



LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

INSTITUT FÜR STATISTIK



Stefanie Kalus, Philipp Sämann, Michael Czisch
& Ludwig Fahrmeir

fMRI activation detection with EEG priors

Technical Report Number 146, 2013
Department of Statistics
University of Munich

<http://www.stat.uni-muenchen.de>



fMRI activation detection with EEG priors

S. Kalus^{a,*}, P.G. Sämann^b, M. Czisch^b, L. Fahrmeir^a

^a*Department of Statistics, Ludwig-Maximilians-University, Ludwigstr. 33, 80539 Munich, Germany*

^b*Max Planck Institute of Psychiatry, Kraepelinstr. 2, 80804 Munich, Germany*

Abstract

The purpose of brain mapping techniques is to advance the understanding of the relationship between structure and function in the human brain in so-called activation studies. In this work, an advanced statistical model for combining functional magnetic resonance imaging (fMRI) and electroencephalography (EEG) recordings is developed to fuse complementary information about the location of neuronal activity. More precisely, a new Bayesian method is proposed for enhancing fMRI activation detection by the use of EEG-based spatial prior information in stimulus based experimental paradigms. I.e., we model and analyse stimulus influence by a spatial Bayesian variable selection scheme, and extend existing high-dimensional regression methods by incorporating prior information on binary selection indicators via a latent probit regression with either a spatially-varying or constant EEG effect. Spatially-varying effects are regularized by intrinsic Markov random field priors. Inference is based on a full Bayesian Markov Chain Monte Carlo (MCMC) approach. Whether the proposed algorithm is able to increase the sensitivity of mere fMRI models is examined in both a real-world application and a simulation study. We observed, that carefully selected EEG-prior information additionally increases sensitivity in activation regions that have been distorted by a low signal-to-noise ratio.

Keywords: EEG, fMRI, activation probabilities, Bayesian spatial modeling, Markov Chain Monte Carlo sampling, hierarchical model

*Corresponding author. Phone: +49 89 2180 6254

Email address: stefanie.kalus@stat.uni-muenchen.de (S. Kalus)

1. Introduction

Recording electrophysiological data simultaneously with functional magnetic resonance imaging (fMRI) targets at combining the advantages of both modalities to gain deeper insight into brain functioning. Recently, extensive experimental and methodological research has been done to achieve this goal (see Mulert and Lemieux, 2010, for a compelling review). Both measurement techniques capture neuronal activity. Scalp electroencephalography (EEG) measures the summed activity of postsynaptic currents. Hereby, time series of electrical activity are recorded from multiple electrodes placed on the scalp. fMRI is one of the most recently developed and at present most popular form of noninvasive imaging of human brain activity. Just recently, celebrating 20 years of fMRI, an entire special issue was released by the *NeuroImage* journal (Bandettini, 2012) demonstrating the popularity of the fMRI technique. fMRI also captures neuronal activity but in a more indirect way than EEG: Electrically active neurons are subject to a higher energy demand leading to an increase in blood flow into active brain regions. fMRI is based on the blood oxygenation level dependent (BOLD) effect and mirrors the changes in cerebral perfusion of capillaries with oxygenated blood. The fMRI technique provides a time series of three-dimensional images of the brain. Both modalities, EEG and fMRI, are based on brain processes related to neuronal activity. Therefore, activated brain regions can be located by analysing both fMRI and EEG data.

EEG and fMRI data differ substantially by their corresponding generative process and also capture different aspects of neuronal activity, but neither captures all. For being recorded at the scalp surface, electric events in deeper brain regions, which are detected by fMRI, emerge as damped oscillations in the EEG and are hard to localize. On the contrary, signals recorded by EEG need not necessarily be identifiable in the simultaneously recorded fMRI time series. For instance, no explicit BOLD response was identified for some short-lasting EEG signals with high amplitude, as e.g. K-complexes. See Ritter and Villringer (2006) and Daunizeau et al. (2010) for a more detailed discussion of situations where discrepancies between concurrent fMRI and EEG recordings emerge.

The main motivation for developing statistical methods for the fMRI-EEG data fusion lies in compensating shortcomings—related to technical issues as well as the biophysical generation process—of one technique by adding information from the other to gain a deeper insight into brain func-

tioning. A close coupling between measuring EEG and fMRI data seems indispensable for this. Therefore, in the last couple of years, recording EEG simultaneously with fMRI has become firmly established. For this to happen, several technical challenges had to be overcome (Ritter and Villringer, 2006).

Statistical methods that aim at combining EEG and fMRI face the challenge of integrating two data types that differ strongly in their structure as well as their generative process. Basically, so far, three different analysis approaches to multimodal integration have been proposed (Laufs et al., 2008; Daunizeau et al., 2010; Rosa et al., 2010). We distinguish between (i) EEG-to-fMRI approaches, where the fMRI signal is correlated with an EEG-defined event or feature to gain information on the location of electrophysiological phenomena, (ii) fMRI-to-EEG approaches, where the spatial information of the fMRI is used for a (spatiotemporal) source reconstruction of the EEG and (iii) symmetrical approaches, usually referring to the use of a common forward or a generative model that explains both EEG and fMRI data.

In this work, this set of approaches is extended by introducing a new type of method: an EEG-informed fMRI activation detection method. It is based on a high-dimensional fMRI regression model, which is also the basis for (i). In preparation for incorporating EEG information, a suitable uninformed fMRI activation detection method was developed. It grounds on ideas of Smith et al. (2003) and Smith and Fahrmeir (2007), who suggested to use a Bayesian variable selection approach and model averaging techniques to assess brain activity. Activity is, thereby, assessed locally at each voxel. A voxel is said to be active—in a narrower sense, to be responsive—if its fMRI signal trajectory responds to stimuli. That is, a voxel is active when a predictor component describing stimulus presentation has a significant effect on the fMRI signal. Alternatively, in terms of Smith et al. (2003), a voxel is said to be active if the inclusion of the stimulus predictor component plays a role in explaining the fMRI signal. To assess voxelwise activity, Smith et al. (2003) introduced voxelwise binary activation parameters in the form of selection indicators for the stimulus component and inspected the size of corresponding selection probability estimates, which serve as activation probability estimates. To account for the correlation between neighbouring voxels, estimation of activation probabilities was regularized by an Ising prior. The authors even provided a way to incorporate external prior information. This prior information, however, had to be available in the form of prior probability maps.

In this work, we extend the work of Smith et al. (2003) by providing means to include general prior activation information in continuous or binary form. For this, we replace the Ising prior with a latent, spatial probit regression stage that accomplishes the mapping between general variable values and probability estimations. For uninformed fMRI activation detection, the probit regression predictor consists of a spatially-varying intercept only (similar to the work of Smith and Smith, 2006). For EEG-informed fMRI activation detection schemes, spatial EEG information can be included as a further variable—possibly with a spatially-varying effect. To adjust for correlations within spatially-varying effect vectors, an intrinsic Gaussian Markov random field (IGMRF) prior was used.

Bayesian posterior analysis is based on a Markov Chain Monte Carlo (MCMC)-approach that allows to directly calculate all parameters and features of interest. In particular, the marginal posterior probability for activation at each voxel is of interest, because it can be thresholded to provide an activation map.

In comparison to previous methods, our approach has several significant advantages. It shares several advantages of the Ising model (see Smith et al., 2003): First, our Bayesian formulation allows the explicit modelling of the probability that a voxel is activated, which circumvents either the problematic interpretation of frequentist p -values in classical approaches or the counterintuitive selection of an activation threshold on the level of Bayesian activation amplitude effect estimates (Friston et al., 2002). Second, it incorporates spatial correlations at the level of activation probabilities, which are the parameters of interest, and not just indirectly on the level of activation effects as, for example, in Gössl et al. (2001), Penny et al. (2005) and Groves et al. (2009). Third, because the posterior distribution incorporates the spatial structure of an IGMRF prior, there is no need for spatial adjustments in a postprocessing step.

Above that, our approach possesses a major advantage over the Ising model. Our modelling formulation allows to incorporate very general forms of prior information, which can be used to enhance fMRI activation detection. The intended use of our approach is to incorporate EEG-prior activation information in the form of 3D source reconstruction maps (Michel et al., 2004), though any kind of external prior information can be used as long as it is available as a single 3D map with activation information. The scale of the contained measurements is quite arbitrary: Voxel values can be binary or continuous. We generally assume, however, that larger values indicate

activation. As a EEG-fMRI data fusion technique, our approach is related to (i) for being based on fMRI regression, but does not simply search for brain regions showing a correlation with EEG features. Our approach extends fMRI regression by adding spatial EEG information to activation detection.

In Kalus (2012) and Kalus et al. (2013), we have seen that our uninformed fMRI activation scheme has superior performance compared to classical SPM and the Ising model. They have highly increased sensitivity without losing their high specificity level and possess excellent edge-preserving properties. Hence, they can be considered to be an appropriate reference for examining the usefulness of EEG-enhanced schemes, which we propose in this work.

We will show that our EEG-enhanced fMRI activation schemes are approximately as sensitive as our uninformed algorithms. This hints at the usefulness of our enhanced approaches and is a prerequisite for their intended purpose. These procedures are developed to compensate a low signal-to-noise ratio in the fMRI signal, which is often a problem in event-related designs, to (a) make EEG phenomena visible otherwise not detected and (b) bring out activation regions more clearly when activation regions are blurred by noise. Though 3D maps derived by EEG source reconstruction methods have a low spatial resolution, they provide complementary information on the location of neuronal activity. To make use of EEG information without being biased by inaccurate location information, EEG-enhanced detection schemes need to be robust against prior misspecifications. Our developed algorithms ensure this for adapting to the level of EEG and fMRI congruency within the brain. This is achieved at the expense of a sensitivity increase. However, we identified data settings in which a substantial performance gain could be achieved compensating activation loss due to noise.

For researchers that want to use either the uninformed or EEG-enhanced fMRI activation scheme, we have implemented a user-friendly, computationally efficient software library making all discussed Bayesian algorithms available. The software is freely available as R package `Rfmrieeg`. Alternatively, C++ source code is provided for a binary `CfmrieegMain` program. Both kinds of software packages can be obtained from the author.

The rest of the paper is organised as follows. In the Methods section (Sect. 2), we first describe the basic fMRI general linear model used conventionally for activation detection (Sect. 2.1) and then describe how this model can be extended to incorporate external EEG information within a Bayesian hierarchical model by specifying suitable priors (Sect. 2.2). Bayesian model inference is outlined in Sect. 2.3. After providing the methodological back-

ground, we present results obtained on an event-related fMRI data set from an acoustic two-tone oddball design (Sect. 3). Simulation results derived from a modified version of the application dataset are shown in Sect. 4. In the Discussion section, we outline the main qualities and findings of our model and suggest starting points for further work.

2. Methods

2.1. The fMRI regression stage

fMRI data consist of signal time series $y_{i,t}, t = 1, \dots, T$, recorded at each voxel $i = 1, \dots, N$ of a three-dimensional brain image. Detection of brain activity is usually based on voxelwise regression models of the form

$$y_{i,t} = f_{base}(i, t) + f_{con}(i, t) + f_{stim}(i, t) + \epsilon_{i,t}. \quad (1)$$

In (1), $f_{base}(i, t)$ is the baseline trend, $f_{con}(i, t)$ is the effect of confounding covariates, $f_{stim}(i, t)$ is the effect of the (transformed) stimulus, and $\epsilon_{i,t}$ is the random error term at voxel i and time t . In the following, we discuss the components of the fMRI regression model (1) in detail.

The baseline term $f_{base}(i, t)$ corrects for slow periodic variations and drift either inherent to the scanning procedure or connected to non-paradigm correlated periodic variations. Thus, $f_{base}(i, t)$ serves as a highpass filter. In this work, it is chosen to consist of a discrete cosine transform (DCT) set (Friston et al., 2008, p. 123) as in SPM. In contrast to SPM, but conceptually equivalent, the highpass filter in our modelling approach enters the regression stage directly as a linear combination of basis functions $f_{base}(i, t) = \mathbf{x}^{(1)}(t)' \boldsymbol{\beta}_i^{(1)}$ with basis functions $\mathbf{x}^{(1)}(t) = (x_k^{(1)}(t), k = 1, \dots, p_1)'$ and weights $\boldsymbol{\beta}_i^{(1)} = (\beta_{i,k}^{(1)}, k = 1, \dots, p_1)'$.

The second term $f_{con}(i, t)$ accounts for further confounding effects like for example movement related artifacts or brain tissue-specific properties capturing effects from cardiac and respiratory cycles, which are not captured by the highpass filter. It is assumed, that according information is available in form of several univariate, global variables with value $\mathbf{x}^{(2)}(t) = (x_k^{(2)}(t), k = 1, \dots, p_2)'$, at time t . The corresponding voxelspecific effect vector $(\beta_{i,k}^{(2)}, k = 1, \dots, p_2)'$ is denoted as $\boldsymbol{\beta}_i^{(2)}$, so that $f_{con}(i, t) = \mathbf{x}^{(2)}(t)' \boldsymbol{\beta}_i^{(2)}$.

The regression component $f_{stim}(i, t)$ includes the transformed stimulus time series of a given type. The fMRI signal represents aggregated and time delayed neuronal activity, which in turn has a close correspondence

to the stimulus presentation. Therefore, the stimulus time series has to be transformed to the level of the fMRI response to obtain models that more closely resemble the observed fMRI signal thus yielding a better fit. In this work, the focus is on modelling event related stimuli, although only minor changes have to be made to incorporate block designs into the framework.

We follow an approach proposed by Josephs et al. (1997) using the concept of mathematical convolution:

$$f_{stim}(i, t) = \int_0^{\tau_{max}} hrf(i, \tau) u(t - \tau) d\tau$$

where the function $u(t)$ describes the given time course of neuronal activity and $hrf(i, \tau)$ is the unknown hemodynamic response function (HRF) at voxel i . To estimate $hrf(i, \tau)$, a flexible modelling strategy with basis functions B_k and corresponding voxelspecific weights $\beta_{i,k}^{(3)}$ is applied:

$$hrf(i, \tau) = \sum_{k=1}^{p_3} \beta_{i,k}^{(3)} B_k(\tau).$$

This approach leads to a flexible and data driven estimation of the voxelspecific functional form of the hemodynamic response. Different choices of basis sets exist (Henson et al., 2001). We focus on the canonical basis function set (Friston et al., 2008, pp. 181), which is the default choice in SPM. We use the functional form of the canonical HRF as written down in Kalus (2012), pp. 41.

The time series of neuronal activity $u(t)$ is set equal to the signal time series, which is modelled as follows: Suppose we have event related stimuli of one type at times $\tau_1, \tau_2, \dots, \tau_M$. A stimulus at time τ_m is modelled via a dirac delta function $\delta(t - \tau_m)$, i.e. a stick-function, so that $u(t) = \sum_{m=1}^M \delta(t - \tau_m)$.

With this modelling strategy, the stimulus predictor for all presented stimuli is linearized with respect to unknown HRF effects:

$$f_{stim}(i, t) = \sum_{k=1}^{p_3} \beta_{i,k}^{(3)} \underbrace{\sum_{m=1}^M B_k(t - \tau_m)}_{=x_k^{(3)}(t)} = \mathbf{x}^{(3)}(t)' \boldsymbol{\beta}_i^{(3)},$$

where $\mathbf{x}^{(3)}(t) = (x_k^{(3)}(t), k = 1, \dots, p_3)'$ and $\boldsymbol{\beta}_i^{(3)} = (\beta_{i,k}^{(3)}, k = 1, \dots, p_3)'$.

Linearization of additive regression components leads to a reformulation of model (1) into voxelwise linear models

$$\begin{aligned} y_{it} &= \mathbf{x}^{(1)}(t)' \boldsymbol{\beta}_i^{(1)} + \mathbf{x}^{(2)}(t)' \boldsymbol{\beta}_i^{(2)} + \mathbf{x}^{(3)}(t)' \boldsymbol{\beta}_i^{(3)} + \epsilon_{i,t} \\ &= \mathbf{x}(t)' \boldsymbol{\beta}_i + \epsilon_{i,t} \end{aligned} \quad (2)$$

where $\mathbf{x}(t)' = (\mathbf{x}^{(1)}(t)', \mathbf{x}^{(2)}(t)', \mathbf{x}^{(3)}(t)')$ and $\boldsymbol{\beta}_i' = (\boldsymbol{\beta}_i^{(1)'}, \boldsymbol{\beta}_i^{(2)'}, \boldsymbol{\beta}_i^{(3)'})$ are concatenated vectors. Collecting all observations and design vectors for voxel i in observation vectors and design matrices, we obtain a voxelwise linear model of the form

$$\mathbf{y}_i = \mathbf{X} \boldsymbol{\beta}_i + \boldsymbol{\epsilon}_i.$$

Our model depends on the assumption $\boldsymbol{\epsilon}_i \sim \mathcal{N}_T(\mathbf{0}, \sigma_i^2 \mathbf{I})$, where \mathbf{I} is the identity matrix of size T . So far, our model does not account for serial correlations, because in this paper we focus on the signal model. Additionally, note that the evaluation of posterior activation probabilities (see below) does not rely on the exact modelling of the error process traditionally needed for significance testing.

2.2. Incorporating EEG information via a spatial Bayesian variable selection

We follow Kalus et al. (2013)—which builds upon Smith and Fahrmeir (2007)—and define a spatial Bayesian activation detection model based on Bayesian variable selection in common linear regression (Smith and Kohn, 1996; George and McCulloch, 1997). For this, voxelwise binary indicator variables γ_i are introduced that are 1 if the stimulus regressor $f_{stim}(i, t)$ is selected, i.e. if there is a relationship between stimulus and signal time series, and 0 otherwise. This corresponds to

$$\gamma_i = \begin{cases} 1 & \text{if voxel } i \text{ is activated,} \\ 0 & \text{if voxel } i \text{ is not activated.} \end{cases}$$

For selecting f_{stim} when there is evidence for $\boldsymbol{\beta}_i^{(3)} \neq \mathbf{0}$, the scalar binary parameter γ_i controls the simultaneous selection of all p_3 basis functions for the stimulus component at voxel i . Selection of baseline and confounder components is not of interest, so that these components are always kept in the model. The vector of binary indicators $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_N)'$ represents an activation surface. Main interest lies in the estimation of voxelwise activation probabilities, $p(\gamma_i = 1), i = 1, \dots, N$. As in Kalus et al. (2013), suitable

spatial prior distributions, like a spatial random field prior, can be used to account for spatial dependencies. In this work, we modify and extend their approach to incorporate continuous spatial EEG information $J_i, i = 1, \dots, N$, into the estimation of activation probabilities. We assume that larger values of J_i are indicative for activation. For this, we use a prior motivated through a probit regression model. Prior activation probabilities are related to a latent predictor stage $\boldsymbol{\eta} = (\eta_1, \dots, \eta_N)'$ via a probit link. More exactly, we assume

$$p(\gamma_i = 1 | \eta_i) = \Phi(\eta_i)$$

where γ_i 's are conditionally independent given the value of $\eta_i, i = 1, \dots, N$. We consider additive predictors $\eta_i = \eta_{0,i} + \eta_{EEG,i}$, which consist of an intercept term and an flexible EEG effect. More precisely, we consider the following three predictor forms:

Predictor	$\eta_{0,i}$	+	$\eta_{EEG,i}$
0	$\alpha_{0,i}$		
glob	$\alpha_{0,i}$	+	$\alpha_G J_i$
flex	$\alpha_{0,i}$	+	$\alpha_i J_i$

Predictor 0 contains a spatially-varying intercept $\boldsymbol{\alpha}_0 = (\alpha_{0,i}, i = 1, \dots, N)'$ only. It is included to model regional differences in the basal response behaviour to experimental stimuli. For not incorporating EEG information and, hence, being a mere fMRI activation detection scheme, models with predictor 0 serve as a benchmark for model choices with EEG effect. Its performance was thoroughly evaluated in Kalus et al. (2013). Predictor glob additionally contains a global EEG effect α_G , which is an aggregated effect over all brain voxels. Its intended use is to uniformly increase the prior probability of activation across the brain proportionally to the local EEG value. That is, if a specific EEG value J is recorded for, e.g., a voxel A in the auditory cortex, and a voxel B in the motor cortex, then the prior influence is $\alpha_G J$ for both of these voxels and, thus, independent from the exact voxel locations (and only dependent on the value J). In contrast to this, predictor flex contains a spatially-varying EEG effect $\boldsymbol{\alpha} = (\alpha_i, i = 1, \dots, N)'$ instead of a global effect α_G . With this, the EEG influence for a specific value J may vary across the brain. E.g., for a voxel A in the auditory cortex, the influence may be $\alpha_A J$, whereas for a voxel B in the motor cortex, it may have a different strength $\alpha_B J$ if $\alpha_A \neq \alpha_B$. Hence, with this predictor type EEG influence may be downweighted if it contradicts fMRI information in one brain region,

but it is strong in congruent areas. Evaluation of a wide range of predictor forms (not shown) revealed that these three predictor types are superior to others in terms of performance, interpretability and algorithmic stability.

2.2.1. Priors for parameter α_0 in predictor θ

Spatial correlations and neighbourhood information for spatially-varying effects are incorporated through an IGMRF prior for α_0 . This spatial prior type imposes an estimable degree of smoothness reducing the number of effective parameters and enhancing clustering of activation regions. With an IGMRF, the prior of α_0 is of the following form:

$$p(\alpha_0|\xi_0^2) \propto (\xi_0^2)^{-(N-1)/2} \exp\left(-\frac{1}{2\xi_0^2}\alpha_0'\mathbf{Q}\alpha_0\right), \quad (3)$$

where \mathbf{Q} is the precision matrix of the IGMRF. Following the literature for first order IGMRFs on regular lattices (Rue and Held, 2005, Chap. 3), the precision matrix \mathbf{Q} is the Laplacian matrix of the underlying voxel grid, i.e. \mathbf{Q} has elements

$$Q_{ij} = \begin{cases} n_i, & i = j, \\ -1, & i \sim j, \\ 0, & \text{else,} \end{cases}$$

with n_i being the number of neighbours of voxel i and $i \sim j$ denoting i and j are neighbours. This choice of \mathbf{Q} has a rank deficiency of 1 with constant vectors lying in the null space of the matrix. This implies that departures from a hyperplane in α_0 are penalized, but not the absolute level.

The degree of smoothness is determined by the variance parameter ξ_0^2 . Complete smoothness, i.e. the case of a constant intercept $\alpha_{0,1} = \dots = \alpha_{0,N}$, is included as a limiting case as $\xi_0^2 \rightarrow 0$. The degree of smoothness is not set fixed in advance, but is estimated within the derived procedure. We make the common assumption of an inverse gamma prior

$$\xi_0^2 \sim IG(a_0, b_0), \quad (4)$$

where a_0 and b_0 are set to user choice.

Because constant vectors lie in the null space of the precision matrix \mathbf{Q} , implicitly an improper prior is used for the global level. We have seen some improvements in the convergency of parameter trajectories within the Markov chain Monte Carlo (MCMC) sampling procedures (see below) if

proper priors are used for every parameter in the proposed model. For this, we add a global intercept $\alpha_{0,G}$ and center $\boldsymbol{\alpha}_0$ to mean $\mathbf{0}$. The global intercept then is modelled to follow an informative univariate normal distribution

$$\alpha_{0,G} \sim \mathcal{N}(\mu_{0,G}, \xi_{0,G}^2), \quad (5)$$

whereas its variance parameter

$$\xi_{0,G}^2 \sim IG(a_{0,G}, b_{0,G}). \quad (6)$$

The mean $\mu_{0,G}$ is set fixed to user choice.

2.2.2. Priors for predictor glob-parameters $\boldsymbol{\alpha}_0$ and α_G

The prior for $\boldsymbol{\alpha}_0$ is chosen as in (3) and hyperparameters are modelled accordingly with (4), (5) and (6).

We constrain the global EEG coefficient α_G to be non-negative by using a log normal prior, i.e. $\alpha_G \sim \mathcal{LN}(\mu_G, \xi_G^2)$. For its variance parameter ξ_G^2 , we use an $IG(a_G, b_G)$ prior. The mean μ_G is set fixed to user choice.

2.2.3. Priors for predictor flex-parameters $\boldsymbol{\alpha}_0$ and $\boldsymbol{\alpha}$

Like before, the prior for $\boldsymbol{\alpha}_0$ is chosen as in (3) and its variance parameter has an inverse gamma prior as in (4).

The prior for $\boldsymbol{\alpha}$ is set likewise to

$$p(\boldsymbol{\alpha}|\xi^2) \propto (\xi^2)^{-(N-1)/2} \exp\left(-\frac{1}{2\xi^2}\boldsymbol{\alpha}'\mathbf{Q}\boldsymbol{\alpha}\right),$$

with $\xi^2 \sim IG(a, b)$. Again, we center $\boldsymbol{\alpha}$ to mean $\mathbf{0}$ and assume that the global intercept $\alpha_G \sim \mathcal{N}(\mu_G, \xi_G^2)$ and its variance parameter $\xi_G^2 \sim IG(a_G, b_G)$.

2.2.4. Prior distributions for common parameters

To complete the Bayesian model, priors are required for all the remaining parameters, which are common to all three predictor types. Given γ_i , let $\boldsymbol{\beta}_i(\gamma_i)$ be the vector of nonzero regression coefficients from regression i , i.e. $\boldsymbol{\beta}_i(\gamma_i = 1)' = (\boldsymbol{\beta}_i^{(1)'}, \boldsymbol{\beta}_i^{(2)'}, \boldsymbol{\beta}_i^{(3)'})$ and $\boldsymbol{\beta}_i(\gamma_i = 0)' = (\boldsymbol{\beta}_i^{(1)'}, \boldsymbol{\beta}_i^{(2)'})$. The matrix $\mathbf{X}(\gamma_i)$ denotes the corresponding design matrix. From the definition of the variable selection scheme, a prior is only required for the nonzero coefficients $\boldsymbol{\beta}_i(\gamma_i)$. This prior, however, has to be proper, because otherwise they will

be estimated as $\mathbf{0}$. Following Smith and Fahrmeir (2007), and references therein, we specify

$$\boldsymbol{\beta}_i(\gamma_i) | \mathbf{y}_i, \sigma_i^2, \gamma_i \sim \mathcal{N} \left(\hat{\boldsymbol{\beta}}_i(\gamma_i), T \sigma_i^2 (\mathbf{X}(\gamma_i)' \mathbf{X}(\gamma_i))^{-1} \right), \quad (7)$$

where $\hat{\boldsymbol{\beta}}_i(\gamma_i) = (\mathbf{X}(\gamma_i)' \mathbf{X}(\gamma_i))^{-1} \mathbf{X}(\gamma_i)' \mathbf{y}_i$.

For the variance parameter vector $\boldsymbol{\sigma}^2 = (\sigma_1^2, \dots, \sigma_N^2)'$, we assume standard independent noninformative priors. That is,

$$p(\boldsymbol{\sigma}^2) = \prod_{i=1}^N \frac{1}{\sigma_i^2}. \quad (8)$$

2.3. Posterior inference

Bayesian model estimation and inference is based on posterior quantities, i.e. quantities of the posterior multivariate probability distribution of all unknown parameters conditional on the data $\mathbf{y} = (\mathbf{y}'_1, \dots, \mathbf{y}'_N)'$ collected:

$$p(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\sigma}^2, \boldsymbol{\theta} | \mathbf{y}) \propto \prod_{i=1}^N [p(\mathbf{y}_i | \gamma_i, \boldsymbol{\beta}_i(\gamma_i), \sigma_i^2) p(\boldsymbol{\beta}_i(\gamma_i) | \gamma_i, \sigma_i^2, \mathbf{y}_i) p(\sigma_i^2)] p(\boldsymbol{\gamma} | \boldsymbol{\theta}) p(\boldsymbol{\theta}),$$

where $\boldsymbol{\beta} = (\boldsymbol{\beta}'_1, \dots, \boldsymbol{\beta}'_N)'$ and $\boldsymbol{\theta}$ contains all parameters in lower hierarchical levels—which depends on the exact predictor choice. For ease of exploration, the marginal posteriors, i.e. the probability distribution of each parameter(-subvector) given \mathbf{y} , are examined. In the case of fMRI analysis, the quantities of immediate interest are the marginal posterior probabilities of activation $p(\gamma_i = 1 | \mathbf{y})$ for all voxels i , $i = 1, \dots, N$, if γ_i is the indicator for selecting $f_{stim}(i, t)$. Closed form calculation of this quantity is not possible because it involves integration out of the binary variables that have support on 2^N possible indicator combinations—amongst other parameters of lower hierarchical levels of our model. Thus, inference relies on an adequate computational strategy that allows exploring features of the posterior distribution.

In this work, we base full Bayesian inference on a MCMC scheme (see Smith and Fahrmeir, 2007, and references therein). More precisely, we use a Gibbs Sampler (Gelfand and Smith, 1990), which involves repeated sampling of full conditional distributions of parameter(-subvectors) given the actual state of all other parameters in the model and the data \mathbf{y} . The generated Markov chain is designed to converge within a burn-in phase to the desired

marginal distributions, i.e. the equilibrium distributions of corresponding Markov chains.

In the following, the full conditional distributions of the different model components are discussed. We start with outlining a basic sampling scheme for estimating marginal posteriors $p(\gamma_i = 1|\mathbf{y})$ of the latent binary MRF $\boldsymbol{\gamma}$.

A convenient posterior sampling scheme can be derived by the following well-known data augmentation (Albert and Chib, 1993): We introduce N latent variables U_1, \dots, U_N with $U_i \stackrel{\text{ind.}}{\sim} \mathcal{N}(\eta_i, 1)$, where η_i is the value of the chosen predictor. Hence, $\mathbf{U} = (U_1, \dots, U_N) \sim \mathcal{N}_N(\boldsymbol{\eta}, \mathbf{I})$ with $\boldsymbol{\eta} = (\eta_1, \dots, \eta_N)'$. Define $\gamma_i = 1$ if $U_i > 0$ and $\gamma_i = 0$ otherwise. It can be shown that this goes along with the assumption of the γ_i 's being independent Bernoulli random variables with $p(\gamma_i = 1|\eta_i) = \Phi(\eta_i)$. From this definition, γ_i 's are conditionally independent given U_i .

In the following, $\boldsymbol{\gamma}_{j \neq i}(\mathbf{U}_{j \neq i})$ denotes $\boldsymbol{\gamma}(\mathbf{U})$ without the i -th component. As in Smith and Smith (2006), (U_i, γ_i) are generated as a pair to avoid a reducible sampling scheme by first generating from the marginalized distribution $p(\gamma_i|\mathbf{U}_{j \neq i}, \boldsymbol{\eta}, \boldsymbol{\beta}, \boldsymbol{\sigma}^2, \mathbf{y})$ and then generating from $p(U_i|\gamma_i, \mathbf{U}_{j \neq i}, \boldsymbol{\eta}, \boldsymbol{\beta}, \boldsymbol{\sigma}^2, \mathbf{y}) = p(U_i|\mathbf{U}_{j \neq i}, \boldsymbol{\eta}, \boldsymbol{\gamma}, \boldsymbol{\beta}, \boldsymbol{\sigma}^2, \mathbf{y})$. The latter coincides with the full conditional distribution for U_i , which can be derived as a truncated normal distribution with mean η_i and variance 1. If $\gamma_i = 1$, then U_i is constrained such that $U_i > 0$, and if $\gamma_i = 0$, U_i is constrained such that $U_i < 0$.

The update for γ_i is based on two marginalization steps to avoid reducibility: Marginalizations with respect to (M1) fMRI regression parameters and (M2) the voxel-specific auxiliary variable U_i .

For (M1), the proposed Bayesian variable selection scheme with priors (7) and (8) belongs to the family of conjugate hierarchical variable selection setups (George and McCulloch, 1997), for which the full conditional distribution of γ_i , $p(\gamma_i|\boldsymbol{\gamma}_{j \neq i}, \boldsymbol{\beta}, \boldsymbol{\sigma}^2, \mathbf{U}, \boldsymbol{\eta}, \mathbf{y})$, can analytically be marginalized with respect to $\boldsymbol{\beta}$ and $\boldsymbol{\sigma}^2$ to yield $p(\gamma_i|\boldsymbol{\gamma}_{j \neq i}, \mathbf{U}, \boldsymbol{\eta}, \mathbf{y})$ up to a proportionality constant. This can be used to derive an efficient and non-reducible computational routine in which a sample of γ_i does not depend on the actual values of $\boldsymbol{\beta}$ and $\boldsymbol{\sigma}^2$. In analogy to Smith and Kohn (1996) and Smith and Fahrmeir (2007), it can be shown that $p(\gamma_i|\boldsymbol{\gamma}_{j \neq i}, \mathbf{U}, \boldsymbol{\eta}, \mathbf{y}) \propto p(\mathbf{y}_i|\gamma_i)p(\gamma_i|\boldsymbol{\gamma}_{j \neq i}, \mathbf{U}, \boldsymbol{\eta})$, where

$$p(\mathbf{y}_i|\gamma_i) \propto S_i(\gamma_i)^{-T/2}(1+T)^{-q_i/2}.$$

Here, q_i is the number of nonzero regressors in regression i (i.e. either $p_1 + p_2$

if $\gamma_i = 0$ or $p = p_1 + p_2 + p_3$ if $\gamma_i = 1$), and

$$S_i(\gamma_i) = \mathbf{y}'_i \mathbf{y}_i - \mathbf{y}'_i \mathbf{X}(\gamma_i) (\mathbf{X}(\gamma_i)' \mathbf{X}(\gamma_i))^{-1} \mathbf{X}(\gamma_i)' \mathbf{y}_i \quad (9)$$

is the sum of squares in regression i corresponding to the value of the selection indicator γ_i .

Then, the marginalized conditional probability for activation can be derived as

$$p(\gamma_i = 1 | \boldsymbol{\gamma}_{j \neq i}, \mathbf{U}, \boldsymbol{\eta}, \mathbf{y}) = \frac{1}{1 + \mathcal{H}}, \quad \mathcal{H} = \exp(l_i) \frac{p(\gamma_i = 0 | \boldsymbol{\gamma}_{j \neq i}, \mathbf{U}, \boldsymbol{\eta})}{p(\gamma_i = 1 | \boldsymbol{\gamma}_{j \neq i}, \mathbf{U}, \boldsymbol{\eta})} \quad (10)$$

with

$$l_i = \log \left\{ \left(\frac{S_{i1}}{S_{i0}} \right)^{T/2} (1 + T)^{(p-p_1-p_2)/2} \right\}.$$

The conditional independence of the γ_i 's given \mathbf{U} can be used to note that $p(\gamma_i = 1 | \boldsymbol{\gamma}_{j \neq i}, \mathbf{U}, \boldsymbol{\eta}) = p(\gamma_i = 1 | U_i, \mathbf{U}_{j \neq i}, \boldsymbol{\eta})$. Marginalizing this quantity with respect to U_i (M2) yields $p(\gamma_i = 1 | \mathbf{U}_{j \neq i}, \boldsymbol{\eta}) = \Phi(\eta_i)$, which can be plugged in into (10) to yield a non-reducible update scheme.

For updating probit regression coefficients and related variance parameters, we propose the following sampling procedures. Derivation of update distributions, which define the specific algorithm, can be found in Kalus (2012). Here, we state the results. In the following, we denote the derived algorithms as iMRF_0 , $\text{iMRF}_{\text{glob}}$ and $\text{iMRF}_{\text{flex}}$ depending on the used predictor type.

2.3.1. iMRF_0

The full conditional of $\boldsymbol{\alpha}_0$ can be derived as

$$\boldsymbol{\alpha}_0 | \cdot \sim \mathcal{N} \left((\mathbf{I} + \boldsymbol{\Sigma}_0^-)^{-1} (\mathbf{U} - \boldsymbol{\eta}_{-\boldsymbol{\alpha}_0}), (\mathbf{I} + \boldsymbol{\Sigma}_0^-)^{-1} \right) \quad (11)$$

Thereby, $\boldsymbol{\Sigma}_0^-$ is the precision matrix of the multivariate normal spatial prior, i.e. $\boldsymbol{\Sigma}_0^- = \frac{1}{\xi_0^2} \mathbf{Q}$. In this equation, $\boldsymbol{\eta}_{-\boldsymbol{\alpha}_0}$ denotes the predictor value without the $\boldsymbol{\alpha}_0$ component, i.e. $\boldsymbol{\eta}_{-\boldsymbol{\alpha}_0} = \boldsymbol{\eta} - \mathbf{I} \boldsymbol{\alpha}_0$.

Hence, all components of $\boldsymbol{\alpha}_0$ are updated in one block to support mixing. The generated $\boldsymbol{\alpha}_0$ is centered to mean $\mathbf{0}$ immediately after sampling. An efficient sampling scheme from these multivariate normal distributions utilizes

the sparse matrix nature of the variance-covariance matrices at hand as described in Rue (2001). The procedure is based on a Cholesky decomposition of the precision matrix $(\mathbf{I} + \boldsymbol{\Sigma}_0^-)$ for the iGMRF update of $\boldsymbol{\alpha}_0$. Fill-reducing permutations yield an astonishing increase in speed. Lang and Brezger (2004) suggested to use the Cuthill-McKee algorithm to reduce the bandwidth of sparse matrices. We decided to calculate Cholesky decompositions efficiently using the routines of the C-library CHOLMOD (Davis and Hager, 1999). In an analyze step, CHOLMOD automatically determines the best reordering strategy. In our analyses, we have encountered that permutations found by CHOLMOD lead to smaller bandwidths than the Cuthill-McKee algorithm. For an in-depth discussion of available reordering strategies see the CHOLMOD user manual ¹.

For modelling the global level, iMRF₀ contains an global intercept $\alpha_{0,G}$. It can be updated by sampling from an univariate normal distribution with variance $(N + 1/\xi_{0,G}^2)^{-1}$ and mean $(N + 1/\xi_{0,G}^2)^{-1} (\mathbf{U} - \boldsymbol{\eta}_{-\alpha_{0,G}})' \mathbf{1} + \frac{\mu_{0,G}}{\xi_{0,G}^2}$.

Updates for variance parameters are based on the full conditional distributions $\xi_0^2 | \cdot \sim IG(a_0 + 0.5(N - 1), b_0 + 0.5\boldsymbol{\alpha}'_0 \mathbf{Q} \boldsymbol{\alpha}_0)$ and $\xi_{0,G}^2 | \cdot \sim IG(a_{0,G} + 0.5, b_{0,G} + 0.5(\alpha_{0,G} - \mu_{0,G})^2)$.

2.3.2. iMRF_{glob}

Additionally to the updates for iMRF₀, the global EEG coefficient α_G has to be updated for iMRF_{glob}.

For the positively restricted α_G coefficient, the full conditional distribution is of an unknown form, hence, it is updated via a Metropolis-Hastings step (see Gilks et al., 1996; Gelman et al., 2004), where the transition kernel of the corresponding Markov chain is based on the following: A proposal α_G is generated, but retained only with a specific acceptance probability. As proposal density, we take a distribution with support on positive α_G values. We decided to use a log-normal distribution with density

$$J(\alpha_{G,new} | \alpha_{G,old}) \propto \frac{1}{\alpha_{G,new}} \exp \left(-\frac{1}{2\xi_{prop,G}^2} (\log(\alpha_{G,new}) - \log(\alpha_{G,old}))^2 \right).$$

The corresponding acceptance probability—that is derived to yield a station-

¹<http://www.cise.ufl.edu/research/sparse/cholmod/>

ary Markov chain—is based on

$$\alpha(\alpha_{G,new}|\alpha_{G,old}) = \min \left\{ \frac{p(\boldsymbol{\gamma}, \boldsymbol{\theta}_{-\alpha_G}, \boldsymbol{\beta}, \boldsymbol{\sigma}^2, \alpha_{G,new}|\mathbf{y})/J(\alpha_{G,new}|\alpha_{G,old})}{p(\boldsymbol{\gamma}, \boldsymbol{\theta}_{-\alpha_G}, \boldsymbol{\beta}, \boldsymbol{\sigma}^2, \alpha_{G,old}|\mathbf{y})/J(\alpha_{G,old}|\alpha_{G,new})}, 1 \right\}.$$

The full conditional for the related variance parameter ξ_G^2 is an inverse gamma distribution with shape parameter a_G+1 and scale parameter $(\log(\alpha_G) - \mu_G)^2/2 + b_G$, which can be used for updating ξ_G^2 in a Gibbs sampling step.

2.3.3. iMRF_{flex}

The iMRF_{flex} algorithm consists of updating $\boldsymbol{\alpha}_0$ and $\boldsymbol{\alpha}$ and related hyperparameters. The spatially-varying intercept $\boldsymbol{\alpha}_0$ is again updated via (11). Analogously, the full conditional distribution of $\boldsymbol{\alpha}$ is derived as

$$\boldsymbol{\alpha}|\cdot \sim \mathcal{N}((\mathbf{J}'\mathbf{J} + \boldsymbol{\Sigma}^-)^{-1}\mathbf{J}'(\mathbf{U} - \boldsymbol{\eta}_{-\boldsymbol{\alpha}}), (\mathbf{J}'\mathbf{J} + \boldsymbol{\Sigma}^-)^{-1})$$

Thereby, $\boldsymbol{\Sigma}^-$ is the precision matrix of the multivariate normal spatial prior, i.e. $\boldsymbol{\Sigma}^- = \frac{1}{\xi^2}\mathbf{Q}$. Again, all components of $\boldsymbol{\alpha}$ are updated in one block to support mixing. The generated $\boldsymbol{\alpha}$ is centered to mean $\mathbf{0}$ immediately after sampling.

The global parameters that arise from centering $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}_0$, are updated in one block by sampling from a bivariate normal distribution, as described in Kalus (2012), p. 66.

Additional to the variance parameter updates listed for iMRF₀, the variance parameters for the local and global EEG effect have to be updated via $\xi^2|\cdot \sim IG(a+0.5(N-1), b+0.5\boldsymbol{\alpha}'\mathbf{Q}\boldsymbol{\alpha})$ and $\xi_G^2|\cdot \sim IG(a_G+0.5, b_G+0.5(\alpha_G - \mu_G)^2)$.

2.4. Monte Carlo Estimates

Posterior probability maps are obtained directly from the MCMC trajectories with sample index $l = 1, \dots, L$ as Rao-Blackwellized estimates

$$\hat{p}(\gamma_i = 1|\mathbf{y}) = \frac{1}{L} \sum_{l=1}^L p(\gamma_i = 1|\boldsymbol{\gamma}_{j \neq i}^{(l)}, \mathbf{U}^{(l)}, \boldsymbol{\eta}^{(l)}, \mathbf{y}), i = 1, \dots, N$$

of posterior probabilities $p(\gamma_i = 1|y)$ of activation. They provide visual evidence of brain regions with peak, high, low and practically no activation. Moreover, they are the basis for binary activation maps, in which voxels are

classified as active if posterior activation probabilities exceed some threshold, see the next section.

Estimates for regression stage parameters can be derived as in Smith et al. (2003). Marginalization yields $E(\boldsymbol{\beta}_i|\gamma_i, \mathbf{y}_i) = \hat{\boldsymbol{\beta}}_i(\gamma_i)$ and $E(\sigma_i^2|\gamma_i, \mathbf{y}_i) = S_i(\gamma_i)/(T - 2)$. With this, model average estimates can be based on

$$\begin{aligned} E(\boldsymbol{\beta}_i|\mathbf{y}) &\approx \hat{\boldsymbol{\beta}}_i(\gamma_i = 1)\hat{p}(\gamma_i = 1|\mathbf{y}) + \hat{\boldsymbol{\beta}}_i(\gamma_i = 0)\hat{p}(\gamma_i = 0|\mathbf{y}) \\ E(\sigma_i^2|\mathbf{y}) &\approx \frac{S_i(\gamma_i = 1)}{T - 2}\hat{p}(\gamma_i = 1|\mathbf{y}) + \frac{S_i(\gamma_i = 0)}{T - 2}\hat{p}(\gamma_i = 0|\mathbf{y}). \end{aligned}$$

These two estimates can be evaluated after the last iteration of the MCMC algorithm.

2.5. Activation Classification

We construct a binary activation map $\hat{\boldsymbol{\gamma}} = (\hat{\gamma}_i, i = 1, \dots, N)$ from a posterior probability map by defining voxel i as active (i.e. $\hat{\gamma}_i = 1$) if $\hat{p}(\gamma_i = 1|\mathbf{y}) > t$ for some threshold t . A general strategy for choosing t is to consider false positive counts $\text{FD} = \sum_{i=1}^N (1 - \gamma_i)d_i$, false negative counts $\text{FN} = \sum_{i=1}^N \gamma_i(1 - d_i)$, or rates, where d_i denotes the binary decision to classify a voxel as active ($d_i = 1$) or not ($d_i = 0$), and define loss functions as linear combinations of related posterior expectations, e.g.

$$c \text{E}(\text{FD} | y) + \text{E}(\text{FN} | y). \quad (12)$$

Müller et al. (2004) show that all optimal decisions for loss functions of this type are of the form $d_i = I(p(\gamma_i = 1|\mathbf{y}) > t)$. In particular, they show that for $c = 1$ the optimal threshold for (12) is $t = 0.5$. In this work, we use the threshold $t = 0.8722$ of Smith and Fahrmeir (2007), which corresponds to $c = 6.87$ and, hence, puts more weight on controlling the specificity of the algorithm.

3. Application to data from an acoustic oddball design

In this section, we present results obtained from an event-related fMRI data set.

3.1. Material and configurations

3.1.1. Event-related fMRI data

The fMRI data were acquired on a clinical 3-Tesla scanner (General Electric MR750) from a healthy male volunteer during an active two-tone-oddball paradigm (Kiehl et al., 2005). In this paradigm, rare (high-pitched) odd tones (1 500 Hz, duration 50 ms) appeared with 10% probability against the background of frequent (low-pitched) tones (1 000 Hz, duration 50 ms). The interstimulus interval (ISI) was set to an average of 1 000 ms. The subject was instructed to continuously pay attention to the tones and press the response button immediately after recognizing an odd (high-pitched) tone. Whole brain fMRI time series were acquired using an echoplanar imaging (EPI) sequence (time of repetition [TR] 2 000 ms, time of echo [TE] 40 ms, slice orientation according to anterior-commissure/posterior-commissure landmarks, 28 slices, slice thickness 3.5 mm, 0.5 mm gap, in-plane resolution $3.125 \times 3.125 \text{ mm}^2$) while the acoustic oddball paradigm was applied. A total of 307 image volumes was recorded over 10 minutes.

The dataset was postprocessed using the SPM software ². First, data were corrected for slice time differences to compensate for different acquisition times due to the interleaved slice acquisition scheme in each volume. Second, images were motion corrected by rigid-body coregistration with the mean image of the uncorrected time series. Third, images were spatially normalized using linear and non-linear transformations to an EPI wholehead template in standard MNI (Montreal Neurological Institute) space with default settings of the SPM8 distribution. Intrinsic to the spatial normalization step is an interpolation step that was set to gain voxels sized $4 \times 4 \times 4 \text{ mm}^3$. Additionally, the image trajectories were smoothed using a symmetric 3D-Gaussian kernel with $8 \times 8 \times 8 \text{ mm}^3$ FWHM (full width half maximum).

3.1.2. Event-related EEG data

For generating the EEG information map, a group of 9 subjects underwent the experimental procedure described above. EEG recording was parallel and synchronized to fMRI measurements. An Easy Cap (<http://www.easycap.de/easycap>) with 64 EEG electrodes was used, including an electrocardiogram electrode placed on the left side of the participants' back close to their spinal cord, referenced against the fronto-central electrode

²<http://www.fil.ion.ucl.ac.uk/spm>, version SPM8

FCz. Data were continuously sampled throughout the experiment at 5 kHz. To allow for optimal artifact correction, EEG recordings and fMRI were synchronized using the scanner’s 10 MHz master clock (Mandelkow et al., 2006). Electrode impedance was below 5 k Ω . Additionally, trigger pulses from the MRI system were recorded for subsequent off-line MRI artifact correction.

EEG data was corrected for gradient induced and cardiobalistic artifacts (Czisch et al., 2009) using Vision Analyzer 1.05 (Brain Products). That is, after MRI artifact correction an independent component analysis (ICA) (Beckmann and Smith, 2005) was performed to clean data from cardiobalistic artifacts. Afterwards data was bandpass-filtered with typical settings for evoked potential analysis (low cutoff 0.5 Hz, high cutoff 30 Hz). After baseline correction and DC detrending, EEG time series were segmented into 1 000 ms segments according to odd and even tone onset time points (-200 ms to +800 ms), which after exclusion of obviously corrupted data segments resulted in an average of \sim 490 segments for even tones and \sim 45 segments for odd tones for each electrode (time resolution 250 Hz). For further use, segments of each stimulus type are averaged to gain average event-related potential (ERP) trajectories. These are assumed to contain an estimate of the amplitude and morphology of the electrophysiological response. For the odd minus even contrast being of special interest in oddball studies, we calculated a differential ERP by taking the pointwise differences of odd and even ERPs.

Based on the derived differential ERP, spatial EEG maps were then calculated for each time point using the sLORETA software (<http://www.uzh.ch/keyinst/loreta.htm>). Then a 3D prior information map in form of a group and time aggregated sLORETA map was derived. For this, a spatial Tensor-PICA component (Beckmann and Smith, 2005) was selected that showed a spatial pattern and time course associated with stimulus presentation. From its spatial pattern, it was supposed to add information to fMRI activation detection especially in upper central parts of the brain. Because we do not differentiate between activation and deactivation (with regard to the sign of effects) we took the absolute value of all map values.

3.1.3. Analysis configurations

For each run, we used 6 000 MCMC iterations including a burnin phase of 1 000 iterations. To remove strong autocorrelations, we thinned out resulting parameter trajectories with a stepping of 5. Hyperparameter values were chosen in such a way that convergency to the equilibrium distribu-

tion was acceptable within burnin. If contained in the model, ξ_0^2 and ξ^2 follow $IG(902, 4054.5)$, $\xi_{0,G}^2$ and ξ_G^2 follow $IG(3, 1)$. Corresponding global prior means (μ_G resp. $\mu_{0,G}$) were set to 0. Convergency was assessed by visual inspection of trace plots and corresponding diagnostic tools. Note that variance parameters priors can be considered to be informative. Robustness studies (not shown here) revealed, however, that activation estimates of interest were robust against changes in hyperparameter specifications. The neighbourhoods underlying both random field priors incorporated the next 6 direct 3D voxel neighbours. We restricted analysis to all voxels lying within a contiguous brain mask defined upon a global mean threshold in analogy to the selection procedure in SPM (cf. the SPM8 user manual). That is, a voxel is selected if all its time series values are large enough to be brain-tissue specific.

3.2. Results

To evaluate whether the proposed models are able to increase sensitivity by using EEG information, we compare our EEG-enhanced activation detection algorithms iMRF_{glob} and iMRF_{flex} with the corresponding uninformed algorithm iMRF_0 . The figures used in this section to display results comprise the following contents:

- $\log(\text{LR}) = (\log(\text{LR}_i))$ is the log-likelihood ratio statistic map with $\log(\text{LR}_i) = T \log(S_{i0}/S_{i1})$, which is a measure for the amount of activation information contained in the fMRI data. It is directly available from our algorithm via (9) and is an alternative to the SPM F -map.
- $\mathbf{J} = (J_i)$ is the EEG-based prior information map.
- $\hat{\boldsymbol{\eta}}_{EEG} = (\hat{\alpha}_i J_i)$ resp. $\hat{\boldsymbol{\eta}}_{EEG} = (\hat{\alpha}_G J_i)$ is the estimated EEG contribution to the latent probit predictor. It visualizes the strength and location of mere EEG effects.
- $\hat{\Phi}_{\text{diff}} = \left(\Phi(\hat{\eta}_{i,\text{iMRF}_{glob/flex}}; 0, 1) - \Phi(\hat{\eta}_{i,\text{iMRF}_0}; 0, 1) \right)$ is the marginal prior probability difference map. To enhance the understanding of these, note the following: It holds that $\hat{p}(\gamma_i = 1 | \mathbf{U}_{j \neq i}, \boldsymbol{\theta}_{-\mathbf{U}}) = \Phi(\hat{\eta}_i; 0, 1)$ for each voxel i . Being based on the posterior estimate $\hat{\eta}_i$ of $\eta_i = \eta_0 + \eta_{EEG}$, this quantity can be interpreted as an adaptive prior probability for activation—taking into account the latent effects arising from the fMRI and EEG activation information, and the congruency of both. In

the figures, the difference maps visualize how the adaptive prior map changes when EEG information is incorporated.

- $\hat{\boldsymbol{\pi}}_{\text{diff}} = (\hat{p}(\gamma_i = 1 | \boldsymbol{\theta}, \mathbf{y}, \text{iMRF}_{glob/flex}) - \hat{p}(\gamma_i = 1 | \boldsymbol{\theta}, \mathbf{y}, \text{iMRF}_0))$ is the posterior activation probability difference map. Posterior activation probability maps are thresholded (cf. Sect. 2.5) to gain activation maps and, hence, visualize the EEG impact on a continuous level.
- $\hat{\boldsymbol{\gamma}}_{\text{diff}} = (\hat{\gamma}_{i, \text{iMRF}_{glob/flex}} - \hat{\gamma}_{i, \text{iMRF}_0})$ is the categorical activation difference map between iMRF_{glob} resp. iMRF_{flex} and iMRF_0 . The impact of incorporating EEG information can most easily be seen with this kind of map, where yellow indicates voxel found to be activated with both algorithms, red indicates voxels only found by the EEG enhancement and blue indicates voxels only found by the uninformed reference model.

In Fig. 1, the results of iMRF_{flex} are depicted. We select result maps from brain slices 21, 23, 25 and 27, because these slices comprise the parts of the brain where the EEG information is strong (compare \mathbf{J} maps) and is supposed to especially add value to the activation detection (compare low values in $\log(\text{LR})$ maps).

For $\hat{\boldsymbol{\eta}}_{EEG}$, each J_i value is multiplied with a spatially-varying effect α_i to form $\hat{\eta}_{EEG,i}$. We see, that $\hat{\boldsymbol{\eta}}_{EEG}$ is zero (medium grey areas), if $J_i = 0$. In other regions (with $J_i > 0$), deviations to negative (dark-grey) or positive (light-grey) values arise, because the spatially-varying effect α_i is non-restricted. If $\hat{\alpha}_i < 0$, then the EEG effect speaks for non-activation, if $\hat{\alpha}_i > 0$, otherwise. It appears that the α_i -values depend on the grade of congruency between EEG (\mathbf{J}) and fMRI ($\log(\text{LR})$) information. The influence of $\hat{\boldsymbol{\eta}}_{EEG}$ can be compensated by the spatially-varying intercept. Both terms sum up to form the whole predictor value $\hat{\boldsymbol{\eta}}$. Structural effects on the whole predictor level are examined via $\hat{\boldsymbol{\Phi}}_{\text{diff}}$, which we discuss next.

Marginal prior probability difference maps $\hat{\boldsymbol{\Phi}}_{\text{diff}}$ can be used to assess the latent predictor estimation. Here, we observe changes in predictor estimation mainly in parts of the brain where fMRI activation is present (non-white areas). This hints at the stability of the algorithm, because predictor values are not changed in arbitrary parts of the brain. The marginal prior activation probability is especially increased at activation region boundaries, e.g. at the central activation focus in slices 25 and 27 (red-scale areas). In adjacent regions, the marginal prior probability of voxels is decreased (blue-scale areas),

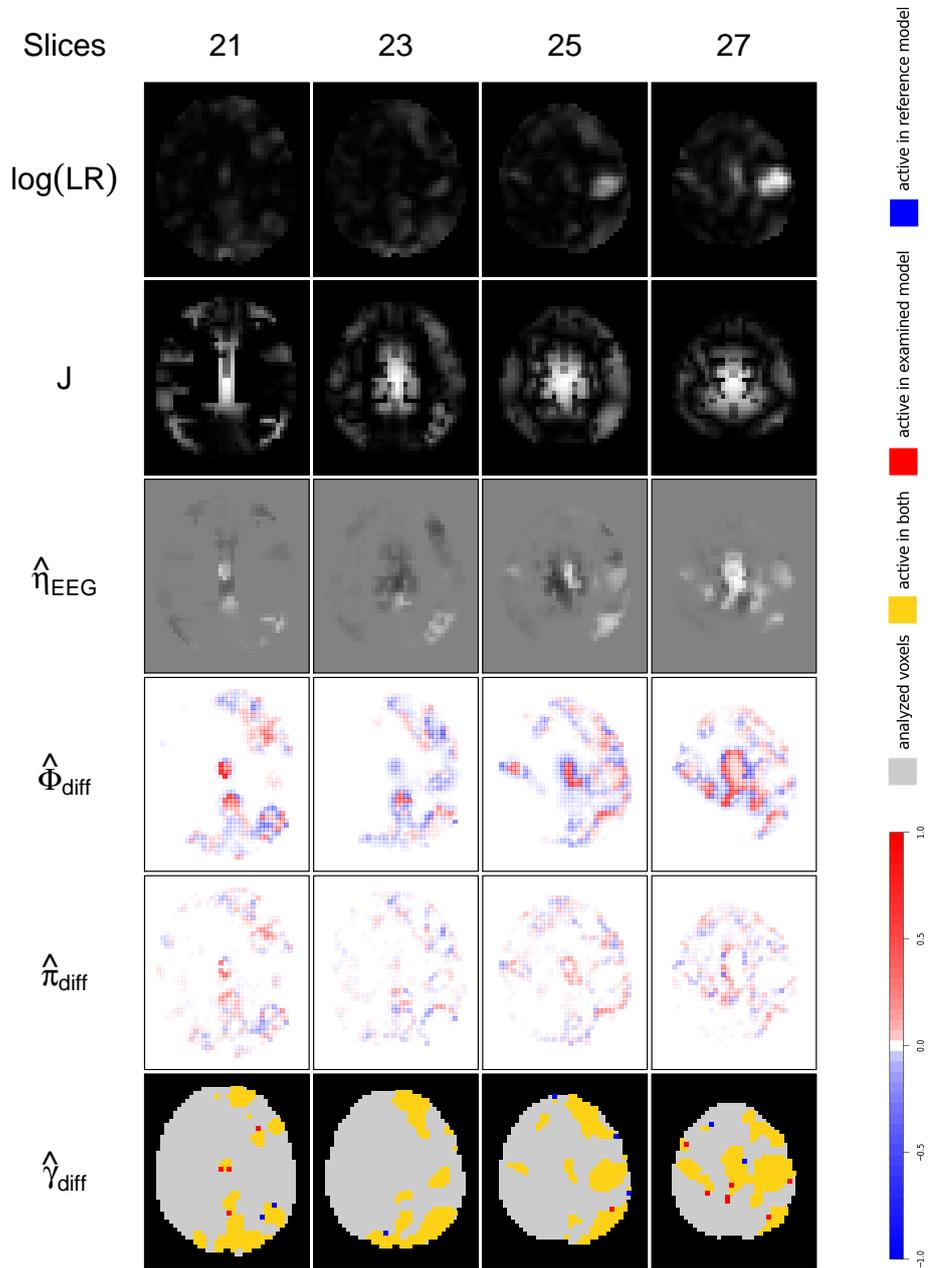


Figure 1: Summary of results from the iMR_{flex} model in comparison to iMR_0 (selected brain slices: 21, 23, 25, 27).

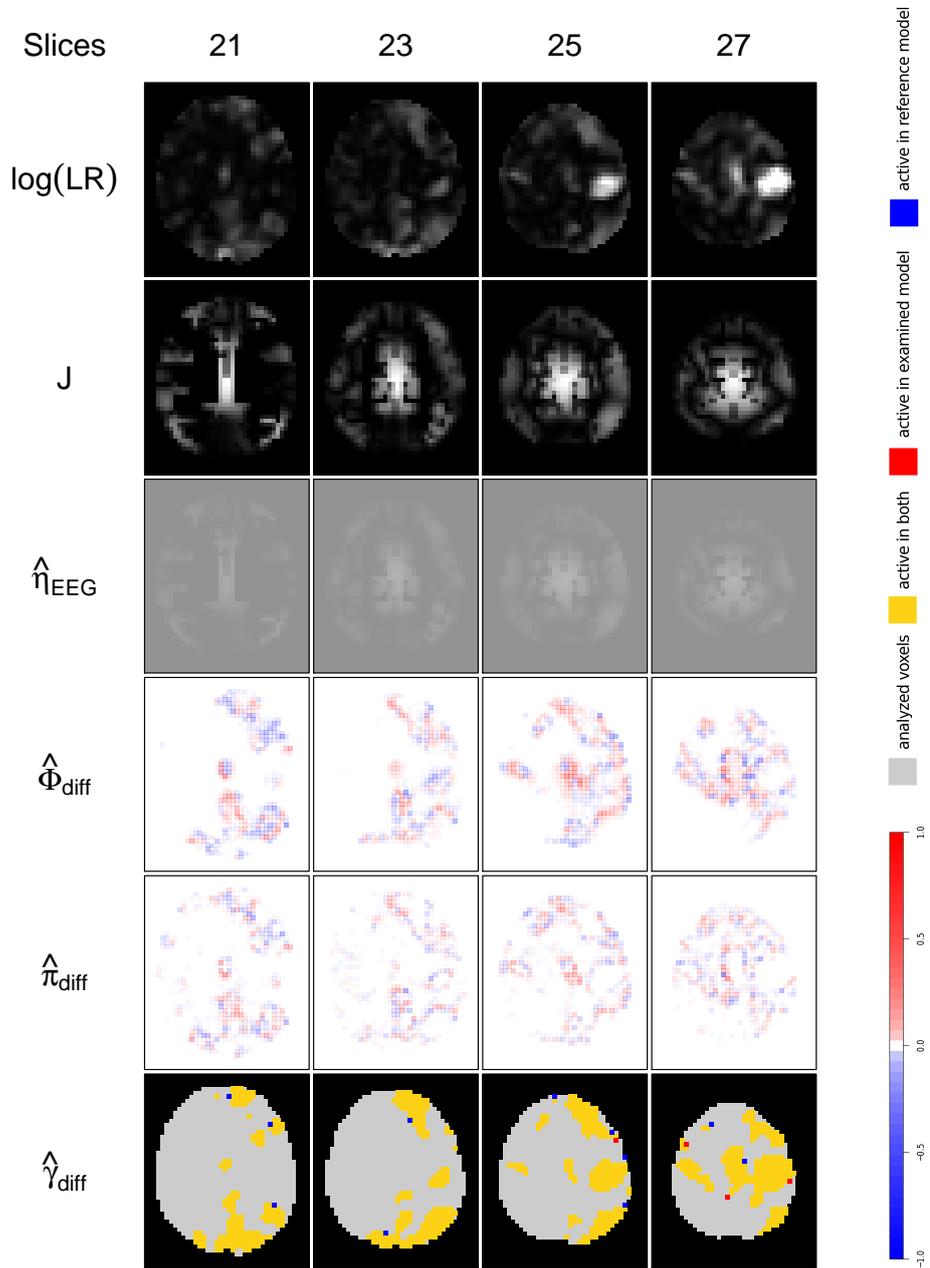


Figure 2: Summary of results from the $\text{iMRF}_{\text{glob}}$ model in comparison to iMRF_0 (selected brain slices: 21, 23, 25, 27).

which indicates that EEG contradicts fMRI activation information there. That is, the EEG activation focus is not congruent to the fMRI activation focus: The EEG peak covers the fMRI peak, but it exceeds the fMRI peak by far. That iMRF_{flex} detects these edges can be regarded as a feature of the algorithm, because it indicates robustness against non-congruency of the EEG prior. In summary, on the one hand, when EEG information contradicts fMRI information, the EEG contribution seems to be either downweighted or to be used as information on deactivation. On the other hand, when EEG information is congruent to fMRI information, the EEG contribution is increased. Regarding the sharp boundaries between red- and blue-scale areas, the algorithm appears to possess good edge-preserving properties—which, however, may prevent EEG-enhancement (see Discussion).

So far, we see that marginal prior probability maps $\hat{\Phi}_{\text{diff}}$ possess a structurally sensible form of EEG dependency. Whether this is carried through to activation classifications can be deferred from posterior activation probability maps $\hat{\pi}_{\text{diff}}$. The iMRF_{flex} algorithm locates largest posterior probability differences at activation region boundaries. Posterior activation probabilities are either increased or decreased. Particularly, with supporting EEG information, i.e. in central parts of the brain, we observe an increase in posterior activation probabilities.

In $\hat{\gamma}_{\text{diff}}$, we see that changes in posterior activation probabilities are large enough to slightly increase region sizes in targeted areas (e.g. in the central part in layer 27). However, a few voxels also become non-active in regions where algorithms emphasize sharp edges between activation and non-activation zones.

In Fig. 2, the results of the iMRF_{glob} analysis can be found. For the $\hat{\eta}_{EEG}$, each J_i value is multiplied with a global effect α_G to form $\hat{\eta}_{EEG,i}$. We see again, that $\hat{\eta}_{EEG}$ is zero (medium grey areas), if $J_i = 0$. In other regions (with $J_i > 0$), positive values arise (light-grey areas), because the global effect α_G is restricted to non-zero values. Note that the colour scale of the $\hat{\eta}_{EEG}$ -maps are chosen with respect to the range of the whole predictor. Hence, if the EEG component is not sufficiently strong, the contrast of the image is low. This happens here, because a global EEG effect is calculated that poorly stands out from the intercept component. Nevertheless, as it can be seen from $\hat{\Phi}_{\text{diff}}$ and $\hat{\pi}_{\text{diff}}$ maps, it is strong enough to affect posterior activation probabilities. However, due to its weakness it is not able to substantially increase sensitivity in this particular application. As soon as EEG contradicts fMRI in larger parts of the brain, the global EEG effect is estimated to be so small that it

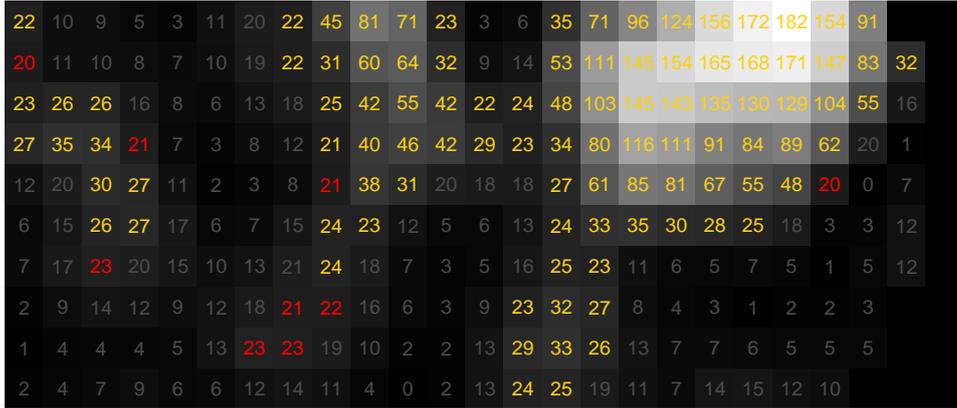
hardly influences activation probability and, hence, activation classification ($\hat{\gamma}_{\text{diff}}$).

Thus, we have seen that the global and local degree of congruency has a strong effect on the performance of the EEG enhancement. To elaborate on this, in Fig. 3a, we visualized the overlap of EEG and fMRI information of a region of interest (center of slice 27). As before, voxelwise fMRI activation information is provided through $\log(\text{LR})$ statistics, whereas for EEG, the corresponding \mathbf{J} -values are plotted. Additionally in Fig. 3b, the $\log(\text{LR})$ values are plotted, with the values being colour-coded as in the activation difference maps—displaying the activation differences between iMRF_{flex} and iMRF_{glob} . From these plots, we see that the fMRI peaks are well-pronounced and $\log(\text{LR})$ values decrease to very small values in in-between activation regions. Only at region borders, there are some voxels with medium sized $\log(\text{LR})$ values indicating that some amount of activation is contained in corresponding time series. Hence, the number of voxels, where the EEG information is supposed to enhance activation information, is limited. Additionally, from Fig. 3b, we see that the overlap between EEG and fMRI information is not sufficient to increase sensitivity of activation detection to a large extent. For example, a small cluster of four voxels can be forced to be found active, but not surrounding voxels with nearly as high $\log(\text{LR})$ -values, because the EEG is not sufficiently strong in this region.

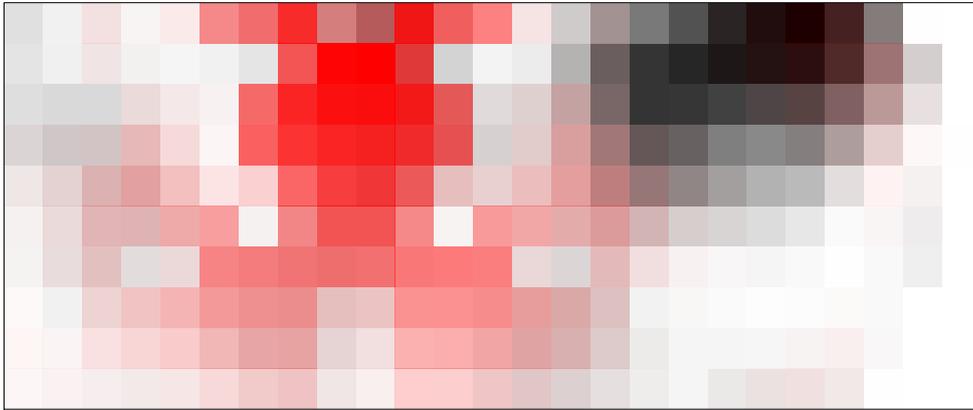
4. Simulation study

Although, iMRF_{flex} and iMRF_{glob} models did only achieve a minor gain in sensitivity in our particular application, analysis results gave rise for some assumptions about suitable data structures, where EEG-enhanced fMRI activation detection algorithms are expected to lead to a substantial gain in sensitivity. Assume that we have a ROI where EEG information should enhance sensitivity in fMRI activation detection. We have, for example, seen that the EEG prior and fMRI data must be fairly congruent in this ROI and that the fMRI time series of corresponding voxels must contain a sufficient amount of activation information. Moreover, the ROI is preferably not just a border, but a de-noised peak structure. To yield some support for these assumptions, we set up the following simulation study based on modifications of the original dataset analysed in the previous section.

To decrease computation time, we took a subset of the data (with 3D coordinates $x \in [15, 38]$, $y \in [14, 33]$ and $z \in [22, 33]$) incorporating the

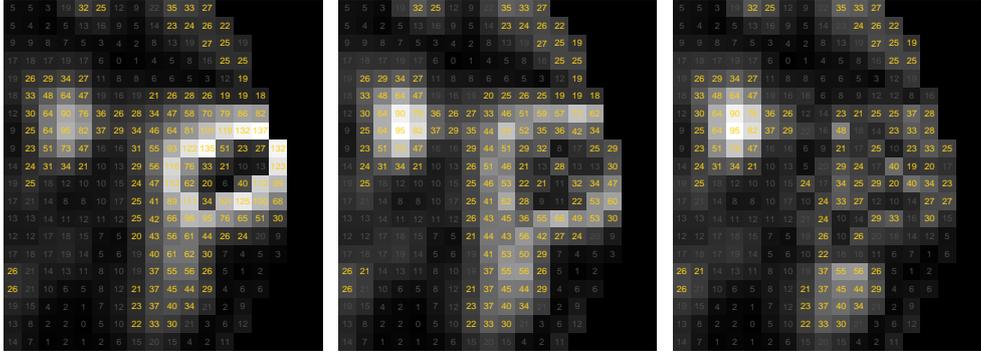


(a) Map with log(LR)-statistic information (grey-scale and numbers). Font colour indicates activation classification differences between $iMRF_{flex}$ and $iMRF_0$.



(b) Overlap of EEG (red-scale) and fMRI (grey-scale) information. Here: the higher the colour saturation, the higher activation information.

Figure 3: Informational content of combined EEG and fMRI dataset in the center of slice 27.



(a) Data scenario 1 (b) Data scenario 2 (c) Data scenario 3

Figure 4: Simulation scenarios: Activation information contained in fMRI data.

(neurologically) right part of Brodmann areas 1-4, 6 (containing the primary somatosensory, primary motor and premotor cortex), which showed a strong fMRI signal in prior analyses. Since the original data structure possessed well pronounced, smooth activation peaks, we decided to de-noise the observed peak structure. For this, we used three different strategies (cf. Fig. 4). Data scenario 1 was generated to contain a hole within the activation structure with surrounding strongly activated voxels. Data scenario 2 was generated to contain a slightly larger hole, but surrounding voxels were forced to obtain much smaller $\log(\text{LR})$ values than in scenario 1. In data scenario 3, the whole activation region was distorted by noise. As artificial prior EEG information, we used two congruent information maps in form of the binary extended fMRI activation map (denoted as J_1 , cf. Fig. 5a) and in form of the continuous $\log(\text{LR})$ -map (denoted as J_2 , cf. Fig. 5b) from the original dataset without additional noise. As a form of local prior information, we additionally used the $\log(\text{LR})$ map that was set to zero outside a rectangular region of 8×9 voxels covering the targeted activation focus across all slices (denoted as J_3 , cf. Fig. 5c). This prior introduced non-congruency in large parts of the analysed brain, whereupon non-congruent regions adjoined the targeted activation focus with congruent information at some borders.

In Fig. 6, the results in form of activation difference maps in reference to iMRF_0 are given. Each two-column block depicts the results from the three different data scenarios, whereas results from using the three different EEG

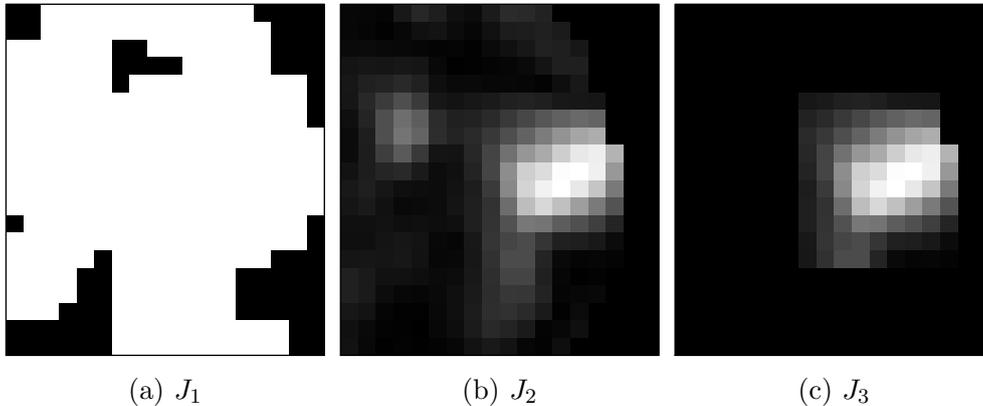


Figure 5: Simulation scenarios: Different types of EEG prior information based on the non-denoised data. Prior J_1 is the underlying activation map from an initial run. Prior J_2 is a congruent, continuous prior in form of the underlying log(LR)-map. Prior J_3 is a non-congruent, continuous prior in form of a cut log(LR)-map.

prior types are presented in each row.

With J_1 , some voxels at the borders of the activation region change their state, but no changes occur in the targeted region inside the activation focus. Particularly, no sensitivity increase is achieved. Hence, we conclude that if the EEG-prior information is not strong enough, EEG-enhanced activation schemes cannot make use of it.

With J_2 , iMRF_{flex} and iMRF_{glob} perform very well compensating the lost activation in any case. Note that the iMRF algorithms incline to smooth across region edges enlarging peak regions as well.

With J_3 , iMRF_{glob} performs well in data scenarios 1 and 2, but it cannot compensate the loss in data scenario 3—indicating that if the activation structure is too de-noised, non-congruent EEG information cannot compensate the loss. This might be traced back to the fact that in the non-congruent setting at hand, there is not enough information in the data to achieve a sufficiently large global EEG effect estimate. In contrast to this, iMRF_{flex} indeed compensate the loss in data scenarios 1 and 2, but outside the activation focus, all voxels become non-activated. This might indicate that iMRF_{flex} uses EEG prior information also on non-activation. If peak structures are too de-noised (data scenario 3), iMRF_{flex} cannot fill up activation gaps and even

deteriorates the activation structure.

5. Discussion

In this work, we proposed a novel strategy for the fMRI-EEG data fusion—unequal to any fusion approach found in the literature. We extended the Bayesian fMRI activation detection approach from Smith et al. (2003) to incorporate spatial EEG information in form of 3D prior activation information to enhance sensitivity. More precisely, a high-dimensional Bayesian variable selection approach was used to relate spatial EEG-prior information to voxelwise selection probabilities of a global stimulus regressor component. These probabilities were then used as indicators for activation, i.e. as activation probabilities. Continuous or binary EEG information was thereby connected to activation probabilities by using a latent probit regression stage. The probit predictor was chosen to consist of at least one spatially-varying effect (intercept or EEG component or both) to adapt to local brain response. Spatially-varying effects were regularized by an intrinsic Gaussian Markov random field (IGMRF) prior to ensure identifiability of voxelwise effects and to impose a dependency structure on neighbouring voxels. Results from a Gaussian conditional autoregressive (CAR) prior can be found in Kalus (2012). Inference was based on a Markov chain Monte Carlo (MCMC) sampling scheme. Two algorithms, iMRF_{glob} and iMRF_{flex} , were proposed, which differ in the way how they incorporate EEG information. Their performance was examined in reference to an activation detection scheme without EEG-enhancement (iMRF_0).

In case of iMRF_{flex} , an spatially-varying EEG coefficient regulates the EEG influence on activation detection and adapts it to local conditions. That is, it increases EEG influence in brain parts where EEG and fMRI are in accordance with each other and downweights EEG influence where it contradicts the fMRI signal. This implicates that the algorithm is quite robust against misspecifications in EEG-prior choice. In Sect. 4, we have also seen that iMRF_{flex} might use non-congruent prior information as information on non-activation. This finding, however, has to be corroborated in further studies. Generally, a requirement for increasing sensitivity in regions of special interest seems to be that these are not too close to regions with non-overlapping fMRI and EEG information. If these areas are distant enough, the downweighting mechanism, which is related to good edge-preserving properties, does not interfere with the positive EEG influence in

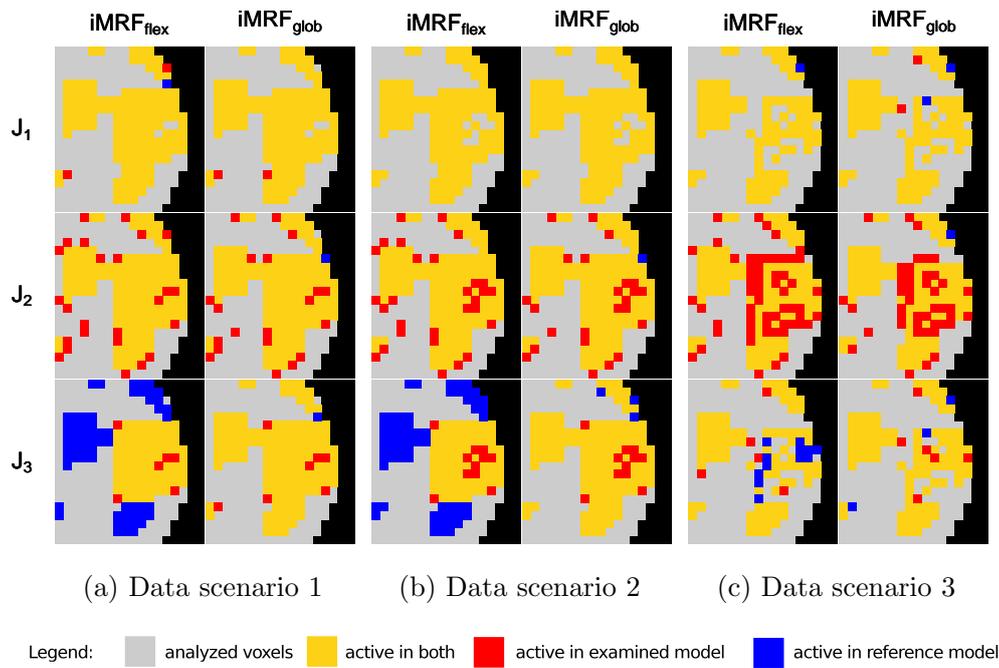


Figure 6: Simulation results: Activation difference maps of $iMRF_{flex}$ resp. $iMRF_{glob}$ in reference to $iMRF_0$.

other regions. Thereby, the corresponding minimal distance depends on the degree of smoothness in the data and is likely to span only a small number of voxels. To assess congruency, EEG effect map estimates can give information about the overlap of fMRI and EEG data and, hence, give hints on possible problems.

In case of iMRF_{glob} , the global EEG effect possesses the ability to increase prior activation probabilities proportionally to the observed EEG measurements as soon as a reasonable large, positive EEG effect is retrieved. If necessary, the spatially-varying intercept can have a region-specific compensating effect to improve the goodness-of-fit. The size of the EEG effect depends on the degree of congruency of fMRI and EEG data. If EEG information does not match fMRI information in large parts of the analysed brain, a virtually zero global EEG effect estimate is recovered. Hence, this algorithm is equally robust against misspecifications, but, in contrast to iMRF_{flex} , it loses its ability to increase sensitivity in remaining congruent areas.

Comparing the performance of both predictor types, we generally expect that iMRF_{flex} models possess a better performance than iMRF_{glob} models when non-congruency is present in the data. The iMRF_{glob} algorithm retrieves a global EEG effect by averaging over all analysed brain parts, so that the effect may be decreased (even to virtually zero estimates) if too much non-congruency is in the data. In contrast to this, iMRF_{flex} is able to adapt the EEG effect to local conditions. As long as regions are sufficiently separated, effects in non-congruent areas do not substantially influence effects in congruent areas, which can retain their high sensitivity level. Generally, we observed that both algorithms are approximately as sensitive as their corresponding uninformed fMRI activation detection algorithm iMRF_0 in most analyses.

Indeed, iMRF_{glob} and iMRF_{flex} had a promising performance, because they appeared to recover interpretable estimates by retaining the high sensitivity level of iMRF_0 . However, despite our expectations, EEG-enhanced fMRI activation detection schemes only revealed their benefit in data settings with certain properties. From these findings, we suppose that they are useful in regions where fMRI data possess activation structures of reasonable large size (not just isolated peaks) with gaps in the structure (e.g. due to noise). Within this region, non-active voxels containing a damped activation signal can then be classified as active if prior information—in form of sufficiently strong EEG measurements—indicates activation at these locations. Thereby the prior information must be fairly congruent to fMRI activation

information overall the analysed brain.

In our real-world example, the data structure deviated from these guidelines. Therefore, only a relatively small number of voxels switched their status from being non-active to active. Hence, we saw indeed that our EEG enhancement possesses some potential to increase sensitivity, but in this particular application, it had not lived up to expectations. In controlled settings, like in the simulation study, we observed, however, a substantial performance increase with data that possessed the above properties.

We conclude with some remarks on current or future research: The assumption of independent error terms might be relaxed by incorporating temporal correlations in voxelwise fMRI regressions. We do not expect that the performance of proposed algorithms is strongly influenced by this, but a more realistic model might be achieved through this. Marginalization steps to fMRI regression parameters, however, might not be feasible anymore—leading to a far inefficient sampling procedure.

Incorporation of brain partitioning information is of particular interest (see, for example, the description in Appendix B of Daunizeau et al., 2007). Estimation of random fields and corresponding variance parameters then could be broken down into parts and be accomplished within (more) homogeneous brain areas. This would be advantageous in several ways: The speed of the algorithms could be increased, smoothness of random field may be allowed to vary across the brain and for iMRF_{glob} non-congruency issues (leading to small global EEG effect estimates) can be resolved by this as well.

The existing approaches can further be modified by changing the fitting algorithm. Though the full Bayesian MCMC approach finds the exact posterior distribution of model parameters, it is rather slow. Another popular approach in neuroimaging is a variational Bayes approach to Bayesian inference (see e.g. Titterton, 2004). It approximates the joint posterior distribution of all unknown parameters with a simpler distribution usually positing further independence assumptions than those implied by the original generative model. By this, the computational demand can be decreased substantially. However, a comparison between full and approximate Bayesian approaches has to be conducted to check whether the speed is achieved at the expense of accuracy.

Note that in this work, we focussed on modelling activation provoked by just one stimulus type. Nevertheless, it is conceivable that other stimulus types, which are presented within the same experimental procedure, are either included as confounders or as a further additive stimulus component

that is subject to Bayesian activation detection. Generally, our methodology is not applicable to stimulus-free experimental procedures like resting state observations. For being inevitably connected to a regressor modelling some kind of stimulus that provokes neuronal activation, an extension to stimulus-free settings is not conceivable.

As noted above, the EEG-enhanced activation detection schemes have not lived up to our expectations in the application at hand. Nevertheless, in the simulation study $iMRF_{flex}$ and $iMRF_{glob}$ yielded promising results. Hence, demonstrating their potential in other data settings, with suitable structure and carefully selected 3D prior knowledge (either derived from EEG or other external information), will be of great interest.

6. Acknowledgements

We are grateful to Sara A. Kiem for supporting data acquisition of the acoustic oddball paradigm, and Manuela Hölzhammer and Ludwig Bothmann for assisting the EEG source localization process. This research has been funded by the German Science Foundation, grant FA 128/6-1.

References

- Albert, J.H., Chib, S., 1993. Bayesian analysis of binary and polychotomous response data. *J Am Stat Assoc* 88, 669–679.
- Bandettini, P.A., 2012. 20 years of fMRI [Special issue]. *NeuroImage* 62, 575–1324.
- Barbieri, M.M., Berger, J.O., 2004. Optimal predictive model selection. *The Annals of Statistics* 32, 870–897.
- Beckmann, C.F., Smith, S.M., 2005. Tensorial extensions of independent component analysis for multisubject FMRI analysis. *NeuroImage* 25, 294–311.
- Czisch, M., Wehrle, R., Stiegler, A., Peters, H., Andrade, K., Holsboer, F., Sämann, P.G., 2009. Acoustic oddball during NREM sleep: a combined EEG/fMRI study. *PLoS ONE* 4, e6749.

- Daunizeau, J., Grova, C., Marrelec, G., Mattout, J., Jbabdi, S., Péligrini-Issac, M., Lina, J.M., Benali, H., 2007. Symmetrical event-related EEG/fMRI information fusion in a variational Bayesian framework. *NeuroImage* 36, 69–87.
- Daunizeau, J., Laufs, H., Friston, K.J., 2010. EEG-fMRI information fusion: Biophysics and data analysis, in: Mulert, C., Lemieux, L. (Eds.), *EEG - fMRI: Physiological Basis, Technique, and Applications*. Springer-Verlag, Berlin Heidelberg. chapter 25, pp. 511–526.
- Davis, T.A., Hager, W.W., 1999. Modifying a sparse Cholesky factorization. *SIAM. J. Matrix Anal. & Appl.* 20, 606–627.
- Friston, K.J., Ashburner, J.T., Kiebel, S.J., Nichols, T.E., Penny, W.D., 2008. *Statistical parametric mapping - the analysis of functional brain images*. Academic Press.
- Friston, K.J., Penny, W.D., Phillips, C., Kiebel, S.J., Hinton, G., Ashburner, J.T., 2002. Classical and Bayesian inference in neuroimaging: Theory. *NeuroImage* 16, 465–483.
- Gelfand, A.E., Smith, A.F.M., 1990. Sampling-based approaches to calculating marginal densities. *JASA J Am Stat Assoc* 85, 398–409.
- Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 2004. *Bayesian data analysis*. Chapman & Hall/CRC, Boca Raton. 2nd editio edition.
- George, E.I., McCulloch, R.E., 1997. Approaches for Bayesian variable selection. *Statistica Sinica* 7, 339–373.
- Gilks, W.R., Richardson, S., Spiegelhalter, D.J. (Eds.), 1996. *Markov chain Monte Carlo in practice*. Chapman & Hall/CRC, London.
- Gössl, C., Auer, D.P., Fahrmeir, L., 2001. Bayesian spatiotemporal inference in functional magnetic resonance imaging. *Biometrics* 57, 554–562.
- Groves, A.R., Chappell, M.A., Woolrich, M.W., 2009. Combined spatial and non-spatial prior for inference on MRI time-series. *NeuroImage* 45, 795–809.
- Henson, R., Rugg, M.D., Friston, K.J., 2001. The choice of basis functions in event-related fMRI. *NeuroImage* 13, 149.

- Josephs, O., Turner, R., Friston, K.J., 1997. Event-related fMRI. *Hum Brain Mapp* 5, 243–248.
- Kalus, S., 2012. Biostatistical Modeling and Analysis of Combined fMRI and EEG Measurements. Ph.d. thesis. Ludwig-Maximilians University, Munich.
- Kalus, S., Sämann, P.G., Fahrmeir, L., 2013. Classification of brain activation via spatial Bayesian variable selection in fMRI regression. Accepted for publication in *Adv Data Anal Classif*.
- Kiehl, K.A., Stevens, M.C., Laurens, K.R., Pearlson, G., Calhoun, V.D., Liddle, P.F., 2005. An adaptive reflexive processing model of neurocognitive function: Supporting evidence from a large scale ($n = 100$) fMRI study of an auditory oddball task. *NeuroImage* 25, 899–915.
- Lang, S., Brezger, A., 2004. Bayesian P-Splines. *J Comput Graph Stat* 13, 183–212.
- Laufs, H., Daunizeau, J., Carmichael, D.W., Kleinschmidt, A., 2008. Recent advances in recording electrophysiological data simultaneously with magnetic resonance imaging. *NeuroImage* 40, 515–528.
- Mandelkow, H., Halder, P., Boesiger, P., Brandeis, D., 2006. Synchronization facilitates removal of MRI artefacts from concurrent EEG recordings and increases usable bandwidth. *NeuroImage* 32, 1120–1126.
- Michel, C.M., Murray, M.M., Lantz, G., Gonzalez, S., Spinelli, L., Grave de Peralta, R., 2004. EEG source imaging. *Clinical Neurophysiology* 115, 2195–2222.
- Mulert, C., Lemieux, L. (Eds.), 2010. EEG - fMRI: Physiological Basis, Technique, and Applications. Springer-Verlag, Berlin Heidelberg.
- Müller, P., Parmigiani, G., Robert, C., Rousseau, J., 2004. Optimal Sample Size for Multiple Testing. *Journal of the American Statistical Association* 99, 990–1001.
- Penny, W.D., Trujillo-Barreto, N.J., Friston, K.J., 2005. Bayesian fMRI time series analysis with spatial priors. *NeuroImage* 24, 350–62.

- Ritter, P., Villringer, A., 2006. Simultaneous EEG-fMRI. *Neuroscience and Biobehavioral Reviews* 30, 823–838.
- Rosa, M.J., Daunizeau, J., Friston, K.J., 2010. EEG-fMRI integration: A critical review of biophysical modeling and data analysis approaches. *Journal of Integrative Neuroscience* 9, 453–476.
- Rosen, B.R., Buckner, R.L., Dale, A.M., 1998. Event-related functional MRI: Past, present, and future. *PNAS* 95, 773–780.
- Rue, H., 2001. Fast sampling of Gaussian Markov random fields. *J Roy Stat Soc B* 63, 325–338.
- Rue, H., Held, L., 2005. *Gaussian Markov random fields: Theory and applications*. Chapman & Hall, London.
- Smith, D., Smith, M., 2006. Estimation of binary Markov random fields using Markov chain Monte Carlo. *J Comput Graph Stat* 15, 207–227.
- Smith, M., Fahrmeir, L., 2007. Spatial Bayesian variable selection with application to functional magnetic resonance imaging. *J Am Stat Assoc* 102, 417–431.
- Smith, M., Kohn, R., 1996. Nonparametric regression using Bayesian variable selection. *J Econometrics* 75, 317–343.
- Smith, M., Pütz, B., Auer, D.P., Fahrmeir, L., 2003. Assessing brain activity through spatial bayesian variable selection. *NeuroImage* 20, 802–815.
- Titterton, D.M., 2004. Bayesian methods for neural networks and related models. *Stat Sci* 19, 128–139.