

Anatomically Informed Bayesian Spatial Priors for FMRI Analysis

David Abramian, Per Sidén, Hans Knutsson, Mattias Villani and Anders Eklund

The self-archived postprint version of this journal article is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-165856>

N.B.: When citing this work, cite the original publication.

Abramian, D., Sidén, P., Knutsson, H., Villani, M., Eklund, A., (2020), Anatomically Informed Bayesian Spatial Priors for FMRI Analysis, *ISBI 2020*.

<https://doi.org/10.1109/ISBI45749.2020.9098342>

Original publication available at:

<https://doi.org/10.1109/ISBI45749.2020.9098342>

Copyright:

Institute of Electrical and Electronics Engineers (IEEE)

<http://www.ieee.org/index.html> ©2018 IEEE.

Personal use of this material is permitted. However, permission to reprint/republish

this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.



ANATOMICALLY INFORMED BAYESIAN SPATIAL PRIORS FOR FMRI ANALYSIS

David Abramian^{*†}, Per Sidén[‡], Hans Knutsson^{*†}, Mattias Villani^{‡†}, Anders Eklund^{*‡}

^{*} Division of Medical Informatics, Department of Biomedical Engineering

[†] Center for Medical Image Science and Visualization (CMIV)

[‡] Division of Statistics & Machine learning, Department of Computer and Information Science
Linköping University, Linköping, Sweden

[†] Department of Statistics, Stockholm University, Stockholm, Sweden

ABSTRACT

Existing Bayesian spatial priors for functional magnetic resonance imaging (fMRI) data correspond to stationary isotropic smoothing filters that may oversmooth at anatomical boundaries. We propose two anatomically informed Bayesian spatial models for fMRI data with local smoothing in each voxel based on a tensor field estimated from a T_1 -weighted anatomical image. We show that our anatomically informed Bayesian spatial models results in posterior probability maps that follow the anatomical structure.

Index Terms— Bayesian statistics, functional MRI, activation mapping, adaptive smoothing

1. INTRODUCTION

The analysis of functional magnetic resonance imaging (fMRI) data has generally relied on frequentist statistics to perform inferences about brain activity. After fitting a general linear model (GLM) to the individual voxel time series, t-values are calculated for contrasts of interest at every voxel, resulting in t-maps that can be thresholded at an appropriate significance level. Isotropic Gaussian smoothing is the most common way of preprocessing fMRI data, but several adaptive smoothing approaches for detecting brain activity have been proposed [1, 2, 3, 4, 5].

The Bayesian framework for fMRI analysis developed by Penny et al. [6, 7, 8, 9] provides increased flexibility compared to the classical frequentist approach, by allowing the estimation of individual smoothness parameters for each regressor and autoregressive (AR) noise coefficient. Sidén et al. [10, 11] extended this framework with an efficient Markov Chain Monte Carlo (MCMC) implementation for both 2D and 3D inference. In this work we propose two approaches for performing 2D Bayesian spatial modeling in an anatomically-adaptive way within this framework. Our contribution relates to past work involving anatomical priors [12, 13] and spatially adaptive Bayesian models [14, 15] for fMRI data analysis.

2. METHODS

2.1. Bayesian spatial GLM

The most common way to model fMRI data is as a voxel-wise general linear model with serial correlations in the residuals, which are modeled as an AR(p) process. This combined model is referred to as GLM-AR(p). In the case of single subject data, with T volumes, N voxels, K regressors and an AR model of order p , this model is written as

$$\begin{aligned} \mathbf{Y} &= \mathbf{X} \mathbf{W} + \mathbf{E}, \\ \mathbf{E} &= \tilde{\mathbf{E}} \mathbf{R} + \mathbf{Z}, \end{aligned}$$

where \mathbf{Y} is the observation matrix, \mathbf{X} is the design matrix, \mathbf{W} is the regressor matrix and \mathbf{E} is the residual matrix. The residual matrix \mathbf{E} is itself expressed as the product of a lagged prediction error matrix $\tilde{\mathbf{E}}$ with an AR coefficient matrix \mathbf{R} plus a matrix of i.i.d. zero mean Gaussian errors \mathbf{Z} , where $\mathbf{Z}_{:,n} \stackrel{iid}{\sim} N(\mathbf{0}, \lambda_n^{-1} \mathbf{I}_T)$, and λ_n is the noise precision at voxel n . In this Bayesian framework, smoothness of the regression coefficients enters the model through a spatial precision matrix \mathbf{D}_w in their prior distribution, according to

$$\mathbf{W}'_{k,\cdot} \sim \mathcal{N}(\mathbf{0}, \alpha_k^{-1} \mathbf{D}_w^{-1}),$$

where $\mathbf{W}'_{k,\cdot}$ is the transpose of the k -th row of the regression coefficient matrix \mathbf{W} , representing the k -th regression coefficient at every voxel, and α_k is a smoothness hyperparameter for the k -th regressor (since the following also applies to the AR coefficients, we refer to their precision matrices generically as \mathbf{D}). The precision matrix \mathbf{D} encodes the conditional dependencies between every pair of voxels. The smoothness assumption constrains these dependencies to exist only between a voxel and its immediate neighbors, which has the advantage of making this matrix very sparse.

2.2. Uniform graph Laplacian precision matrix

Graphs constitute a natural way of describing relationships between sets of elements. Therefore, the precision matrices

This work was supported by the Swedish Research Council, grant 2017-04889, and by the Center for Industrial Information Technology (CENIIT) at Linköping University

employed by Penny et al. and Sidén et al. take the form of graph Laplacian matrices. In their formulation the precision matrix \mathbf{D} is an unweighted graph Laplacian (UGL), where $d_{i,i}$ equals the number of pixels neighboring pixel i and $d_{i,j} = -1$ if pixels i, j are cardinal neighbors. For inner pixels this corresponds to the Laplacian operator, which in 2D takes the form

$$\mathbf{H}_{\text{UGL}} = \begin{pmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{pmatrix}.$$

Generically, graph Laplacian matrices can be constructed with [16]

$$\mathbf{L} = \mathbf{B} - \mathbf{A},$$

where \mathbf{A} is an *adjacency matrix* whose $a_{i,j}$ element, called a *weight*, represents in this context the strength of the conditional dependence between pixels i and j , and \mathbf{B} is a diagonal *degree matrix* with $b_{i,i} = \sum_j a_{i,j}$, that is, the sum of all the incoming weights to pixel i . Since \mathbf{B} can be constructed from \mathbf{A} , the adjacency matrix is sufficient for generating the graph Laplacian. Under this formulation the same UGL prior can be constructed by considering its dependency structure (neighborhood), which for inner pixels takes the shape

$$\mathbf{N}_{\text{UGL}} = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix},$$

that is, $a_{i,j} = 1$ if pixels i, j are cardinal neighbors. It can be seen that this results in the same precision matrix $\mathbf{D} \equiv \mathbf{L}$.

Due to the uniform neighborhood shape and constant weights between neighbors this type of prior is incapable of encoding anatomical information. Such a prior will not respect anatomical tissue boundaries, mixing, for example, signals from white and gray brain matter. Conversely, in order for the prior to encode relevant anatomical information it is necessary for the dependencies between neighboring pixels to be specified independently at each location.

2.3. The structure tensor

In order to establish pixel relationships that respect anatomical boundaries, we must first estimate the position and orientation of these boundaries. Our main tool for achieving this is the structure tensor, a tensor estimated at every pixel (using quadrature filters along different directions [17, 18]) whose eigenvalues and eigenvectors reveal the degree to which spatial structure is present and the orientation of any such structure, here understood as lines or edges. A detailed explanation of the estimation and interpretation of structure tensors falls outside the scope of this paper, see [19, 17, 20] for a thorough treatment of this topic. For our purposes it is sufficient to point out that the eigenvector of the structure tensor corresponding to the largest eigenvalue is perpendicular to the main structure orientation at the given pixel, while the eigenvalue is proportional to the degree of certainty of this orientation. The




Fig. 1. Local orientation represented using the structure tensor. Red vectors indicate the main local orientation, while green vectors indicate the second (perpendicular) orientation. If no vector is present, there is no orientation information available (e.g. due to uniform intensity).

second eigenvalue indicates the amount of structure in the orientation perpendicular to the first eigenvector. The structure tensor is, therefore, a flexible model capable of representing structure in zero (noise, uniform data), one (line, edge), or two (crossing lines, crossing edges) orientations (for the 2D case).

As fMRI data lacks contrast, the structure tensor is calculated from a registered T_1 -weighted volume. Figure 1 shows the eigenvectors of the structure tensor for a typical brain image. The tensor is small and isotropic in very uniform regions, while next to lines or edges it becomes anisotropic and oriented perpendicularly to the line or edge.

2.4. Spatial model with four orientations

In our first proposed solution we assign to every pixel one of four oriented neighborhood structures according to the local structural orientation in that pixel. Such a model offers limited angular resolution, since only four different orientations can be represented. However, it simplifies the formulation of the possible neighborhoods and provides increased sparsity in the precision matrix \mathbf{D} . We refer to this model as 4DIR. We start by considering four possible orientations: horizontal, vertical, and two diagonals. The corresponding orientation vectors are

$$\mathbf{d}_x = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \mathbf{d}_y = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

$$\mathbf{d}_{xy} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad \mathbf{d}_{-xy} = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

For computational efficiency we want to avoid having to calculate the eigenvectors of the structure tensor. We can find

which of the four orientations is closest to the main tensor orientation at each point by first defining tensors corresponding to each of the four orientations

$$\begin{aligned} \mathbf{T}_x &= \mathbf{d}_x \mathbf{d}_x^T, & \mathbf{T}_y &= \mathbf{d}_y \mathbf{d}_y^T, \\ \mathbf{T}_{xy} &= \mathbf{d}_{xy} \mathbf{d}_{xy}^T, & \mathbf{T}_{-xy} &= \mathbf{d}_{-xy} \mathbf{d}_{-xy}^T, \end{aligned}$$

and then projecting, through an inner product, the structure tensor at each point onto the four orientation tensors. The maximum projection value will correspond to the orientation closest to that of the structure tensors.

Having found at each pixel which of the four orientations is closest to that of the structure tensor, we define neighborhoods (filters) along each of the four orientations

$$\begin{aligned} \mathbf{N}_x &= \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, & \mathbf{N}_y &= \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \\ \mathbf{N}_{xy} &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, & \mathbf{N}_{-xy} &= \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}. \end{aligned}$$

As the structure tensor is aligned across the orientation of lines and edges (see Figure 1), which we are trying to preserve, we associate to each pixel a neighborhood perpendicular to the orientation of the structure at that point (e.g. if \mathbf{d}_x is the orientation for some pixel, then the appropriate neighborhood would be \mathbf{N}_y).

Having assigned a specific neighborhood to each pixel, a new adjacency matrix $\mathbf{A}_{4\text{DIR}}$ can be constructed and used to define a graph Laplacian $\mathbf{D}_{4\text{DIR}}$. However, due to the specific process for calculating the orientation at each position, it cannot be guaranteed that the relationship between a pair of pixels i, j is symmetric, i.e., that $a_{i,j} = a_{j,i}$, which would preclude $\mathbf{D}_{4\text{DIR}}$ from being used as a precision matrix. Symmetry can be enforced by applying the procedure

$$\mathbf{A}_{4\text{DIR}}' = \frac{\mathbf{A}_{4\text{DIR}} + \mathbf{A}_{4\text{DIR}}^T}{2},$$

which effectively sets the value of $a_{i,j}$ and $a_{j,i}$ to the average of the two. This correction has little effect on the orientations encoded in the graph and allows the new matrix $\mathbf{L}'_{4\text{DIR}}$ to be used as a precision matrix $\mathbf{D}_{4\text{DIR}}$.

2.5. Spatial model with arbitrary orientations

The previously presented model is limited in angular resolution, as it can only encode four orientations. In order to overcome this limitation we consider a 3×3 neighborhood around every pixel and determine the weight for each of the 8 possible neighbors by sampling from a continuous function at discrete positions corresponding to the centers of the neighboring pixels. We refer to this model as ANYDIR. The sampled function is

$$a_{i,j} = \mathbf{f}(i, j) = \frac{|\sin(\phi_{\text{pix}_j} - \phi_{\text{tensor}_i})|^\alpha}{r_{\text{pix}_j}^\beta}, \quad \alpha, \beta > 0,$$




Fig. 2. Variables involved in weighting of the neighborhoods for the ANYDIR model. The structure tensor is shown in the middle.

where ϕ_{pix_j} is the angle of the line connecting the central pixel i with neighboring pixel j , ϕ_{tensor_i} is an angle representing the main orientation of the structure tensor at pixel i , r_{pix_j} is the distance between pixels i and j (see Figure 2). Additionally, α and β are adjustable parameters taking non-negative real values, where α controls the width of the distribution around ϕ_{tensor_i} and β penalizes the values in the diagonal neighbors with respect to the horizontal and vertical ones. In this work we use $\alpha = 12$ and $\beta = 5$.

This neighborhood formulation generates neighborhoods that vary continuously with respect to ϕ_{tensor_i} , which allows the representation of arbitrary orientations in the structures around pixels, at the cost of reduced sparsity in the precision matrix. As in the previous case the adjacency matrix $\mathbf{A}_{\text{ANYDIR}}$ is not guaranteed to be symmetric. The same kind of correction can be applied in order to generate a graph Laplacian matrix $\mathbf{L}'_{\text{ANYDIR}}$ suitable for use as a precision matrix $\mathbf{D}_{\text{ANYDIR}}$.

3. RESULTS

We compared the results produced by all three precision matrices in the context of Bayesian fMRI analysis. We used preprocessed data from subject 100307 of the Human Connectome Project [21], specifically the T_1 -weighted volume and the motor task data. The analysis was performed slice-by-slice using the SPM package for Matlab together with the extension developed by Sidén et al. [10]. The only modification to the code was the addition of the proposed precision matrices. Our modified code can be found at (<https://github.com/DavidAbramian/adaptiveBayesianPrior>).

As a preprocessing step, the T_1 -weighted volume was downsampled to match the resolution of the fMRI data. This is necessary, since the precision matrix has to be estimated from the T_1 -weighted volume and later applied to the fMRI data.




Fig. 3. Comparison of neighborhood structures implied by both of the proposed models. Lines represent the orientation of spatial prior dependencies at the given point. Both methods adapt these dependencies at each point in accordance with the anatomical structure. The 4DIR method only makes use of four possible orientations, while with ANYDIR the orientations vary continuously.

Due to the time-consuming nature of the MCMC analysis it was carried out on a single representative axial slice. The Gibbs sampling algorithm was iterated 10,000 times, with 1,000 warmup iterations, and with a thinning factor of 5.

3.1. Anatomical adaptiveness

Figure 3 illustrates the orientation of the pixel neighborhoods generated by the two proposed approaches, which determine the conditional dependencies encoded in prior precision matrix D . In both cases the models have successfully adapted to the anatomical structure given by the T_1 -weighted volume, as spatial prior dependence is placed along lines and edges and not across them. The 4DIR method shows limited angular resolution, as it can only represent neighborhoods in four orientations, while the ANYDIR produces neighborhoods in arbitrary orientations.

3.2. Functional MRI results

Figure 4 shows, for all three models, the regression coefficients obtained for a right hand motor task, as well as posterior probability maps (PPMs) quantifying the probability of the effect of said task exceeding 0.2% of the global mean signal and thresholded at 0.8. The regression coefficients for both of the proposed models clearly reflect anatomical spatial patterns absent from the UGL results. The patterns are highly angular for the 4DIR model as a result of the limited angular resolution, while the ANYDIR model results in more natural curved patterns.

The PPMs from all three methods are similar, showing large activation in the motor cortex and close to the central sulcus, and smaller activations in the somatosensory cortex. However, the activations detected by both of the proposed methods are slightly narrower and extend further along




Fig. 4. Bayesian GLM regression results obtained using MCMC. Left: regression coefficient corresponding to right hand motor task. Right: PPMs for the probability of effect of the right hand motor task exceeding 0.2% of the global mean signal, thresholded at 0.8. Top: UGL model. Middle: 4DIR model. Bottom: ANYDIR model.

anatomical lines, indicating that the priors respect anatomical boundaries.

4. DISCUSSION

We have proposed two new Bayesian spatial priors for fMRI analysis that allow for a locally anisotropic spatial dependence over voxels. These priors were used to encode anatomical structure, resulting in anatomically-adaptive smoothing for fMRI data. The priors can be easily incorporated into the existing framework for Bayesian fMRI analysis, requiring minimal modifications.

While our presentation is centered on 2D priors, both of the proposed approaches can be extended to 3D without requiring significant modifications. The priors can also be improved by incorporating additional information from the structure tensor to, for example, use the UGL prior in areas with uniform intensity. These adaptive spatial models can also be applied in Bayesian frameworks for diffusion MRI data [22].

5. REFERENCES

- [1] Joakim Rydell, Hans Knutsson, and Magnus Borga, "Bilateral filtering of fMRI data," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 6, pp. 891–896, 2008.
- [2] Anders Eklund, Mats Andersson, and Hans Knutsson, "Fast random permutation tests enable objective evaluation of methods for single-subject fMRI analysis," *International journal of biomedical imaging*, vol. 2011, 2011.
- [3] Hamid Behjat, Nora Leonardi, Leif Sörnmo, and Dimitri Van De Ville, "Anatomically-adapted graph wavelets for improved group-level fMRI activation mapping," *NeuroImage*, vol. 123, pp. 185–199, 2015.
- [4] Xiaowei Zhuang, Zhengshi Yang, Tim Curran, Richard Byrd, Rajesh Nandy, and Dietmar Cordes, "A family of locally constrained CCA models for detecting activation patterns in fMRI," *NeuroImage*, vol. 149, pp. 63–84, 2017.
- [5] Gabriele Lohmann, Johannes Stelzer, Eric Lacosse, Vinod J Kumar, Karsten Mueller, Esther Kuehn, Wolfgang Grodd, and Klaus Scheffler, "LISA improves statistical analysis for fMRI," *Nature communications*, vol. 9, no. 1, pp. 4014, 2018.
- [6] Will Penny and Guillaume Flandin, "Bayesian analysis of fMRI data with spatial priors," in *Proceedings of the Joint Statistical Meeting (JSM)*. American Statistical Association. Citeseer, 2005.
- [7] Will Penny, Guillaume Flandin, and Nelson Trujillo-Barreto, "Bayesian comparison of spatially regularised general linear models," *Human brain mapping*, vol. 28, no. 4, pp. 275–293, 2007.
- [8] Lee M Harrison, W Penny, Jean Daunizeau, and Karl J Friston, "Diffusion-based spatial priors for functional magnetic resonance images," *Neuroimage*, vol. 41, no. 2, pp. 408–423, 2008.
- [9] Lee M Harrison, Will Penny, Guillaume Flandin, Christian C Ruff, Nikolaus Weiskopf, and Karl J Friston, "Graph-partitioned spatial priors for functional magnetic resonance images," *NeuroImage*, vol. 43, no. 4, pp. 694–707, 2008.
- [10] Per Sidén, Anders Eklund, David Bolin, and Mattias Villani, "Fast Bayesian whole-brain fMRI analysis with spatial 3D priors," *NeuroImage*, vol. 146, pp. 211–225, 2017.
- [11] Per Sidén, Finn Lindgren, David Bolin, Anders Eklund, and Mattias Villani, "Spatial 3D Matérn priors for fast whole-brain fMRI analysis," *arXiv preprint arXiv:1906.10591*, 2019.
- [12] Wanmei Ou, William M Wells III, and Polina Golland, "Combining spatial priors and anatomical information for fmri detection," *Medical image analysis*, vol. 14, no. 3, pp. 318–331, 2010.
- [13] Ville-Veikko Wettenhovi, Ville Kolehmainen, Joanna Huttunen, Mikko Kettunen, Olli Gröhn, and Marko Vauhkonen, "State estimation with structural priors in fmri," *Journal of Mathematical Imaging and Vision*, vol. 60, no. 2, pp. 174–188, 2018.
- [14] Thomas Vincent, Laurent Risser, and Philippe Ciuciu, "Spatially adaptive mixture modeling for analysis of fmri time series," *IEEE transactions on medical imaging*, vol. 29, no. 4, pp. 1059–1074, 2010.
- [15] Lotfi Chaari, Thomas Vincent, Florence Forbes, Michel Dojat, and Philippe Ciuciu, "Fast joint detection-estimation of evoked brain activity in event-related fmri using a variational approach," *IEEE transactions on Medical Imaging*, vol. 32, no. 5, pp. 821–837, 2012.
- [16] Fan RK Chung, *Spectral graph theory*, American Mathematical Society, 1997.
- [17] Gösta H Granlund and Hans Knutsson, *Signal processing for computer vision*, Kluwer, 1995.
- [18] Hans Knutsson and Mats Andersson, "What's so good about quadrature filters?," in *Proceedings 2003 International Conference on Image Processing (Cat. No. 03CH37429)*. IEEE, 2003, vol. 3, pp. III–61.
- [19] Hans Knutsson, "Representing local structure using tensors," in *Scandinavian Conference on Image Analysis*, 1989.
- [20] Hans Knutsson, Carl-Fredrik Westin, and Mats Andersson, "Representing local structure using tensors II," in *Scandinavian Conference on Image Analysis*. Springer, 2011, pp. 545–556.
- [21] David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, et al., "The WU-Minn human connectome project: an overview," *Neuroimage*, vol. 80, pp. 62–79, 2013.
- [22] Xuan Gu, Per Sidén, Bertil Wegmann, Anders Eklund, Mattias Villani, and Hans Knutsson, "Bayesian diffusion tensor estimation with spatial priors," in *International Conference on Computer Analysis of Images and Patterns*. Springer, 2017, pp. 372–383.