional MRI and postoperative outcome in focal epilepsy. *J. Neurol. Neurosurg. Psychiatry* 81 (8), 922–927.
- Tousseyn, S., Dupont, P., Goffin, K., Sunaert, S., Van Paesschen, W., 2014. Sensitivity and specificity of interictal EEG-fMRI for detecting the ictal onset zone at different statistical thresholds. *Front. Neurol.* 5, 131.
- Tousseyn, S., Dupont, P., Goffin, K., Sunaert, S., Van Paesschen, W., 2015. Correspondence between large-scale ictal and interictal epileptic networks revealed by single photon emission computed tomography (SPECT) and electroencephalography (EEG)-functional magnetic resonance imaging (fMRI). *Epilepsia* 56 (3), 382–392.
- Tousseyn, S., Dupont, P., Robben, D., Goffin, K., Sunaert, S., Van Paesschen, W., 2014. A reliable and time-saving semiautomatic spike-template-based analysis of interictal EEG-fMRI. *Epilepsia* 55 (12), 2048–2058.
- Valdes-Sosa, P.A., Sanchez-Bornot, J.M., Sotero, R.C., Iturria-Medina, Y., Aleman-Gomez, Y., Bosch-Bayard, J., Carbonell, F., Ozaki, T., 2009. Model driven EEG/fMRI fusion of brain oscillations. *Hum. Brain Map.* 30 (9), 2701–2721.
- Van Eyndhoven, S., Hunyadi, B., De Lathauwer, L., Van Huffel, S., 2017. Flexible fusion of electroencephalography and functional magnetic resonance imaging: Revealing neural-hemodynamic coupling through structured matrix-tensor factorization. In: *Proceedings of the 25th European Signal Processing Conference (EUSIPCO)*. IEEE, pp. 26–30.
- Van Eyndhoven, S., Hunyadi, B., Dupont, P., Van Paesschen, W., Van Huffel, S., 2019. Semi-automated EEG enhancement improves localization of ictal onset zone with EEG-correlated fMRI. *Front. Neurol.* 10.
- Van Eyndhoven, S., Vervliet, N., De Lathauwer, L., Van Huffel, S., 2019. Identifying stable components of matrix/tensor factorizations via low-rank approximation of inter-factorization similarity. In: *Proceedings of the 27th European Signal Processing Conference (EUSIPCO)*, A Coruna, Spain, pp. 2–6.
- Vervliet, N., Debals, O., De Lathauwer, L., 2016. Tensorlab 3.0 Numerical optimization strategies for large-scale constrained and coupled matrix/tensor factorization. In: *Proceedings of the 50th Asilomar Conference on Signals, Systems and Computers*. IEEE, pp. 1733–1738.
- Vincent, T., Risser, L., Ciuciu, P., 2010. Spatially adaptive mixture modeling for analysis of fMRI time series. *IEEE Trans. Med. Imaging* 29 (4), 1059–1074.
- Vulliemoz, S., Thornton, R., Rodionov, R., Carmichael, D.W., Guye, M., Lhatoo, S., McEvoy, A.W., Spinelli, L., Michel, C.M., Duncan, J.S., et al., 2009. The spatio-temporal mapping of epileptic networks: combination of EEG-fMRI and EEG source imaging. *Neuroimage* 46 (3), 834–843.

- Waites, A.B., Shaw, M.E., Briellmann, R.S., Labate, A., Abbott, D.F., Jackson, G.D., 2005. How reliable are fMRI-EEG studies of epilepsy? A nonparametric approach to analysis validation and optimization. *Neuroimage* 24 (1), 192–199.
- Whitfield-Gabrieli, S., Nieto-Castanon, A., 2012. Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. *Brain Connect.* 2 (3), 125–141.
- Wilson, S.B., Turner, C.A., Emerson, R.G., Scheuer, M.L., 1999. Spike detection II: automatic, perception-based detection and clustering. *Clin. Neurophysiol.* 110 (3), 404–411.
- Woolrich, M.W., Behrens, T.E., Smith, S.M., 2004. Constrained linear basis sets for HRF modelling using Variational Bayes. *NeuroImage* 21 (4), 1748–1761.
- World Health Organization, 2019. Epilepsy. <https://www.who.int/en/news-room/factsheets/detail/epilepsy>. [Online; accessed 14/03/2019].
- Yuan, M., Lin, Y., 2006. Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc.: Series B (Stat. Methodol.)* 68 (1), 49–67.
- Zijlmans, M., Huiskamp, G., Hersevoort, M., Seppenwoolde, J.-H., van Huffelen, A.C., Leijten, F.S., 2007. EEG-fMRI in the preoperative work-up for epilepsy surgery. *Brain* 130 (9), 2343–2353.