

Fahrmeir, Ludwig; Auer, Dorothee P.; Pütz, Benno; Smith, Michael

Working Paper

Assessing Brain Activity through Spatial Bayesian Variable Selection

Discussion Paper, No. 316

Provided in Cooperation with:

Collaborative Research Center (SFB) 386: Statistical Analysis of discrete structures - Applications in Biometrics and Econometrics, University of Munich (LMU)

Suggested Citation: Fahrmeir, Ludwig; Auer, Dorothee P.; Pütz, Benno; Smith, Michael (2003) : Assessing Brain Activity through Spatial Bayesian Variable Selection, Discussion Paper, No. 316, Ludwig-Maximilians-Universität München, Sonderforschungsbereich 386 - Statistische Analyse diskreter Strukturen, München, <https://doi.org/10.5282/ubm/epub.1697>

This Version is available at:

<https://hdl.handle.net/10419/23874>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Assessing Brain Activity through Spatial Bayesian Variable Selection

Michael Smith ^{a,*} Benno Pütz ^b Dorothee Auer ^b
Ludwig Fahrmeir ^c

^a*University of Sydney, Australia*

^b*NMR Research Group, Max-Planck-Institute of Psychiatry, Munich, Germany*

^c*Institute of Statistics, Ludwig Maximilian University Munich, Germany*

Abstract

Statistical parametric mapping (SPM), relying on the general linear model and classical hypothesis testing, is a benchmark tool for assessing human brain activity using data from fMRI experiments. Friston *et al.* (2002a) discuss some limitations of this frequentist approach and point out promising Bayesian perspectives. In particular, a Bayesian formulation allows explicit modeling and estimation of activation probabilities. In this paper, we directly address this issue and develop a new regression based approach using spatial Bayesian variable selection. Our method has several advantages. First, spatial correlation is directly modeled for activation probabilities and indirectly for activation amplitudes. As a consequence, there is no need for spatial adjustment in a post-processing step. Second, anatomical prior information, such as the distribution of grey matter or expert knowledge, can be included as part of the model. Third, the method has superior edge-preservation properties as well as being fast to compute. When applied to data from a simple visual experiment, the results demonstrate improved sensitivity for detecting activated cortical areas and for better preserving details of activated structures.

Key words: Activation probabilities, anatomical prior information, Bayesian spatial modeling, edge-preservation, fMRI experiments, Ising prior, Markov chain Monte Carlo sampling

* Econometrics and Business Statistics,
Merewether Building, H04,
University of Sydney.
Phone (+612) 9351 2787
Fax (+612) 9351 6409
email: mikes@econ.usyd.edu.au

1 Introduction

Functional magnetic resonance imaging (fMRI) has evolved into a powerful tool to non-invasively assess the functionality of neuronal networks. A crucial part in these studies is the reliable distinction of activated from non-activated voxels based on the statistical analysis of a functional time series. Standard validated models exist to perform group analysis for neuroscientific applications aimed at reducing false-positive activations. However, assessing focal brain activity individually is much more demanding and may require a focus on the reduction of false-negative voxel assignment, as in presurgical mapping tasks. The statistical power of such fMRI studies in single neurosurgical cases is further reduced by the typically older age compared to volunteers for neuropsychological studies in view of the known age-dependent reduction in blood oxygenation (or BOLD) response; see Huettel (2001) for a discussion of this issue. To overcome these limitations it would be desirable to increase the sensitivity of the analysis over standard procedures.

Interestingly, standard statistical analysis of fMRI data does not exploit individual anatomic information nor functional neuroanatomic knowledge, both of which may usefully complement the information provided by the functional time series. Local physiological responses upon neuronal activation is known to be restricted to cortical and subcortical neuronal assemblies. Thus, the spatial information about the distribution of cortical and subcortical gray matter (as opposed to white matter and cerebro-spinal fluid structures which both do not show BOLD-responses to neuronal activation) may be helpful.

The package ‘Statistical Parametric Mapping’ (SPM) is in common use as a standard tool for assessing human brain activity using data from fMRI experiments. SPM relies on the parametric general linear model and on the principles of classical statistics as outlined by Friston *et al.* (1995). It calculates t -values for activation effects separately for all voxels within a slice. The theory of stationary Gaussian random fields provides a way to account for spatial correlation between voxels and to adjust the corresponding t -values in a post-processing step. Inferences about regionally specific effects are then based on the resulting spatially smoothed map of t -statistics.

Friston *et al.* (2002b) discuss some limitations of this classical approach and ways in which Bayesian inference can be used to improve the analysis. A fundamental issue is “that the [frequentist] p -value does not reflect the likelihood that the effect is present, but simply the probability of getting the observed data in the effect’s absence.” This implies several shortcomings described in Friston *et al.* (2002b, pp 466). For example, the alternative hypothesis that an activation has not occurred cannot be rejected because the probability for this event is zero. Furthermore, given enough repeated observations, one can

always find a false-positive effect at every voxel. However, in comparison a Bayesian formulation allows the explicit modeling of the probability that a voxel has been activated. After the collection of the data, this probability of activation can be calculated using modern computational techniques as the marginal posterior probability of activation at each voxel.

In this paper, we directly address this question and suggest a new regression based approach using spatial Bayesian variable selection and model averaging. For each voxel, a binary indicator variable explicitly parameterizes whether the voxel is ‘active’ or ‘inactive’. That is, whether the corresponding effect in a regression model is non-zero or zero, respectively. Spatial correlation is introduced into the model by an Ising prior for the set of binary indicator variables. As noted in computational physics (see Swendsen and Wang (1987) and references therein), the model encourages the formation of clusters of like-valued binary variables. Bayesian posterior analysis using a carefully tailored Markov chain Monte Carlo (MCMC) algorithm allows direct calculation of parameters and features of interest, and in particular the marginal posterior probability of activation at each voxel.

In comparison to previous methods, this approach has several significant advantages. First, spatial correlation is directly modeled for the activation probabilities themselves and the activation amplitudes only indirectly. This is in contrast to approaches which model spatial correlation directly in the activation amplitudes, and indirectly in the probabilities of activation; for examples see Gössl *et al.* (2001a; 2001b). Second, because the posterior distribution incorporates the spatial structure of the Ising prior, there is no need for spatial adjustments in a post-processing step. Third, an extremely useful property is that prior information about likely areas of activation can be incorporated in a natural and straightforward way through the ‘external field’ of the Ising prior. This can include anatomical information, such as the distribution of grey matter, or expert knowledge regarding likely areas of activation in response to an experiment. Fourth, the Ising model on the binary activation indicators has substantially improved edge preservation for spatial smoothing of activation probabilities in comparison to approaches that induce probabilities indirectly after spatial smoothing of activation amplitudes. Last, the method proves to be computationally very efficient and fast, providing full processing of slices in under a minute on Pentium III based PCs.

The paper is organized as follows. In section 2 we outline our Bayesian approach in detail. In particular, we discuss regression modeling of fMRI data and how spatial Bayesian variable selection can be applied to such regression models to obtain posterior activation and amplitude maps. We also highlight how expert prior knowledge or anatomical information can be used to inform the maps. In section 3 we demonstrate the features of the method and improved sensitivity of the resulting analysis by applying the results to data from a sim-

ple visual experiment. In section 4 we conclude the paper with a discussion of our Bayesian method in relation to alternative methods of analysis for fMRI data that are current in the literature.

2 Methods

Regression modeling of fMRI data

Let $\{y_{it}, t = 1, \dots, T\}$ denote the MR signal time series at voxel i , for $i = 1, \dots, N$. All regression based approaches assume an additive model of the form

$$y_{it} = a_{it} + f_{it} + \varepsilon_{it}, \quad i = 1, \dots, N; t = 1, \dots, T,$$

where $\{a_{it}, t = 1, \dots, T\}$ is a baseline trend at voxel i , $\{f_{it}, t = 1, \dots, T\}$ is the activation profile and $\{\varepsilon_{it}, t = 1, \dots, T\}$ is the measurement noise. Various models differ in the way these three components are specified and how spatial correlation is taken into account. Conventional parametric voxel-wise regression analysis, such as that in Friston *et al.* (1995), assumes a linear model

$$y_{it} = w_t' \alpha_i + z_{it} \beta_i + \varepsilon_{it}. \quad (1)$$

The baseline trend $a_{it} = w_t' \alpha_i$ is usually modeled as a linear combination of a few simple basis functions w_t , such as the lower frequency terms of a Fourier expansion or piecewise continuous polynomials, with coefficients $\alpha_i = (\alpha_{i1}, \dots, \alpha_{im})'$. The activation profile, $f_{it} = z_{it} \beta_i$, is defined as the product of a scalar activation effect or ‘amplitude’ β_i with a transformed stimulus $\{z_{it}, t = 1, \dots, T\}$. This is a delayed and continuously modified version of the original binary valued ‘boxcar’ stimulus $\{x_t, t = 1, \dots, T\}$. The transformation is motivated by the fact that blood oxygenation is delayed by a short time lag d_i and has a smooth unimodal distribution over time, rather than the distribution of the original boxcar untransformed stimulus.

The transformation is often described by a convolution with a density function, such as a Poisson or gamma density $h(\cdot, \lambda_i)$, so that

$$z_{it} = \sum_{s=0}^{t-d_i} h(s, \lambda_i) x_{t-d_i-s}. \quad (2)$$

Here, the lags $\{d_1, \dots, d_N\}$ and the transformation parameters $\{\lambda_1, \dots, \lambda_N\}$ are usually estimated in a pre-processing step.

Other approaches which incorporate estimation of the transformation of the stimulus jointly with the regression model have also been proposed— see Gössl *et al.* (2001b). In addition, more flexible nonparametric fitting of the baseline trends a_{it} and time-varying activation amplitudes can also be performed with additive regression models; see Genovese (2000) and Gössl *et al.* (2000). While these approaches provide a more complete statistical solution to the modeling of fMRI data, they are not necessarily so attractive to the practitioner for two reasons. First, increased model flexibility in these areas does not appear to substantially improve empirical estimates of activation maps. Second, the prompt processing of each fMRI acquisition is important for diagnostic purposes. There is a trade-off between model complexity and computational feasibility, and introducing flexibility in these areas substantially increases processing time. Therefore, we use the parametric regression model (1) as a basic ingredient of our spatial Bayesian variable selection and model averaging approach, and work with pre-processed transformed stimuli.

Spatial Bayesian variable selection

Defining the vectors $y_i = (y_{i1}, \dots, y_{iT})'$, $z_i = (z_{i1}, \dots, z_{iT})'$, $\varepsilon_i = (\varepsilon_{i1}, \dots, \varepsilon_{iT})'$ of fMRI signals, transformed stimuli, and errors at voxel i , respectively, and the design matrix $W = (w_1, \dots, w_T)'$ for the baseline trend, model (1) with time-constant activation amplitude is given by

$$y_i = W\alpha_i + z_i\beta_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma_i^2 I). \quad (3)$$

We define a spatial Bayesian model by introducing binary indicator variables

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

so that the activation amplitude

$$\begin{aligned} \beta_i &\neq 0 \text{ iff } \gamma_i = 1, \\ \beta_i &= 0 \text{ iff } \gamma_i = 0. \end{aligned}$$

The vector of binary indicators $\gamma = (\gamma_1, \dots, \gamma_N)'$ then represents an activation surface, while $\beta = (\beta_1, \dots, \beta_N)'$ is a vector of corresponding activation amplitudes. We process each slice of an fMRI scan separately, so that these vectors are defined over a regular lattice.

Spatial correlation and anatomical or expert prior information is introduced through an Ising prior for the activation vector γ . Ising priors are commonly used as a Markov random field (MRF) prior in image analysis with binary valued pixels, such as analysis of black and white images; see Higdon (1998) and Swendsen and Wang (1987) and references therein for examples. This is of particular advantage in fMRI analysis because it allows one to incorporate spatial correlation directly on the probabilities of activation. In addition, anatomical or expert prior information on activation of voxels may be incorporated through the ‘external field’ of the Ising prior.

The prior is defined as

$$p(\gamma) \propto \exp \left\{ \sum_{i=1}^N \delta_i \gamma_i + \theta \sum_{i \sim j} w_{ij} I(\gamma_i = \gamma_j) \right\}, \quad (4)$$

where $I(A) = 1$ if A is true, and $I(A) = 0$ if A is false. The sum $\sum_{i=1}^N \delta_i \gamma_i$ is the external field, while the interaction term $\theta \sum_{i \sim j} w_{ij} I(\gamma_i = \gamma_j)$ models spatial correlation. The sum over $i \sim j$ is notation for summation over all neighboring voxels, while the w_{ij} are pre-specified constants that weight the interaction between neighboring locations i and j on the lattice. Here, the neighborhood of i is defined to contain the directly adjacent vertical and horizontal voxels j with weights $w_{ij} = 1$, and the directly adjacent diagonal voxels j with weights $w_{ij} = 1/\sqrt{2}$. The parameter θ controls the amount of spatial smoothing, with the elements of γ being independent if $\theta = 0$. Note that while this prior directly introduces spatial smoothing for the binary activation indicators, it also indirectly introduces spatial smoothing for the activation amplitudes.

To complete the Bayesian model, priors are required for all the remaining parameters in the model. From the definition of the indicators, $\beta_i = 0 | \gamma_i = 0$ and a prior is only required for the conditional distribution of $\beta_i | \gamma_i = 1$. Following extensive discussion in the statistical literature (see Smith and Kohn (1996), George and McCulloch (1997) and references therein) a proper prior must be assigned to the non-zero coefficients to undertake variable selection and model averaging. Following Kohn *et al.* (2001), we define a proper conditional prior by setting

$$p(\beta_i | \gamma_i = 1, \alpha_i, \sigma_i^2, y_i) \propto p(y_i | \beta_i, \gamma_i = 1, \alpha_i, \sigma_i^2)^{1/T}. \quad (5)$$

The prior in (5) is a fractional conditional prior, which uses a fraction $1/T$ of the data to give a proper prior and is uninformative relative to the likelihood.

For the remaining parameters $\alpha_i = (\alpha_{i1}, \dots, \alpha_{im})$ and $\sigma_1^2, \dots, \sigma_N^2$ we assume standard independent non-informative priors, so that

$$p(\alpha_1, \dots, \alpha_N, \sigma_1^2, \dots, \sigma_N^2) \propto \prod_{i=1}^N 1/\sigma_i^2.$$

Incorporating anatomical prior information

To incorporate anatomical prior information, define a grey matter indicator $g_i = 1$ if voxel i is grey matter, and $g_i = 0$ otherwise, with $p(g_i = 1)$ as the grey matter prior obtained from segmentation. Assume also that the grey matter indicators are a priori independent across voxels, so that $p(g_1, \dots, g_N) = \prod_{i=1}^N p(g_i)$. Then it is known that $p(\gamma_i = 1 | g_i = 0) = 0$ because activation cannot occur outside grey matter, and that $p(\gamma_i = 1 | g_i = 1) = a$, where for the visual experiment in our empirical application we assume $a = 0.1$. Note that this is a reasonable assumption, because it is known a priori that only five to ten percent of grey matter is likely to be activated by this experiment. Then

$$p(\gamma_i = 1) = p(\gamma_i = 1 | g_i = 1)p(g_i = 1) = p(g_i = 1)a = c_i.$$

We match the marginal prior probabilities c_i to the external field of the Ising prior when there is no spatial correlation, so that $\theta = 0$. In this case the joint density is $p(\gamma) \propto \exp\{\sum_i \delta_i \gamma_i\}$, with marginals

$$p(\gamma_i = 1) = \exp(\delta_i) / (\exp(\delta_i) + 1) = c_i \text{ for } i = 1, \dots, N.$$

Therefore, solving for δ_i , an anatomically informed Ising prior is given by equation (4) with $\delta_i = \log(c_i / (1 - c_i))$.

Posterior Inference

Bayesian model estimation and inference is based on posterior quantities; that is, inference conditional on the data y collected. In the case of fMRI analysis the quantities of immediate interest are the marginal posterior probabilities of activation $p(\gamma_i = 1 | y)$ and marginal posterior expectations of the activation amplitudes $E(\beta_i | y)$ for all the voxels $i = 1, 2, \dots, N$ in a scan. Closed form calculation of these quantities is not possible because it involves high order integration out of nuisance parameters. However, if a Monte Carlo sample $\{\gamma^{[1]}, \gamma^{[2]}, \dots, \gamma^{[J]}\}$ is available from the joint posterior distribution $p(\gamma | y)$, then it is possible to estimate these two posterior quantities as

$$p(\gamma_i = 1|y) = \int p(\gamma_i = 1|\gamma_{j \neq i}, y)p(\gamma_{j \neq i}|y)d\gamma_{j \neq i} \approx \frac{1}{J} \sum_{k=1}^J p(\gamma_i = 1|\gamma_{j \neq i}^{[k]}, y),$$

$$E(\beta_i|y) = E(\beta_i|\gamma_i = 1, y_i)p(\gamma_i = 1|y) \approx E(\beta_i|\gamma_i = 1, y_i) \left(\frac{1}{J} \sum_{k=1}^J p(\gamma_i = 1|\gamma_{j \neq i}^{[k]}, y) \right).$$

These are Monte Carlo mixture estimates and are in common usage in computational statistics—see Tanner (1996) p.51. We label the two estimates when calculated for the N voxels of an fMRI slice as an ‘activation map’ and an ‘amplitude map’, respectively.

To obtain the Monte Carlo sample, we use a Markov chain Monte Carlo (MCMC) sampling scheme. Such sampling schemes have become the most widely used computational approach in modern statistics. MCMC techniques are based on repeated sampling from conditional posterior distributions, defining a Markov chain with the target distribution as its stationary distribution; see Gamerman (1997) and references therein for an introduction to MCMC schemes. We outline a very efficient and fast MCMC sampling in the Appendix A.2, which produces the required sample from the target distribution $p(\gamma|y)$.

To calibrate the activation maps with maps from the frequentist approach found in SPM, we follow Raftery (1996) and note that $-2 \log((1 - p(\gamma_i = 1|y))/p(\gamma_i = 1|y))$ is on the same scale as a likelihood ratio statistic and is approximately distributed $\chi^2(1)$. A p -value of 0.05 gives a critical value of 3.841, so that solving for the posterior probability $p(\gamma_i = 1|y) = 0.8722$ at this critical value. Therefore, a Bayesian classification rule calibrated to the frequentist p -value of 0.05 is to classify all voxels with $p(\gamma_i = 1|y) > 0.8722$ as active, and inactive otherwise.

Data

MR data were acquired on a clinical 1.5 T scanner (GE Medical Systems) from five healthy male volunteers (age 21–37 years). For each subject a functional time series was obtained for a visual stimulus paradigm (8 Hz alternating black and white checkerboard) and a 3D whole brain anatomical data set were acquired. For the functional acquisition an EPI sequence was employed to obtain 23 contiguous slices with TR=4 s, TE=60 ms and a spatial resolution of $2.19 \times 2.19 \times 5 \text{ mm}^3$. The paradigm consisted of an on-off block design with 9 epochs (5 off, 4 on) lasting 7 images each.

All data sets were post-processed using SPM99. The images of each time series were realigned (motion corrected) and normalized to a template in MNI standard space retaining the original spatial resolution. Activation maps were calculated on these data sets with additional spatial smoothing using SPM’s

Gaussian field correction at a threshold of 0.05. Grey matter prior information (g_i) was obtained from an SPM segmentation of the first image of each functional data set.

3 Results

To explore the effect of the Ising prior at equation (4), we analyze a single slice of the scan from individual A using the spatial Bayesian variable selection procedure (SpBVS) with four different values for the hyperparameters. The four Ising priors used are given in table A.1. Figures A.1 and A.2 provide the Monte Carlo estimates of the marginal posterior probabilities of activation $p(\gamma_i = 1|y)$ and corresponding marginal posterior means of the amplitudes $E(\beta_i|y)$. These are calculated for all voxels in the slice and when plotted over the lattice, they form posterior activation and amplitude maps, respectively. In figure A.1, voxels classified as activated (that is, with estimates of $p(\gamma_i = 1|y) > 0.8772$) are plotted in white. The anatomically informed grey matter prior is plotted in the background in grey. In figure A.2 the amplitude map is thresholded, so that only the amplitudes of the activated voxels are shown.

—TABLE A.1 ABOUT HERE—

Ising model (a) involves no spatial smoothing and no anatomical prior information. The resulting posterior activation map has many false-positives (voxels incorrectly identified as activated), including voxels in areas which have low probability of being grey matter. Ising model (b) involves no spatial smoothing, but incorporates the anatomically informed grey matter prior from automated SPM segmentation using signal intensity and prior knowledge from the template. The anatomical prior reduces the false-positives in the posterior activation map in areas of the brain that have low probability of being grey matter from the segmentation data. Moreover, in areas outside the skull which have zero prior probability of being grey matter, both the marginal prior (and therefore posterior) probability of activation is zero.

Models (c) and (d) involve positive spatial smoothing with $\theta = 0.6$, which thins out many of what are likely to be incorrectly identified activated voxels and bolsters regions of multiple activated voxels. Model (d) also incorporates the grey matter prior and trims the activation profile down to the observed anatomy of the brain. This latter model is the one recommended for analysis, and we note that the maps suggest that there is indeed an acoustic co-activation in response to the visual experiment for this individual. This co-activation is not identified in either the default SPM analysis or using continuous MRF priors on the amplitudes as in Fahrmeir *et al.* (2002). Nevertheless, the acoustic activation does not appear to be a false-positive, with the am-

plitude map at figure A.2(d) identifying these voxels as possessing very high amplitudes. To confirm this analysis, the raw time series from these acoustic voxels were examined and found to display strong reaction to the stimulus.

—FIGURES A.1 AND A.2 ABOUT HERE—

Figure A.3 shows the activation maps resulting from applying the SpBVS procedure to eight contiguous slices from a scan of individual B. This subject possesses a large visual cortex area, and there is a clear activation with little evidence of false-positives. Note that in slice 10 the small lateral geniculate nucleus regions are clearly identified as activate. This activation is unlikely to be a false-positive, given that this is a visual experiment. Figure A.4 contains the corresponding maps resulting from an SPM analysis of the same data. The SPM analysis corresponds to the default smoothing and an applied corrected threshold of $p < 0.05$. Comparing the results, we can see that that the detection sensitivity is lower, best noted in the extrastriatal visual cortex and the loss of structural details within activated areas. In particular, in the SPM analysis the lateral geniculate nucleus is not identified as active, which is a result of the spatial over-smoothing of the method of Friston *et al.* (1995).

—FIGURES A.3 AND A.4 ABOUT HERE—

The edge-preserving properties of the SpBVS methodology result in a spatial smoothing procedure that is more reliable in identifying small areas of activation. To examine this, we investigate whether, or not, the method identifies activation of the lateral geniculate nucleus (LGN) regions in the five individuals. For each individual, the slice that most corresponds to the LGN region was selected and the corresponding SpBVS posterior activation maps plotted in figure A.5. Clear activation of the LGN region was detected for individuals A and B, weak activation for individual D and no activation for individuals C and E. The corresponding SPM maps did not show any geniculate activation in any of the five subjects.

—FIGURE A.5 ABOUT HERE—

For the three individuals without clear activation of the LGN, additional expert prior knowledge was employed to enhance the map. In each of the three slices a neuroradiologist (DPA) identified a 10×10 voxel region R which is likely to contain the LGN. In these areas the prior probability of activation, conditional on the voxel being grey matter, is increased so that $p(\gamma_i = 1 | g_i = 1) = 0.5$, for $i \in R$. For the remaining voxels, we maintained the conditional prior probability, so that $p(\gamma_i = 1 | g_i = 1) = 0.1$ for $i \notin R$. The Ising prior employed is therefore as outlined in row (d) of table A.1, but

where

$$\delta_i = \begin{cases} \log\left(\frac{0.5p(g_i=1)}{1-0.5p(g_i=1)}\right) & \text{for } i \in R \\ \log\left(\frac{0.1p(g_i=1)}{1-0.1p(g_i=1)}\right) & \text{for } i \notin R. \end{cases}$$

Figure A.6 contains the resulting enhanced activation maps for individuals D and E. The anatomically informed prior with enhanced region R is plotted in the background. Note that when $\theta = 0$ the prior $p(\gamma_i = 1) \leq 0.1$ for $i \notin R$, while $p(g_i = 1) \leq 0.5$ for $i \in R$. In the figure, prior values in the range $0.1 < p(\gamma_i = 1) < 0.5$ are painted in shades of red, while prior values $p(\gamma_i = 1) \leq 0.1$ are in greyscale. As before, the voxels with marginal posterior probabilities of activation $p(\gamma_i = 1|y) > 0.8772$ are painted over with white. A comparison with the posterior activation maps in figure A.5 shows that for individual E activation in the LGN region has been enhanced. For individual D weak partial activation is now detected, whereas without the enhanced anatomical prior no activation is detected. However, for individual C no activation was detected even with the enhanced prior.

—FIGURE A.6 ABOUT HERE—

4 Discussion

Overall, the application of the SpBVS approach to the fMRI data from a simple visual experiment shows an increased sensitivity in detecting activated cortical areas when compared to the calibrated SPM99 spatially smoothed analysis. In particular, the proposed approach better preserves anatomical details of activated structures. This increased sensitivity is an important advantage in some fMRI applications. For example, in presurgical brain mapping in tumor patients great emphasis is put on avoiding spatial misregistration of activated regions and detecting any cortical activation in the vicinity of the surgical target. However, compared to the SPM analysis, we observe more isolated pixels suggestive of false-positive activation.

Another important feature of the method is the ability to use prior information to inform the activation and amplitude maps. This is undertaken in a structured way using a prior distribution which, when combined with the information in the data in the form of the likelihood, informs the resulting posterior probabilities of activation. For the simple visual experiment, analysis was based on grey matter priors obtained from segmenting the brain under study, as well as on additional expert knowledge to enhance the grey matter priors in a specific region of interest. Another possibility would be to use infor-

mation from previous multi-subject studies to define anatomically informed priors for a new subject. This flexibility in assigning prior probabilities for activation is a distinct advantage of our Bayesian method that is hard to replicate in non-Bayesian paradigms.

Recently, several alternative Bayesian spatial and spatio-temporal models for analyzing fMRI data have been proposed. Mostly, spatial or spatio-temporal priors are imposed on activation amplitudes, but not on activation indicators. Therefore, the focus is on deriving posterior distributions of the amplitudes, and less on direct calculation of posterior probabilities of activation or non-activation of voxels.

Gössl *et al.* (2001a) developed a hierarchy of Bayesian models, increasing in complexity from simple spatial models for time-constant activation amplitudes to flexible semiparametric spatio-temporal models allowing for time-space interactions. Corresponding spatial or spatio-temporal Gaussian Markov random field (GMRF) priors are assumed for activation amplitudes, but without incorporation of anatomical prior knowledge. A concern with such anatomically noninformed Gaussian priors is the tendency of oversmoothing in regions with high spatial curvature, for example at borders between high and low activation, or to smooth away biologically meaningful small local area effects. Therefore, Fahrmeir *et al.* (2002) replace GMRF priors by more robust and locally adaptive MRF priors. Although some improvement compared to GMRF priors is achieved, the potential to detect small local area effects remains limited. Moreover, in both cases, the Markov chain Monte Carlo algorithms used to estimate the models prove to be substantially slower than that used to estimate the model proposed in the current paper.

A different approach to specify spatial or spatio-temporal models for fMRI data incorporating anatomical prior knowledge has been developed by Kiebel *et al.* (2000), and adapted to multi-subject studies in Kiebel and Friston (2002). The basic idea is as follows: First, an approximate reconstruction of the grey matter surface is calculated. In a second step, this surface is projected onto a plane producing a flattened version of the grey matter surface. Gaussian isotropic basis functions are then defined with centers on a spatial grid of the flattened surface, and transformed back onto the original grey matter surface. The resulting basis functions are non-Gaussian, anisotropic, and locally adaptive through the anatomical information employed in their construction. Finally, spatially smoothed activation maps are obtained as regularized least squares fits, with the design matrix determined by these anatomically informed basis functions.

Based on this work, Friston *et al.* (2002a) recently suggested a hierarchical spatio-temporal Bayesian model which incorporates anatomical prior information. They impose a Gaussian prior on activation effects that is obtained

by modulating the covariance matrix of a stationary Gaussian random field using a grey matter prior. At least in current implementations, a problem with this approach is that it can only be applied to a restricted search volume. To avoid this computational problem, they combine principal component analysis and the use of anatomically informed basis functions developed in Kiebel *et al.* (2000). However, in contrast to our approach the prior probability of activation of a voxel is not part of the model, so that the posterior probability of activation is not directly available through the Bayesian paradigm. However, conceptually it seems possible to combine these approaches with spatial Bayesian variable selection.

To our knowledge, models which assign prior probabilities directly to indicators for activation or non-activation of voxels have been proposed only in a somewhat different and more restricted context. Based on the ideas of Everitt and Bullmore (1999), Hartvig and Jensen (2000) propose a spatial mixture model for the map of voxelwise t -values of activation effects estimated by classical SPM regression in a post-processing step. For each voxel i , a latent binary indicator γ_i is introduced, and the (marginal) distribution of its t -values t_i is specified by a mixture distribution $\pi f(t_i|\gamma_i = 1) + (1 - \pi)f(t_i|\gamma_i = 0)$. Here, π is the fraction of activated voxels, or is obtained through more complex mixture models. Spatial correlations are modeled by priors for activation profiles. Again, this approach is applied only to smaller search volumes, and it is not integrated into a complete spatio-temporal model.

A Appendix

A.1 The fractional conditional prior

We provide some details on the posterior distributions used in the sampling scheme. Note first that the fractional conditional prior in (5) is Gaussian

$$\beta_i|\gamma_i = 1, \alpha_i, \sigma_i^2, y_i \sim N(\mu_i, \sigma_i^2\tau_i),$$

where

$$\mu_i = \frac{z_i'(y_i - W\alpha_i)}{z_i'z_i}, \quad \tau_i = \frac{T}{z_i'z_i}.$$

A.2 Sampling activation probabilities

To draw the sample $\{\gamma^{[1]}, \dots, \gamma^{[J]}\}$ from the joint posterior $p(\gamma|y)$ of activation binary indicators, we construct a ‘single site’ sampler which repeatedly draws

from the full conditional posteriors $p(\gamma_i = 1 | \gamma_{-i}, y)$, where γ_{-i} is the vector γ without γ_i . These conditional posteriors are available in closed form as

$$p(\gamma_i = 1 | \gamma_{-i}, y) = 1/(1 + h),$$

where

$$h = \exp \left\{ -\delta_i + \ell_i + \theta \sum_{k \in \kappa_i} w_{ik} (1 - 2\gamma_k) \right\},$$

and $\kappa_i = \{k | k \sim i\}$ is the set of neighbors of voxel i . The value of ℓ_i can be calculated from

$$\ell_i = \log \left\{ \left(\frac{S_{i1}}{S_{i0}} \right)^{(T-m)/2} \left(\frac{|W' M_i W|}{|W' W|} \right)^{1/2} (T+1)^{1/2} \right\},$$

where

$$\begin{aligned} S_{i0} &= y_i' y_i - y_i' W (W' W)^{-1} W' y_i, \\ S_{i1} &= y_i' M_i y_i - y_i' M_i W (W' M_i W)^{-1} W' M_i y_i, \\ M_i &= I - z_i z_i' / (z_i' z_i), \end{aligned}$$

and $m = \dim(\alpha_i)$.

Proof:

The conditional posterior is obtained by considering

$$p(\gamma_i | \gamma_{-i}, y) \propto p(y | \gamma_{-i}, \gamma_i) p(\gamma_i | \gamma_{-i}) \propto A(\gamma_i),$$

calculating $A(\gamma_i = 1)$ and $A(\gamma_i = 0)$ and normalizing to find $p(\gamma_i = 1 | \gamma_{-i}, y) = 1/(1 + h)$, where $h = A(\gamma_i = 0)/A(\gamma_i = 1)$.

For the case when $\gamma_i = 1$,

$$\begin{aligned} p(\gamma_i = 1 | \gamma_{-i}, y) &\propto \iiint p(y_i | \gamma_i = 1, \beta_i, \sigma_i^2, \alpha_i) p(\beta_i | \gamma_i = 1, \sigma_i^2, \alpha_i) p(\sigma_i^2) d(\beta_i, \sigma_i^2, \alpha_i) p(\gamma_i = 1 | \gamma_{-i}) \\ &\propto (2\pi)^{-(T+1)/2} \iiint (\sigma_i^2)^{-(T+3)/2} (\tau_i)^{-1/2} \\ &\exp \left\{ -\frac{1}{2\sigma_i^2} \left[(y_i - W\alpha_i - z_i\beta_i)' (y_i - W\alpha_i - z_i\beta_i) + \frac{1}{\tau_i} (\beta_i - \mu_i)^2 \right] \right\} d(\beta_i, \sigma_i^2, \alpha_i) p(\gamma_i = 1 | \gamma_{-i}). \end{aligned}$$

The coefficient β_i can be integrated out by recognizing a $N(\mu_i, \frac{\sigma_i^2 \tau_i}{(T+1)})$ density. Therefore,

$$p(\gamma_i = 1 | \gamma_{-i}, y) \propto (2\pi)^{-T/2} \iint (\sigma_i^2)^{-(T+2)/2} (\tau_i)^{-1/2} (\tau_i / (T+1))^{1/2}$$

$$\begin{aligned} & \exp \left\{ -\frac{1}{2\sigma_i^2} \left[(y_i - W\alpha_i)'(y_i - W\alpha_i) + \frac{\mu_i^2}{\tau_i} - \mu_i^2 \frac{(T+1)}{\tau_i} \right] \right\} d(\sigma_i^2, \alpha_i) p(\gamma_i = 1 | \gamma_{-i}) \\ & \quad \propto (2\pi)^{-T/2} (T+1)^{-1/2} \iint (\sigma_i^2)^{-(T+2)/2} \\ & \exp \left\{ -\frac{1}{2\sigma_i^2} \left[(y_i - W\alpha_i)'(y_i - W\alpha_i) - \frac{1}{z_i' z_i} (y_i - W\alpha_i)' z_i z_i' (y_i - W\alpha_i) \right] \right\} d(\sigma_i^2, \alpha_i) p(\gamma_i = 1 | \gamma_{-i}) \end{aligned}$$

The parameter α_i is integrated out by recognizing a multivariate normal density with mean $\xi_i = (W' M_i W)^{-1} W' M_i y_i$ and variance $\Omega_i = \sigma_i^2 (W' M_i W)^{-1}$, with $M_i = I - \frac{1}{z_i' z_i} (z_i z_i')$. Therefore,

$$\begin{aligned} p(\gamma_i = 1 | \gamma_{-i}, y) & \propto (2\pi)^{-(T-m)/2} (T+1)^{-1/2} \int (\sigma_i^2)^{-(T+2)/2} |\Omega_i|^{1/2} \\ & \exp \left\{ -\frac{1}{2\sigma_i^2} \left[y_i' y_i - \frac{1}{z_i' z_i} y_i' z_i z_i' y_i \right] + \frac{1}{2} (\xi_i' \Omega_i^{-1} \xi_i) \right\} d\sigma_i^2 p(\gamma_i = 1 | \gamma_{-i}) \\ & = (2\pi)^{-(T-m)/2} (T+1)^{-1/2} |W' M_i W|^{-1/2} \int (\sigma_i^2)^{-(T-m+2)/2} \exp \left(-\frac{S_{i1}}{2\sigma_i^2} \right) d\sigma_i^2 p(\gamma_i = 1 | \gamma_{-i}) \end{aligned}$$

where S_{i1} is defined as above. The parameter σ_i^2 is integrated out as an inverse gamma $IG(\frac{T-m}{2}, \frac{S_{i1}}{2})$ density, so that

$$\begin{aligned} p(\gamma_i = 1 | \gamma_{-i}, y) & \propto \\ & (2\pi)^{-(T-m)/2} (T+1)^{-1/2} |W' M_i W|^{-1/2} \left(\frac{S_{i1}}{2} \right)^{-(T-m)/2} \Gamma \left(\frac{T-m}{2} \right) p(\gamma_i = 1 | \gamma_{-i}) = A(\gamma_i = 1). \end{aligned}$$

For the case when $\gamma_i = 0$,

$$p(\gamma_i = 0 | \gamma_{-i}, y) \propto \iint p(y_i | \gamma_i = 0, \sigma_i^2, \alpha_i) p(\sigma_i^2) d(\sigma_i^2, \alpha_i) p(\gamma_i = 0 | \gamma_{-i}).$$

Integrating out a normal density in α_i and an inverse gamma in σ_i^2 in a similar manner as above,

$$\begin{aligned} p(\gamma_i = 0 | \gamma_{-i}, y) & \propto \\ & (2\pi)^{-(T-m)/2} |W' W|^{-1/2} \left(\frac{S_{i0}}{2} \right)^{-(T-m)/2} \left(\frac{T-m}{2} \right) p(\gamma_i = 0 | \gamma_{-i}) = A(\gamma_i = 0). \end{aligned}$$

Normalizing the results, we get that

$$h = \exp\{\ell_i\} \frac{p(\gamma_i = 0 | \gamma_{-i})}{p(\gamma_i = 1 | \gamma_{-i})} = \exp\{\ell_i\} \frac{p(\gamma_i = 0, \gamma_{-i})}{p(\gamma_i = 1, \gamma_{-i})}.$$

Evaluating the joint prior defined in equation (4) for $\gamma_i = 1$ and $\gamma_i = 0$ and simplifying, the formula for $p(\gamma_i = 1|\gamma_{-i}, y)$ is obtained.

Computational Note:

Repeated sampling from the conditional posteriors above in a Markov chain Monte Carlo sampling scheme is reasonably fast. First, only the sum $\sum_{k \in \kappa_i} w_{ij}(1 - 2\gamma_k)$ requires repeated computation each sweep of the sampling scheme. This sum involves only a handful of floating point operations, given the small 8-neighborhood structure used in this paper. Second, the component $\ell_i - \delta_i$ requires calculation only once for each location i on the lattice. Last, some matrix computations, such as $(W'W)^{-1}$, need only be evaluated once in the entire estimation procedure. Overall, we find that when coded efficiently in a low level language these computations can be undertaken quickly. For example, even with a conservative Monte Carlo sampling period of 8000 sweeps, the entire posterior activation map for a single slice can be computed in less than one minute on a Pentium III processor.

Acknowledgements

The authors would like to acknowledge additional financial support through a grant from the German National Science Foundation (DFG), Sonderforschungsbereich 386.

References

- B. S. Everitt and E. T. Bullmore. Mixture model mapping of the brain activation in functional magnetic resonance images. *Human Brain Mapping*, 7(1):1-14, 1999.
- L. Fahrmeir, C. Gössl, and A. Hennerfeind. Spatial smoothing with robust priors in functional MRI. In: *Exploratory Data Analysis in Empirical Research*, Springer (eds. M. Schwaiger and O. Opitz), 50-57, 2002.
- K. J. Friston, D. E. Glaser, R. N. A. Henson, S. Kiebel, C. Phillips, and J. Ashburner. Classical and Bayesian inference in neuroimaging: Applications. *Neuroimage*, 16:484-512, 2002a.
- K. J. Friston, A. P. Holmes, J.-B. Poline, P. J. Grasby, S. C. R. Williams, R. S. J. Frackowiak, and R. Turner. Analysis of fMRI time-series revisited. *Neuroimage*, 2:45-53, 1995.
- K. J. Friston, W. Penny, C. Phillips, S. Kiebel, G. Hinton, and J. Ashburner. Classical and Bayesian inference in neuroimaging: Theory. *Neuroimage*, 16:465-483, 2002b.

- D. Gamerman. Efficient sampling from the posterior distribution in generalized linear models. *Statistics and Computing*, 7:57–68, 1997.
- C. R. Genovese. A Bayesian time-course model for functional magnetic resonance imaging data (with discussion). *J. Am. Stat. Ass.*, 95:691–719, 2000.
- E. George and R. McCulloch. Approaches for Bayesian variable selection. *Statistica Sinica*, 7:339–373, 1997.
- C. Gössl, D. P. Auer, and L. Fahrmeir. Dynamic models in fMRI. *Magn. Reson. Med.*, 43:72–81, 2000.
- C. Gössl, D. P. Auer, and L. Fahrmeir. Bayesian spatio-temporal inference in functional magnetic resonance imaging. *Biometrics*, 57:554–562, June 2001a.
- C. Gössl, L. Fahrmeir, and D. P. Auer. Bayesian modeling of the hemodynamic response function in BOLD fMRI. *NeuroImage*, 14(1 Part 1):140–148, July 2001b.
- N. V. Hartvig and J. L. Jensen. Spatial mixture modeling of fMRI data. *Human Brain Mapping*, 11:233–248, 2000.
- D. Higdon. Auxiliary variable methods for Markov chain Monte Carlo with applications. *Journal of the American Statistical Association*, 93:585–595, 1998.
- S. A. Huettel. The effects of aging upon the hemodynamic response measured by functional MRI. *NeuroImage*, 13(1):161–175, January 2001.
- S. Kiebel and K. J. Friston. Anatomically informed basis functions in multi-subject studies. *Human Brain Mapping*, 16:36–46, 2002.
- S. Kiebel, R. Goebel, and K. J. Friston. Anatomically informed basis functions. *Neuroimage*, 11:656–667, 2000.
- R. Kohn, M. Smith, and D. Chan. Nonparametric regression using linear combinations of basis functions. *Statistics and Computing*, 11:313–322, 2001.
- A. Raftery. Hypothesis testing and model selection. In W. Gilks, S. Richardson, and D. Spiegelhalter, editors, *Markov chain Monte Carlo in Practice*, pages 165–187. Chapman and Hall, London, 1996.
- M. Smith and R. Kohn. Nonparametric regression via Bayesian variable selection. *Journal of Econometrics*, 75:317–344, 1996.
- R. Swendsen and J. Wang. Non-universal critical dynamics in Monte Carlo simulations. *Physical Review Letters*, 58:86–88, 1987.
- M. Tanner. *Tools for Statistical Inference*. 3rd. Ed., Springer-Verlag: New York, 1996.

Prior	Hyperparameter Values	Form of Ising Prior
(a)	$\theta = 0, \delta_i = \log(0.1/0.9)$	$p(\gamma) = \prod_{i=1}^N \{(0.1)^{\gamma_i} (0.9)^{(1-\gamma_i)}\}$
(b)	$\theta = 0, \delta_i = \log\left(\frac{0.1p(g_i=1)}{1-0.1p(g_i=1)}\right)$	$p(\gamma) = \prod_{i=1}^N \{(0.1p(g_i=1))^{\gamma_i} (1-0.1p(g_i=1))^{(1-\gamma_i)}\}$
(c)	$\theta = 0.6, \delta_i = \log(0.1/0.9)$	$p(\gamma) \propto \exp\left\{\sum_{i=1}^N \delta_i \gamma_i + 0.6 \sum_{i \sim j} \omega_{ij} I(\gamma_i = \gamma_j)\right\}$
(d)	$\theta = 0.6, \delta_i = \log\left(\frac{0.1p(g_i=1)}{1-0.1p(g_i=1)}\right)$	$p(\gamma) \propto \exp\left\{\sum_{i=1}^N \delta_i \gamma_i + 0.6 \sum_{i \sim j} \omega_{ij} I(\gamma_i = \gamma_j)\right\}$

Table A.1

The four Ising priors corresponding to the posterior activation and amplitude maps in figures 1 and 2.




Fig. A.1. Posterior activation maps for individual A using the four Ising priors in table 1. The grey matter prior is plotted in grey on the background, with lighter values corresponding to higher prior probabilities of activation and black corresponding to very low, or zero, prior probabilities of activation. Voxels with estimated posterior probabilities of activation $p(\gamma_i = 1|y) > 0.8772$ are classified as active and painted over in white.

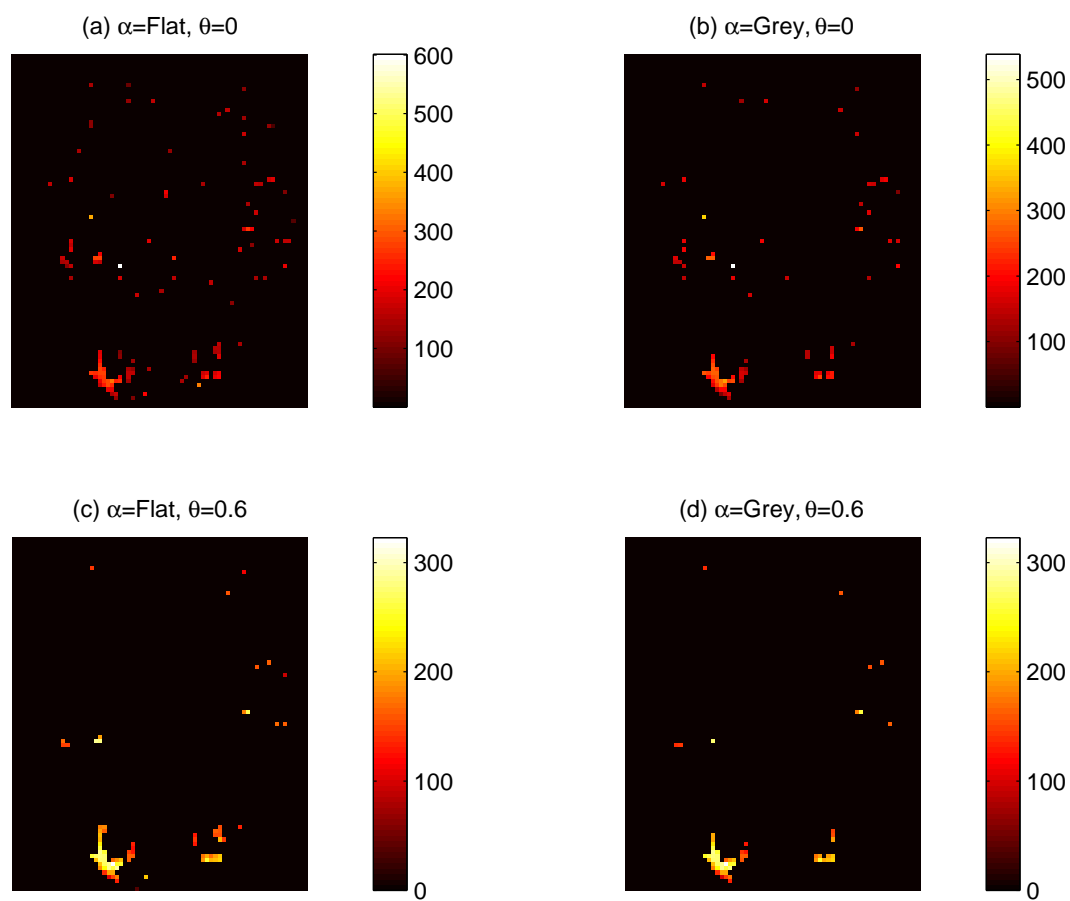


Fig. A.2. Posterior amplitude maps corresponding to figure 1, using the four Ising priors in table 1. The Monte Carlo estimates of the marginal posterior means $E(\beta_i|y)$ are plotted in color, but only for those voxels that are classified as active in figure 1.




Fig. A.3. Posterior activation maps for eight contiguous slices for individual B. The grey matter prior is plotted in grey in the background and voxels classified as active using the posterior probabilities of activation are painted over in white.




Fig. A.4. Activation maps for individual B produced by SPM99 corresponding to those found in figure 3. The method uses spatial smoothing of t-values with a significance level of $p = 0.05$. Areas of statistically significant activation are shown in color.




Fig. A.5. Posterior activation maps for the slice that most corresponds to the LGN for each of the five individuals in our study. Clear activation of the LGN can be observed for individuals A and B, weak activation for individual D and no activation for individuals C and E.




Fig. A.6. Enhanced posterior activation maps for the slices found in figure 5 for individuals E and D. Voxels classified as active are painted over in white. The enhanced prior is plotted in the background in greyscale for prior probabilities less than 0.1 and in shades of red for prior probabilities between 0.1 and 0.5.