

Author Manuscript

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the [Version of Record](#). Please cite this article as [doi: 10.1111/rssc.12320](https://doi.org/10.1111/rssc.12320)

This article is protected by copyright. All rights reserved

Bayesian Analysis of fMRI data with Spatially-Varying Autoregressive Orders

Ming Teng

Department of Biostatistics, University of Michigan, Ann Arbor, Mi., USA.

Farouk S. Nathoo

*Department of Mathematics and Statistics, University of Victoria,
Victoria, British Columbia, Canada.*

Timothy D. Johnson

Department of Biostatistics, University of Michigan, Ann Arbor, Mi., USA.

E-mail: tdjdj@umich.edu

Summary.

Statistical modeling of fMRI data is challenging as the data are both spatially and temporally correlated. Spatially, measurements are taken at thousands of contiguous regions, called voxels, and temporally measurements are taken at hundreds of time points at each voxel. Recent advances in Bayesian hierarchical modeling have addressed the challenges of spatiotemporal structure in fMRI data with models incorporating both spatial and temporal priors for signal and noise. While there has been extensive research on modeling the fMRI signal (i. e., the convolution of the experimental design with the functional choice for the hemodynamic response function) and its spatial variability, less attention has been paid to realistic modeling of the temporal dependence that typically exists within the fMRI noise, where a low order autoregressive process is typically adopted. Furthermore, the AR order is held constant across voxels (e.g. AR(1) at each voxel). Motivated by an event-related fMRI experiment, we propose a novel hierarchical Bayesian model with automatic selection of the autoregressive orders of the noise process that vary spatially over the brain. With simulation studies we show that our model is more statistically efficient and apply it to our motivating example.

1. Introduction

In the analysis of functional magnetic resonance imaging (fMRI) data a key challenge is dealing with spatial and temporal correlation. Temporal correlation can arise from many sources, including scanner drift at very low frequencies, slow vascular/metabolic oscillations that are of moderate to low frequency, and other sources of noise such as breathing and heartbeat. Simply ignoring these sources of temporal correlation may lead to increased false positive discoveries (Makni et al., 2006). To deal with these issues, a variety of approaches have been proposed. One commonly used approach, namely prewhitening, works by estimating the temporal correlation and then decorrelating the noise using the estimates (Bullmore et al., 1996; Locascio et al., 1997). Besides these stationary time series models, non-stationary $1/f$ models have also been proposed (Zarahn et al., 1997; Bullmore et al., 2004). According to Friston et al. (2000), prewhitening

can produce an extraneous source of bias. Alternatively, a band-pass filtering technique known as precoloring can be applied to the data first, followed by statistical modeling that deals with the autocorrelation in the colored data. For a review and discussion of these approaches the reader is referred to Woolrich et al. (2001). While high-pass filtering has proven to be beneficial in increasing the power of the statistical analysis, the low-pass filtering involved in coloring is considered controversial in that it tends to add autocorrelation into the data (Skudlarski et al., 1999; Della-Maggiore et al., 2002).

While accurate temporal modeling is important for estimation of the fMRI signal and, especially, its standard error, traditional approaches apply a temporal model at each voxel, independently. That is, they ignore spatial correlation. More specifically, this mass univariate approach, considered to be the classical approach to the analysis of fMRI data, includes a smoothing step involving a spatial Gaussian filter that is applied to the data first (Friston et al., 1995), followed by model estimation at each voxel, and then statistical inference is based on random field theory (Worsley and Friston, 1995) which is applied to adjust for multiplicity in the spatial domain. While this approach remains the most common approach for analyzing fMRI data it has been criticized on a number of grounds. For example, the Gaussian kernel that is used to smooth the data has to be pre-specified and introduces artificial spatial correlation into the data, similar to how precoloring adds artificial temporal correlation. In addition, this approach does not directly account for spatial correlation in the model.

Partly as a result of these criticisms, Bayesian models with spatially structured priors have been proposed that allow for the calculation of posterior probability maps (PPM) for activation. This Bayesian approach to inference is based on an explicit spatial model and does not require smoothing the data with a Gaussian kernel nor does it require the use of random-field theory-based adjustments for multiplicity. A variety of spatiotemporal Bayesian models have been proposed. One model that is widely used and implemented within the SPM software is the GLM-AR (general linear model, auto-regressive) model (Penny et al., 2003, 2005, 2007). The GLM-AR model assumes that the data can be decomposed into two sources of variability. The first source is the product of the design matrix for the fMRI experiment convolved with a hemodynamic response function (HRF) and experimental factors, and the second source represents temporally correlated noise that is modeled using a low-order AR structure. In addition, the regression coefficients and the autoregressive (AR) coefficients vary across space and are assigned spatial smoothing priors. Gössl et al. (2001) proposed a model where the data are decomposed into three sources, a spatial stimulus, a deterministic trend, and a white noise process. However, this modeling approach may not account for some higher frequency stochastic noise components. Woolrich et al. (2004b) assumed that the temporal noise arises from both large scale and small scale variation, and built a space-time simultaneously auto-regressive model that accounts for both scales of variation. Methods focusing on spatial variable selection have also been proposed (see, e.g., Bezener et al. (2016), Lee et al. (2014), Musgrove et al. (2016)); while Kim et al. (2010) proposed a mixture of experts model to represent spatial activation clusters. While these models have a number of different characteristics which make the approaches unique, most of them commonly assume a homogeneous, low order AR or ARMA process for the temporal noise. By homogeneous, we mean that the order of the AR or ARMA process is assumed constant

across all voxels. This assumption is also made in Penny et al. (2003). However, as we demonstrate using a simple empirical example in the next section, this homogeneous AR order assumption may be violated in real fMRI data.

Instead of formulating the model at each voxel and then adopting spatial smoothing priors for parameters across voxels, another branch of research is based on vector autoregressive (VAR) processes, see Harrison et al. (2003). This approach allows for time-lagged dependence across voxels and spatial-temporal interaction but fitting these models across a large number of voxels is computationally intractable and low-rank approximations have to be used. These models are also useful for studying effective brain connectivity, where one time course is used to predict the other (Castruccio et al., 2016; Chang and Glover, 2010). Another line of work chooses to model the temporal noise as a $1/f$ long memory process, and applies discrete wavelet transforms (DWT) towards fitting the model (see, e.g., Jeong et al. (2013); Bullmore et al. (2004); Fadili and Bullmore (2002); Meyer (2003)). While this approach seems promising, our focus in this paper will be with modeling short term memory using the classical AR process and spatial priors. The reason we choose to work with the AR process is because of its mathematical amenability and simplicity, and its wide use in different areas of science. A novel aspect of our work is that we allow the data to determine the order of the AR process at each voxel, borrowing strength from neighboring voxels. The AR order is determined using ideas from Bayesian spatial variable selection as we expect neighboring voxels to have similar temporal structure.

Computation is an important issue when considering Bayesian spatial-temporal models for fMRI data. While the main focus of this paper lies with the development of a new model, another aspect of this work is the comparison of fully Bayesian and approximate Bayesian computation methods. Due to the computational burden associated with fitting models to high-dimensional brain imaging data, approximate Bayesian methods have received considerable attention in the neuroimaging literature. One such method is the variational Bayes (VB) inference (Penny et al., 2003, 2007; Woolrich et al., 2004a). As there are currently no theoretical results quantifying the accuracy of VB methods (in contrast to MCMC which is justified by large sample theory of stationary Markov chains), the evaluation of VB has to be performed on a case-by-case basis. In some cases, the performance of VB can be quite good and in other cases it can be quite poor. In addition to the implementation of our new model based on a suitably designed MCMC sampler, we also develop an MCMC algorithm to sample from the posterior of the original GLM-AR model. We then compare our model with both the VB implementation of the GLM-AR model (using SPM code) and our MCMC implementation of the GLM-AR model. Our studies indicate that under a low signal-to-noise (SNR) ratio the accuracy of MCMC outperforms VB according to several criteria.

1.1. Motivating example

Our motivating example comes from a single subject in a fMRI experiment examining a face-repetition stimulus. The experiment involves the presentation of either famous faces (F) or non-famous faces (N) with each type of face presented two times. Convolving this experiment design with the canonical hemodynamic response function and its time and dispersion derivatives leads to a design matrix with twelve columns plus one extra

column for an intercept term in the regression model. After performing the necessary pre-processing steps as described in Penny et al. (2005), we fit a simple linear regression at each of the voxels. After obtaining the residuals from each voxel-specific fit, we fit an AR process up to order 12 for each voxel using the “ar” function in R. AIC is used to estimate the true AR orders. Figure 1 displays a pictorial representation of the results.

Figure 1 shows considerable variability in the estimated AR order across voxels. While most of the estimated AR orders are 4 or less, higher orders up to 12 are selected at some of the voxels. Furthermore, these estimated AR orders tend to show some extent of spatial clustering. If, as is often done, we simply model the data using a homogeneous low-order AR process, then the voxels with higher AR orders will be incorrectly modeled, and this inaccuracy in the modeling of temporal noise will have an impact on the inference on the covariates of interest (via underestimated standard errors), resulting in potentially false inferences about brain activation. To address this issue, we propose a spatially varying AR order (SVARO) model, where the AR orders vary spatially across the brain. This is made possible by adopting a spike-and-slab prior with a stochastic search variable selection scheme. Spatial clustering of AR orders is incorporated by imposing an Ising prior (Ising, 1925) as the latent indicator for the spike and slab prior. We update the latent indicators using the Swendsen-Wang algorithm (Swendsen and Wang, 1987) alternating with Gibbs sampling in our MCMC algorithm. To avoid the phase transition problem associated with the Ising model, we derive theoretical bounds as in Li et al. (2015) and use these bounds to prevent critical slowing of the algorithm. We compare our model with the GLM-AR model of Penny et al. (2007) (implemented under two schemes: our self written MCMC sampler and the VB algorithm available in the SPM software) in terms of mean squared error (MSE) and sensitivity. We conduct these comparisons using two simulation studies and then compare results on the motivating data set.

The rest of the paper is organized as follows: In Section 2 we present our model and MCMC sampling scheme. We present results from our simulation studies in Section 3, followed by the analysis of the face-repetition data set in Section 4. Lastly, we provide a discussion and outline some possible directions for future work in Section 5.

2. The Spatially-Varying AR Order Model

First, we will define the model likelihood and then specify the spatial and temporal priors. Then we discuss our posterior sampling scheme, include the construction of bounds for the hyperpriors that we use for the Ising priors used for variable selection, lastly we discuss inference based on posterior probability maps.

2.1. Model likelihood

We let P denote the maximum allowable order of the AR process at every voxel and let K denote the number of regression coefficients at each voxel. Using similar notation as in Teng et al. (2016), for voxel n ($n = 1, \dots, N$), we let \mathbf{y}_n denote the observed time series of length T . For simplicity, our model is specified conditional on the first P observations at each voxel so that the likelihood function is constructed based on the model for the remaining $T - P$ observations in the time series. We let \mathbf{X} denote the $(T - P) \times K$ design matrix, \mathbf{w}_n denote the K -dimensional vector of regression coefficients at voxel n ,

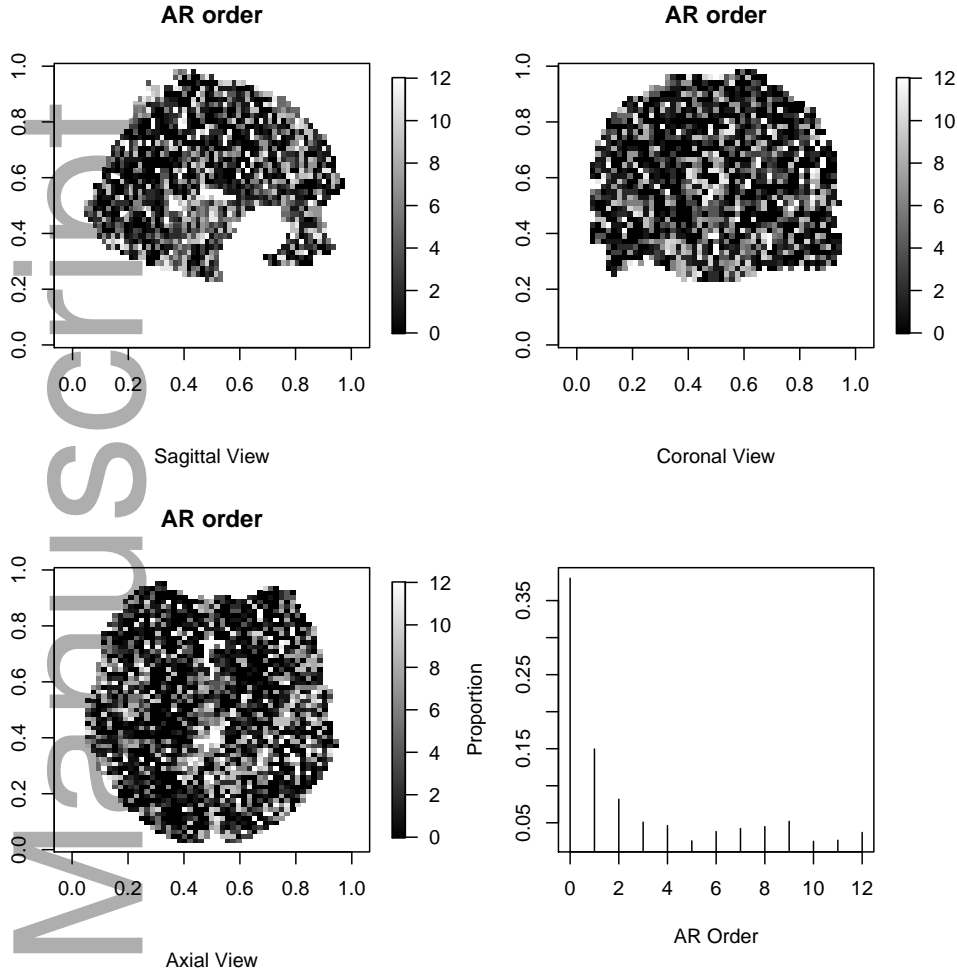


Fig. 1. AIC estimated AR orders. The upper AR order bound is set to 12.

and \mathbf{e}_n denotes the corresponding error term. Define the vector $\mathbf{y}_n \equiv \mathbf{y}_{1:T,n}$, the entire time series observed at voxel n . The hierarchical model is specified in several stages. The first stage is a general linear model:

$$\mathbf{y}_{P+1:T,n} = \mathbf{X}\mathbf{w}_n + \mathbf{e}_n, \quad (1)$$

where we emphasize again the implicit conditioning on $y_{1:P,n}$ ($n = 1, \dots, N$). Let $\tilde{\mathbf{E}}_n$ denote the embedded error (or lagged prediction) matrix of dimension $(T-P) \times P$, with t, p element $(\mathbf{y}_{P+1:T,n} - \mathbf{X}\mathbf{w}_n)_{[t-p]}$ where the notation $[i]$ denotes the i^{th} index of the vector. Let $\mathbf{z}_n \equiv \mathbf{z}_{P+1:T,n}$ denotes a vector of i.i.d mean-zero Gaussian random variables with precision λ_n . The second stage is then an AR model at each voxel:

$$\mathbf{e}_n = \tilde{\mathbf{E}}_n \mathbf{a}_n + \mathbf{z}_n \quad (2)$$

where \mathbf{a}_n is a vector of AR coefficients for the time series at voxel n .

Letting c denote a constant term, the log-likelihood for voxel n , is

$$l_n = -\frac{\lambda_n}{2} \sum_{t=P+1}^T [(y_{tn} - \mathbf{x}_t \mathbf{w}_n) - \tilde{\mathbf{e}}_{tn} \mathbf{a}_n]^2 + \frac{T-P}{2} \log \lambda_n + c, \quad (3)$$

where \mathbf{x}_t is the $(t-P)$ th row of the design matrix \mathbf{X} and $\tilde{\mathbf{e}}_{tn}$ is the $(t-P)$ th row of $\tilde{\mathbf{E}}_n$.

Summing this log-likelihood over the number of voxels, n , we obtain the overall log-likelihood:

$$l = \sum_{n=1}^N \left\{ -\frac{\lambda_n}{2} \sum_{t=P+1}^T [(y_{tn} - \mathbf{x}_t \mathbf{w}_n) - \tilde{\mathbf{e}}_{tn} \mathbf{a}_n]^2 + \frac{T-P}{2} \log \lambda_n + c. \right\} \quad (4)$$

2.2. Spatial prior

At the next level of the model we specify a spatial smoothing prior for the regression coefficients $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_N)$, a $K \times N$ matrix, with k th row $\mathbf{W}_{k,\cdot}$. Following Penny et al. (2005), we assume that the prior for \mathbf{W} takes the form

$$\pi(\mathbf{W}) = \prod_{k=1}^K \pi(\mathbf{W}_{k,\cdot}) \quad (5)$$

$$\mathbf{W}_{k,\cdot} \sim \mathbf{N}(\mathbf{0}, \alpha_k^{-1} (\mathbf{S}^T \mathbf{S})^{-1}), \quad k = 1, \dots, K. \quad (6)$$

A priori the regression coefficients within voxels are independent (5) whilst spatially, the k th, $k = 1, \dots, K$, regression coefficients (across voxels) are modeled dependently through an N -dimensional multivariate normal distribution (6). Here \mathbf{S} is known as a Laplacian matrix. The n th diagonal term of this matrix is equal to the corresponding number of first order neighbors of the voxel n . All off-diagonal terms are zero except for -1 in off-diagonal elements (n, j) and (j, n) if voxel j is a neighbor of voxel n , for $n = 1, \dots, N$. This form for the prior accommodates spatial smoothing while also being sparse and convenient to work with computationally. Also, note that each regression coefficient is smoothed independently of one another and this smoothness is controlled by α_k . In the SPM12 software this prior is referred to as the low resolution tomography (LORETA) prior. Ultimately, what is of primary interest in studies of brain activation is a posterior probability of some function of these regression coefficients, and this posterior probability is computed at each voxel to produce posterior probability maps (PPM, see Penny et al. (2005)). The precision parameter in (6), α_k , is assigned a conditionally conjugate hyper-prior distribution:

$$\alpha_k \stackrel{iid}{\sim} \text{Gamma}(q_1, q_2) \quad (k = 1, \dots, K). \quad (7)$$

2.3. Temporal prior

The key difference between our model and the model of Penny et al. (2007) lies in our modeling of the temporal noise. Rather than assuming AR orders are homogeneous across the brain (we refer the readers to Teng et al. (2016) and Penny et al. (2007) for

model details), we allow for variability in the order of the AR processes across voxels. In addition, we adopt a spatial prior for this variability under the assumption that the order of the AR processes of neighboring voxels will be similar and that their coefficients will also be similar. Specifically, for each voxel n and AR order p , $p = 1, \dots, P$, we assign the latent indicator variable γ_{pn} to the p^{th} AR coefficient a_{pn} , such that given γ_{pn} ($p = 1 \dots P, n = 1 \dots N$), a_{pn} will be conditionally independent. γ_{pn} will take value 1 if the p th AR coefficient is non-zero at voxel n and will take the value 0 otherwise. Conditional on γ_{pn} , a_{pn} will either have a normal distribution or unit mass at 0. This is commonly referred to as the spike-and-slab prior (George and McCulloch, 1993; Mitchell and Beauchamp, 1988), though we note that our formulation is a *spatial* spike-and-slab prior and that this prior is assigned to the coefficients of the AR process governing the temporal noise.

$$\begin{aligned}\pi(\mathbf{a} \mid \boldsymbol{\gamma}) &= \prod_n \prod_p \pi(\mathbf{a}_{pn} \mid \gamma_{pn}) \\ \pi(a_{pn} \mid \gamma_{pn}) &= \gamma_{pn} \phi(a_{pn}; 0, \tau_p^{-1}) + (1 - \gamma_{pn}) I_0(a_{pn})\end{aligned}$$

Here, $\phi(\cdot; a, b)$ is the probability density function (pdf) of a normal distribution with mean a and variance b and $I_0(\cdot)$ is the indicator function that its argument equals 0, and where γ_{pn} is the binary indicator. τ_p is the precision of the normal component and is again given a Gamma prior $\tau_p \sim \text{Gamma}(r_1, r_2)$.

The advantages of introducing such a prior are three fold: First, the coefficients in the AR process at each voxel that lack support from the data can be effectively removed from the model. This allows us to infer which AR coefficients are present in which voxels. Second, the number of voxels with large AR orders is non-zero but expected to be small, which is an aspect of this prior that can be controlled by tuning the hyper-parameters. Third, for some of the voxels it is possible that some AR coefficients are zero while others are non-zero. For example if the AR order at a particular voxel is 5, the second and third coefficient may be zero while the fourth and fifth are non-zero. The proposed model is flexible enough to allow for this behaviour, since we have a total of P independent Ising processes, one for each possible $p \in \{1, \dots, P\}$. Note that this allows for voxels to have independent temporal error structures when $\gamma_{pn} = 0$ for all p .

There are of course other model selection techniques that could have been considered. For example a type of Bayesian lasso could have been used as an alternative to the spike-and-slab prior. Wang et al. (2007) have applied the lasso to the selection of AR processes, and for Bayesian lasso we refer to Schmidt and Makalic (2013). A recent alternative prior known as the non-local prior for variable selection has been proposed by Johnson and Rossell (2012) that has desirable consistency properties and yields smaller prediction errors in large sample settings. A review of Bayesian priors that can be employed for model selection is presented in O'Hara et al. (2009).

We assume that the indicator processes are independent across the set $\{1, \dots, P\}$: $\pi(\boldsymbol{\gamma}) = \prod_{p=1}^P \pi(\boldsymbol{\gamma}_p)$, where $\boldsymbol{\gamma}_p = (\gamma_{p1}, \dots, \gamma_{pN})^T$. The simplest variable selection model would assume γ_{pn} follows a Bernoulli distribution (George and McCulloch, 1993). Here, in order to borrow information across neighbors, we choose to use the Ising prior (Ising,

1925) independently for each $p = 1 \dots P$.

$$\pi(\boldsymbol{\gamma}_p) \propto \exp \left(\beta_{0p} \sum_n \gamma_{pn} + \beta_{1p} \sum_{n_1 \sim n_2} I(\gamma_{pn_1} = \gamma_{pn_2}) \right), \quad (8)$$

where β_{0p} and β_{1p} are two hyper-parameters controlling the sparsity and smoothness of the binary latent field, respectively. Larger values of β_{0p} result in less sparsity and larger values of β_{1p} result in more smoothness. One issue with the Ising model that requires some care is the choice of hyper-parameters. When β_{1p} take values near the phase transition boundary, the mixing of an MCMC sampler will suffer from a critical slow down (Stanley et al., 1987). To avoid the phase transition boundary, we adopt an analytical approach similar to Li et al. (2015) to quantify the value for the bounds of both β_{0p} and β_{1p} . An outline of the bound derivation is given below in Subsection 2.5.

2.4. Posterior sampling scheme

Most parameter updates related to the posterior sampling of our model can be accomplished via Gibbs sampling. One exception is the update to the latent indicators γ_p , $p = 1, \dots, P$. For each γ_p we use the Swendsen-Wang algorithm alternating with Gibbs sampling (Johnson et al., 2013). This strategy, that is, mixing Swendsen-Wang updates with Gibbs updates for γ_p has proven successful in improving the mixing of the Markov chain sampler and results in faster block updates in various studies (Higdon, 1998).

2.5. Bound construction for the Ising Model Hyper-Parameters

The hyper-parameters in the Ising priors play a vital role in posterior estimation. Without careful selection, we are faced with mixing issues associated with “phase-transition” (Stanley et al., 1987). There are various approaches to sampling such hyper-parameters, Johnson et al. (2013) estimated them using path sampling (Gelman and Meng, 1998), Shu et al. (2015) proposed a Monte Carlo EM algorithm to obtain a point estimate of the hyper-parameters, but these procedures are too time consuming for fMRI data and our model, considering that we need to estimate P independent Ising fields and thousands of voxels. Smith and Fahrmeir (2007) proposed to update the hyper-parameters and binary indicators together, but this approach still suffers from potential possibility of sampling over the phase transition boundary. Here, we adopt an approach similar to that considered in Li et al. (2015) and construct some theoretical bounds to prevent phase transition. The resulting hyper-parameter values are then chosen as fixed values within the estimated bounds. This procedure turns out to work well in our analysis and studies.

To construct the bounds, we first write out the posterior conditional density with

respect to γ_p ,

$$\begin{aligned} \pi(\gamma_p | \cdot) \propto & \exp \left[\beta_{0p} \sum_{n=1}^N \gamma_{pn} + \beta_{1p} \sum_{n_1 \sim n_2} \mathbf{I}(\gamma_{pn_1} = \gamma_{pn_2}) \right. \\ & \left. + \sum_{n=1}^N \sum_{t=1}^T \frac{-\lambda_n}{2} \left(e_{tn} - \sum_{j=1}^P \tilde{e}_{tnj} a_{jn} \right)^2 \right] \end{aligned} \quad (9)$$

In our model where multiple AR orders exists across space, it is natural to assume that: 1) there are relatively few time series with large AR order; and 2) $\pi(\gamma_p \neq \mathbf{0} | \cdot) > \pi(\gamma_p = \mathbf{0} | \cdot)$ for small p .

On a 3D lattice, selected voxels give rise to the greatest number of neighboring pairs when they form a cube. Therefore, following Li et al. (2015), the approximate location of the selected voxels are obtained using cubes. Let π_p denote the those voxels whose p th AR coefficient is non-zero. Let $V_p = \lceil (\pi_p N)^{1/3} \rceil$ where $\lceil x \rceil$ is the greatest integer less than x . Then for a cube containing V_p^3 voxels, there are $3V_p^2(V_p - 1)$ neighboring pairs. Based on this we derive

$$\beta_{0p} \sum_{n=1}^N \gamma_{pn} + \beta_{1p} \sum_{n_1 \sim n_2} \mathbf{I}(\gamma_{pn_1} = \gamma_{pn_2}) = \beta_{0p} V_p^3 + 3\beta_{1p} V_p^2 (V_p - 1), \quad (10)$$

for the p th latent Ising image. According to assumption 1) for large p (typically $p > 8$) $\beta_{0p} + 3\beta_{1p} < 0$. According to assumption 2) for small p (typically $p < 4$),

$$\begin{aligned} \sum_{n=1}^N \sum_{t=1}^T -0.5\lambda_n \left(e_{tn} - \sum_{j=1, j \neq p}^P \tilde{e}_{tnj} a_{jn} \right)^2 & \leq \sum_{n=1}^N \sum_{t=1}^T -0.5\lambda_n \left(e_{tn} - \sum_{p=1}^P \tilde{e}_{tnj} a_{jn} \right)^2 \\ & + \beta_{0p} \sum_{n=1}^N \gamma_{pn} + \beta_{1p} \mathbf{I}(\gamma_{pn_1} = \gamma_{pn_2}) \end{aligned} \quad (11)$$

Reorganizing this by moving the first term on right-hand side to left produces:

$$\begin{aligned} \sum_{n=1}^N \sum_{t=1}^T -0.5\lambda_n \left[\left(e_{tn} - \sum_{j=1, j \neq p}^P \tilde{e}_{tnj} a_{jn} \right)^2 - \left(e_{tn} - \sum_{j=1}^P \tilde{e}_{tnj} a_{jn} \right)^2 \right] \\ \leq \beta_{0p} \sum_{n=1}^N \gamma_{pn} + \beta_{1p} \mathbf{I}(\gamma_{pn_1} = \gamma_{pn_2}). \end{aligned} \quad (12)$$

The two terms in the square brackets on the left side can be considered as one with and without $\tilde{e}_{tnp} a_{pn}$. Thus, it can be approximately considered as the residual sum of squares of an AR regression when the p th coefficient a_{pn} is included/excluded in the model. Let R_{pn}^2 denote this coefficient of determination, then we have

$$\beta_{0p} \sum_{n=1}^N \gamma_{pn} + \beta_{1p} \sum_{n_1 \sim n_2} \mathbf{I}(\gamma_{pn_1} = \gamma_{pn_2}) \geq -0.5 \sum_{n=1}^N \sum_{t=1}^T \frac{R_{pn}^2}{1 - R_{pn}^2}. \quad (13)$$

Combined with Equation 10, we have

$$\beta_{0p}V_p^3 + 3\beta_{1p}V_p^2(V_p - 1) \geq -0.5\pi_pNT \frac{R_{pn}^2}{1 - R_{pn}^2}. \quad (14)$$

As an example, let the number of voxels equal 56526, that of the real data we analyze below. Suppose, for example, for the p th AR coefficient we set $\pi_p = 0.1$ and that 5% of the total variation can be explained by the inclusion of the p th AR coefficient, a_{pn} . From Equation 14 we then have the bound: $\beta_{0p} + 2.83\beta_{1p} \geq -9.26$.

Note that the inequality above just gives a range values for the hyperparameters, rather than providing the values directly. In practice, the exact values of hyperparameters are largely determined by the researcher, which should be combined with one's prior experience and an initial analysis of the data. We suggest obtaining such values based on some exploratory ad-hoc approaches, e.g., a linear regression at each voxel followed by fitting an AR process. Then the the estimated AR orders can be used as a reference when determining the hyper-parameters in the Ising model. This method has turned out to work well empirically as we demonstrate in Section 4.

2.6. Posterior probability maps (PPM)

A primary emphasis on fMRI data analysis is inference for activation. For Bayesian modeling of fMRI data this is typically achieved via PPM. Let $\mathbf{c}^T \mathbf{w}_n$ denote a contrast of interest of the regression coefficients. A PPM is a map of the posterior probability of activation for each voxel: $\Pr(\mathbf{c}^T \mathbf{w}_n > \delta_e | \mathbf{y})$. Here δ_e is a pre-specified activation threshold, for example, a value that corresponds to 1% of the global mean value. Thus, PPM looks at the probability of the contrast $\mathbf{c}^T \mathbf{w}_n$ being greater than activation threshold δ_e , given the data.

To formally determine activation in the brain, one can look at a thresholded PPM. This is obtained by exerting a second threshold, namely a probability threshold δ_p , onto the original PPM. Thus, a voxel is considered activate if $\Pr(\mathbf{c}^T \mathbf{w}_n > \delta_e | \mathbf{y}) > \delta_p$. This δ_p reflects the confidence of the inference and usually takes a value above 0.9 (e.g. 0.95 or 0.99). This process discretizes the PPM into null and activate voxels and is commonly used in summarizing a Bayesian analysis for brain activation.

3. Simulation Study

To evaluate the performance of our model, we make comparisons with the standard GLM-AR spatial model. One implementation of this model that we make comparisons to is the Variational Bayes (VB) method available in the SPM12 software. Another implementation is our self-written MCMC sampler for the same model. Although the accuracy of VB has been verified in a setting with high signal-to-noise ratio (SNR) by Teng et al. (2016), under low SNR, MCMC outperforms VB according to certain metrics. This will be illustrated presently. Henceforth, we will refer to the VB implementation and MCMC implementation of the GLM-AR model as PVB and PMCMC, respectively.

3.1. Simulation design

Our design matrix consists of two columns ($K = 2$); the first column is the experimental design (fashioned after the face-repetition data set) convolved with the canonical HRF, and the second column is the intercept. The parameters of interest, corresponding to the experimental design, (one at each voxel) are generated under a mean zero multivariate normal distribution: $\mathbf{W}_{1,\cdot} \sim N(0, (10\mathbf{S}^T\mathbf{S})^{-1/2})$, while the intercepts are generated under a mean 100 multivariate normal distribution: $\mathbf{W}_{2,\cdot} \sim N(100, (10\mathbf{S}^T\mathbf{S}^T)^{-1/2})$. The model noise will have a precision of $\lambda_n = 0.1$ ($n = 1 \dots N$). This corresponds to a fairly small signal-to-noise (SNR) ratio, where the temporal noise will play a greater role in the data. In the following, we will carry out two simulations. In the first case the data will be simulated under our model, and in the second case the data will be generated under the standard spatial GLM-AR model. In these two simulations, we investigate the estimation accuracy of the slopes ($\mathbf{W}_{1,dot}$), intercepts ($\mathbf{W}_{2,dot}$), and autoregressive coefficients (\mathbf{a}_p $p = 1 \dots P$), and we also examine if the difference in inference for these coefficients will lead to a possible difference in the final inference on brain activation. All simulations are based on 100 replicate data sets, and we perform the simulations on a 2-dimensional axial slice of the brain.

3.2. Simulation 1

The purpose of this first simulation study is two-fold. First, since we are generating the data under our model, we are testing our model under an ideal situation to determine its statistical properties and to ensure the algorithm is coded correctly. Second, we compare our model to the standard spatial GLM-AR model under the situation where the AR order and coefficients vary spatially as we have seen empirically (see Figure 1). For the GLM-AR model we use the commonly used assumption that the AR order is small and set it to one. Here we simulate the AR parameters under our SVARO model, that is from the Ising priors. When simulating data we assume that the largest possible AR order at each voxel is 8, however it could be as small as 0. The precision parameters are set as $\tau_p = 400$ ($p = 1 \dots P$). For simplicity, we assume that all AR orders are generated spatially according to the same values for the hyperparameters of the Ising model, i.e., $\beta_{0p} = -0.2$ and $\beta_{1p} = 0.3$. The AR order in PMCMC and PVB are set to $P = 1$ as is standard practice. We note here that the GLM-AR model is misspecified and we expect that its performance will suffer regardless of which posterior sampling algorithm is used. Note that we could have set the AR order to 8 for the GLM-AR model, but this is never done in practice and assumes knowledge of the true maximum value. For our SVARO model, we allow the order to be as large as 12 and let the algorithm determine spatially which of the twelve AR coefficients are set to zero and which ones enter the model via the variable selection mechanism.

We compare SVARO with PVB and PMCMC in estimating the 1st AR coefficient. As shown in Figure 2, SVARO shows little error compared with the truth, indicating that our model has captured the autoregressive parameter quite well. In contrast, PMCMC and PVB exhibit more bias, indicating a lack of fit for the temporal noise. Note that we are only displaying the SVARO estimate for the 1st AR coefficient for simplicity and direct comparison, the other AR coefficients are well-estimated as well.

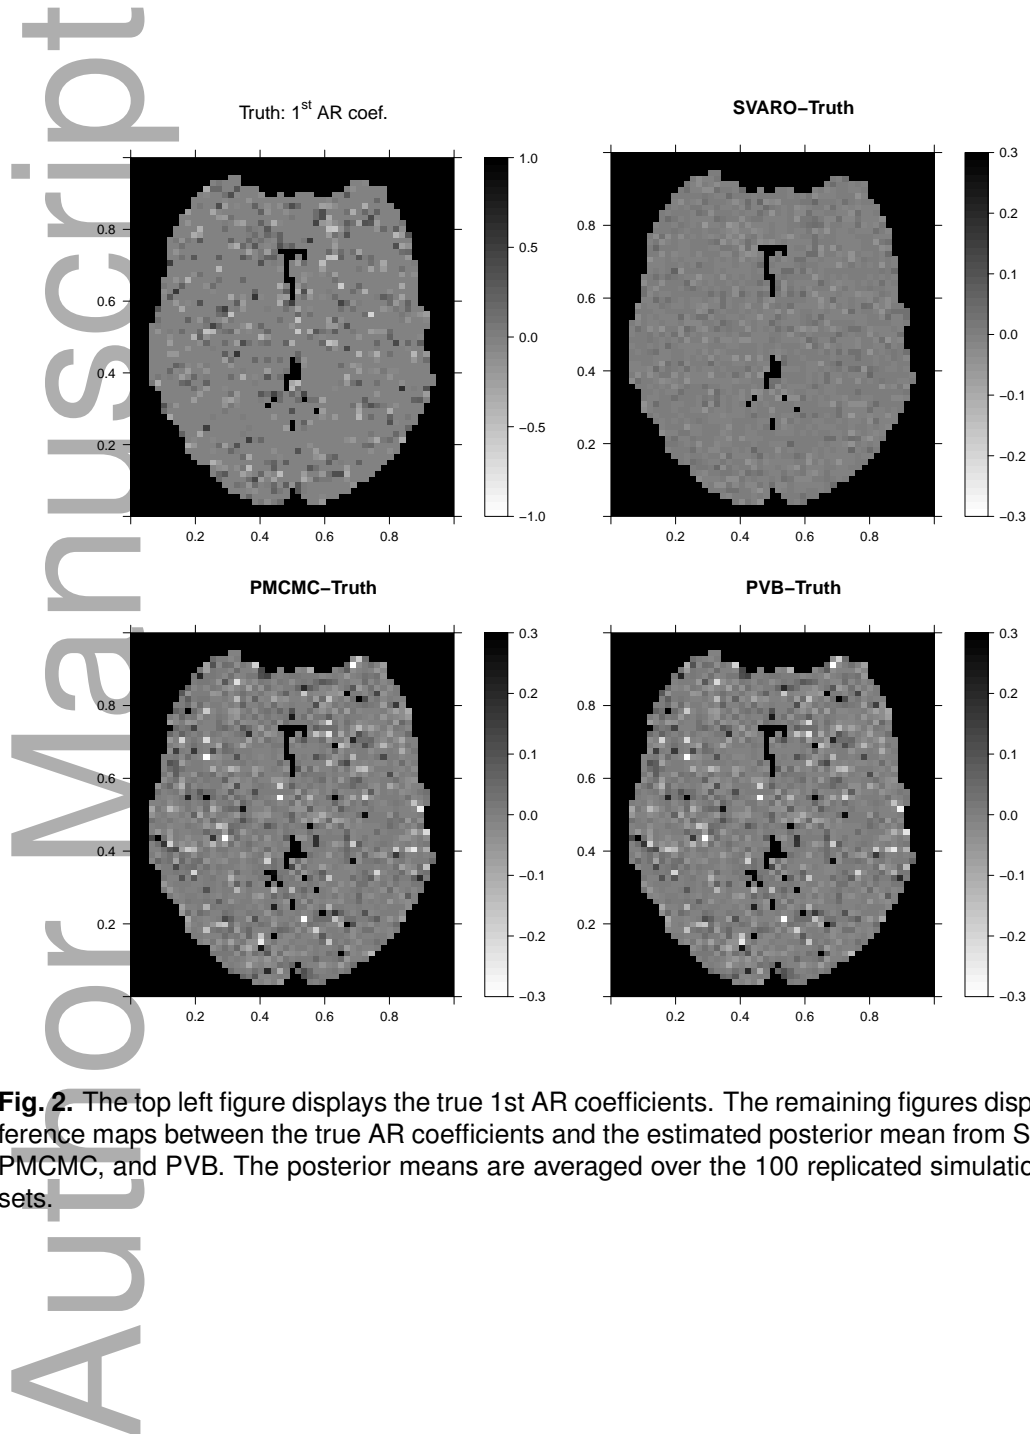


Fig. 2. The top left figure displays the true 1st AR coefficients. The remaining figures display difference maps between the true AR coefficients and the estimated posterior mean from SVARO, PMCMC, and PVB. The posterior means are averaged over the 100 replicated simulation data sets.

Table 1. $\overline{\text{MSE}}$, LPML and Timing for the three models. $\overline{\text{MSE}}$ is calculated by averaging MSE in each voxel and over simulation replicates. The MSE values for PMCMC and PVB are relative to those in SVARO.

	$\overline{\text{MSE}}$			LPML	Timing (min)
	$\mathbf{W}_{1..}$	$\mathbf{W}_{2..}$	\mathbf{a}_1		
SVARO	0.478	0.030	0.001	-1842902	108
PMCMC	113%	135%	509%	-1926620	11
PVB	199%	138%	510%	NA	1

Table 1 summarizes the average MSE ($\overline{\text{MSE}}$) for various parameters. These summaries are obtained by averaging the MSE of the corresponding parameters across all the voxels and over simulation replicates. It is clear that SVARO has the smallest $\overline{\text{MSE}}$ for these three parameters. In addition, PMCMC outperforms PVB in estimating the coefficient of the hemodynamic response, $\mathbf{W}_{1..}$, which is the primary parameter upon which inference is based. This finding is in line with our previous findings in Teng et al. (2016), where we found that a low SNR is one setting where MCMC outperforms VB for this particular model (GLM-AR). We also calculate and present the log-pseudo marginal likelihood, averaged over simulations, (LPML, Gelfand and Dey (1994)) in Table 1. SVARO has a larger LPML than the MCMC implementation of GLM-AR, indicating a better model fit. Note that LPML cannot be obtained from the VB algorithm. In terms of timing, SVARO takes 108 minutes with 10,000 iterations following 10,000 burn-in iterations, PMCMC takes 11 minutes with the same number of iterations, PVB is the fastest at 1 minute computation time.

We next investigate how the differences observed for the individual parameters will impact the overall inference of interest. A sensitivity plot is presented in Figure 3. This figure is obtained by plotting the average sensitivity against a range of marginal posterior probability thresholds from 0.9 to 1. We choose this range because it covers those values most often used in practice.

In terms of the underlying activation threshold, we use two thresholds: the true value of the contrast that corresponds to the top 10% and top 5% of all voxels. Thus, corresponding to a certain activation threshold and a certain probability threshold, the higher the sensitivity, the better the model is in terms of capturing activation. Again, a notable difference is observed when comparing the three methods, with SVARO giving the uniformly highest sensitivity across the entire range of probability thresholds and PVB resulting in the lowest sensitivity. PMCMC is better than PVB but still underperforms relative to SVARO.

We plot the posterior probability maps (PPM) in Figure 4. This figure depicts the locations of the true activations and the posterior probability maps from SVARO. In addition, differences in the probability maps comparing SVARO with PMCMC and PVB are also depicted. Again, SVARO appears to perform the best in producing the highest posterior probabilities for regions that are truly activated. PMCMC is similar to SVARO but its probability on those activated regions are slightly lower than those from SVARO, especially near the boundary of activated regions. PVB underperforms compared with the other two approaches by providing greater posterior probability on

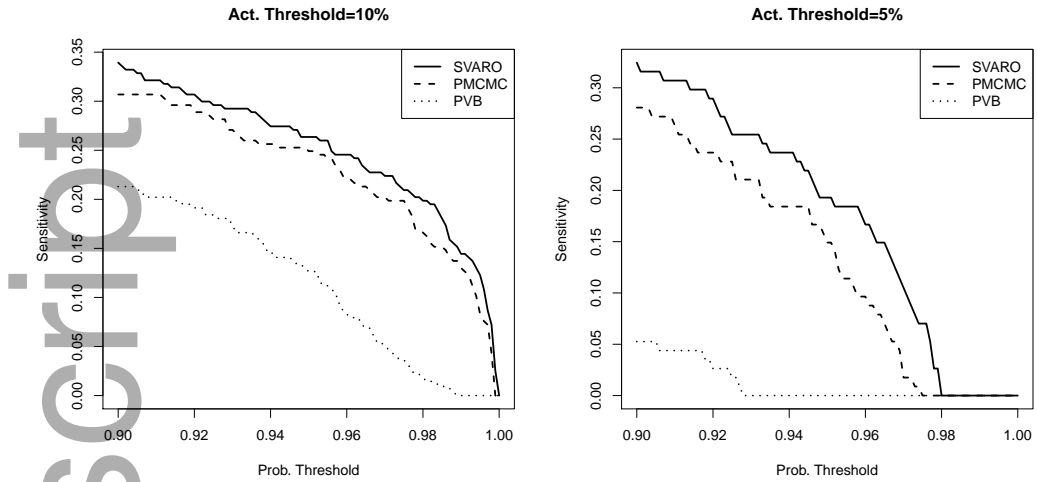


Fig. 3. Thresholded sensitivity curves for the three methods, with two activation thresholds. Left: activation threshold corresponds to the top 10% of the parameter estimates; Right: top 5%.

null locations while providing smaller posterior probability on activated locations.

3.3. Simulation 2

In this simulation study we compare the performance of our model under model misspecification against the correctly specified GLM-AR model using the two algorithms PMCMC and PVB. Here we simulate data under the competing GLM-AR model. The AR coefficients are simulated using the LORETA prior and the AR order is set to 1 for every voxel, a typical assumption used when analyzing fMRI data. The prior precision, τ_{u_1} is set at 400. The maximum AR order allowed for our model is set at 12.

Table 2 shows the average MSE summaries of the estimators. When data are simulated under the competing model, SVARO still exhibits good performance in terms of average MSE for the two regression parameters, suffering only a slight loss of efficiency compared to the GLM-AR model using the PMCMC algorithm. Nevertheless, the average LPML indicates that, overall, SVARO fits the simulated data better than the GLM-AR model. It is worth mentioning that PVB again under performs relative to PMCMC in terms of the hemodynamic response parameter and the AR coefficient.

Figure 5 presents the sensitivity curves. Despite the data being simulated under a constant order AR assumption, SVARO demonstrates similar sensitivity to that of PMCMC. The sensitivity curve for PVB is uniformly smaller than the other two because of the inaccurate estimation of the hemodynamic response parameter.

Overall, these simulation studies show that our SVARO model is more efficient from a statistical point of view than the GLM-AR model when the data display heterogeneous AR orders across the image—a situation we believe to be more common than not. They also show that our model does not suffer much in terms of statistical efficiency when the

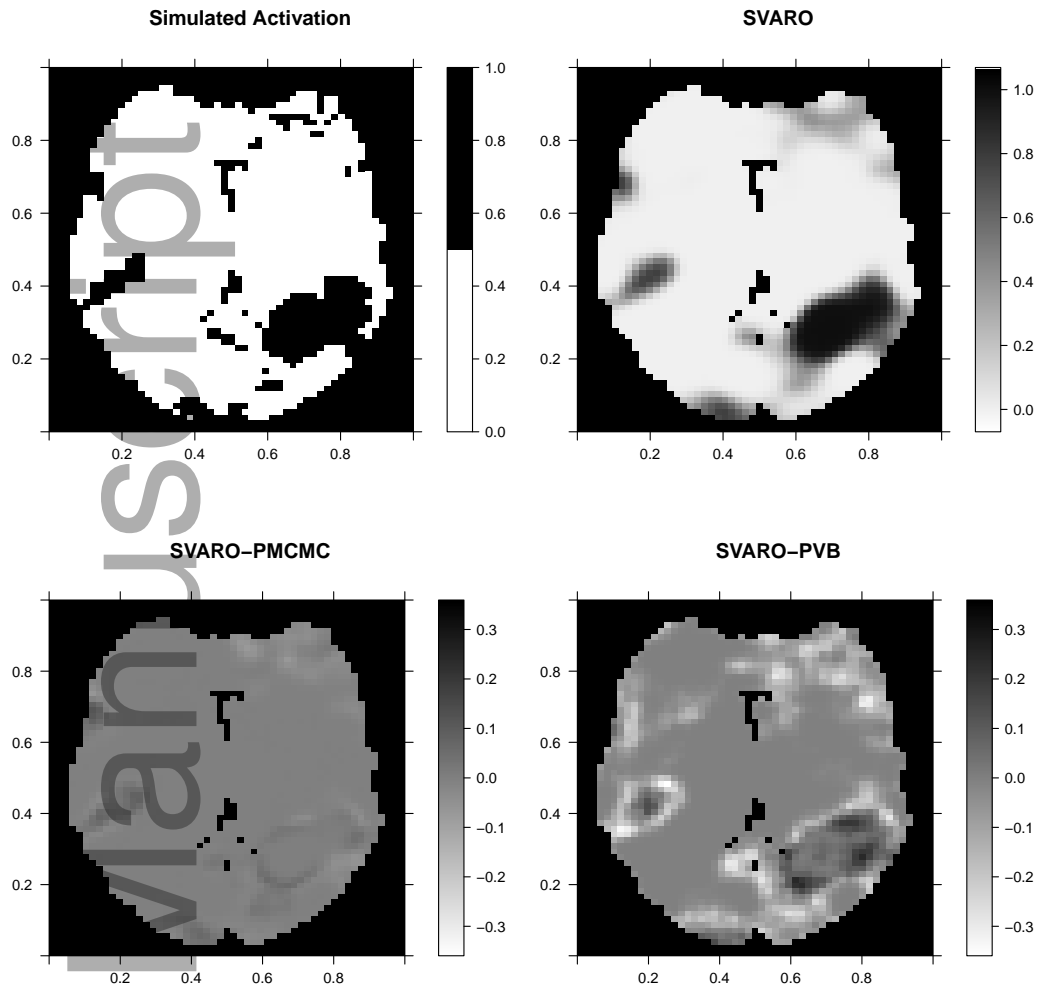


Fig. 4. Top left: the true activation map (red denotes activated voxels). The remaining panels are posterior probability maps (PPM) of activation obtained using SVARO, (SVARO-PMCMC) and (SVARO-PVB). The latter two reflect the difference of the two alternative approaches relative to SVARO.

AR order is small and homogeneous across the image.

4. Analysis of Face Perception fMRI Data

We turn our focus back to the face repetition data set that originally motivated our model development and compare results from the two models and three algorithms. In this analysis we use the complete experimental design consisting of famous faces, repeated twice and unfamiliar faces, repeated twice. These four design vectors are then

Table 2. $\overline{\text{MSE}}$, LPML and computation time for the three methods. $\overline{\text{MSE}}$ is calculated by averaging MSE over voxels and simulation replicates. The MSE values for PMCMC and PVB are relative to those in SVARO.

	$\overline{\text{MSE}}$			LPML	Timing (min)
	$\mathbf{W}_{1..}$	$\mathbf{W}_{2..}$	\mathbf{a}_1		
SVARO	0.502	0.031	0.003	-1817287	206
PMCMC	99%	97%	54%	-1875900	11
PVB	167%	98%	49%	NA	1

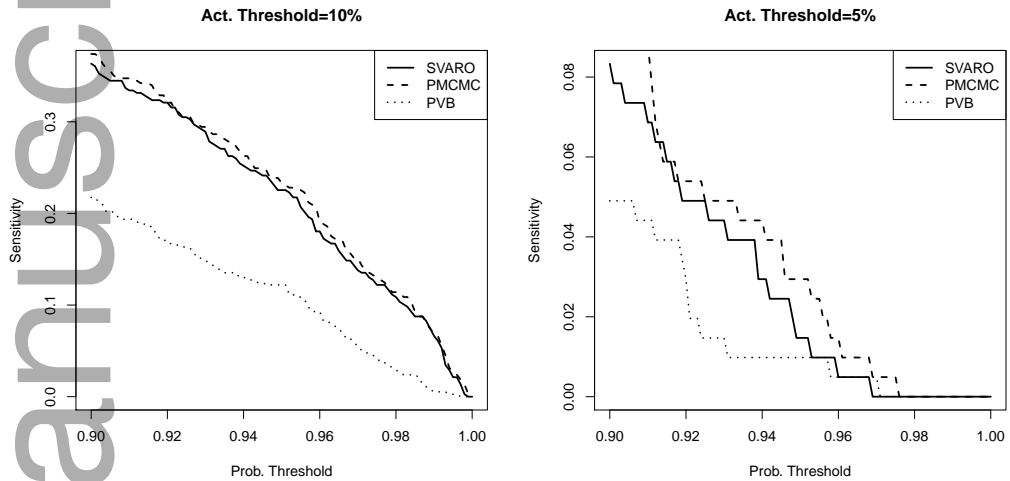


Fig. 5. Thresholded sensitivity curves for the three methods, with two activation thresholds. Left: activation threshold corresponds to the top 10% of the parameter estimates; Right: top 5%.

convolved with the canonical HRF as well as its time and dispersion derivatives. An intercept term is also added to the final design matrix for a total of 13 covariates. We allow an AR order up to a maximum of $P = 12$ when fitting the SVARO model, and an $AR(3)$ for the GLM-AR model using both the PMCMC and PVB algorithms. While the choice of an $AR(3)$ for the latter two approaches may seem arbitrary, this is exactly the justification for the use of the SVARO model where such an arbitrary assumption need not be made. We consider data collected on a single subject in what follows.

Pre-processing steps are applied to the data prior to fitting the Bayesian models. All functional images are aligned to the first image using a six-parameter rigid-body transformation. Then slice-timing correction is performed to set the standard acquisition time as the 12th slice. Images are spatially normalized to a standard EPI image. The global mean, g , is computed and each time series is divided by $100/g$ to represent a percentage of g . Finally, a high-pass filter with cut-off frequency of $1/128\text{Hz}$ is applied to the data and design matrix to remove low frequency signals that arise through scanner drift.

Figure 6 presents the distribution of estimated posterior modal AR orders from our model across voxels. The most frequent order is the zero order, or no autocorrelation in the time series, accounting for approximately 35% of the voxels. Interestingly, the next highest is order 8, with 9.4% of the voxels. Overall, roughly 51% of the voxels exhibit an AR order greater than 3. The existence of these higher orders and the variability in the orders is in general agreement with our exploratory analysis of the face repetition data set (see Section 1.1) and indicates the necessity of our proposed model.

Next, we compare models/algorithms with respect to a particular contrast of interest, namely the effect of fame. That is, famous faces versus unfamiliar faces. The estimated posterior mean and standard deviation (SD) maps for the fame contrast are displayed in Figure 7. While the posterior SD estimated from SVARO and PMCMC are very close, the estimated posterior means are different between the two approaches. Also, the estimated SD obtained from PVB shows apparent discretization. This is due to a graph-partitioning that is incorporated in the algorithm for the sake of computational speed (Penny et al., 2007). It is clear that the boundaries of these graph-partitioned regions have substantially larger estimated SD than the interior locations. The posterior mean of PVB also seems to exhibit artifacts at the partition boundaries, though the effect is not as pronounced. The estimated LPML is -4.66×10^7 under our proposed model and is -4.82×10^7 under GLM-AR with MCMC sampling. According to this model selection criterion, our proposed model is preferred.

Finally, we look at the effect of fame using thresholded PPMs. The activation threshold is set to 0.2% of the global mean value, and the probability threshold is set to 0.95. Figure 8 shows the activation regions projected onto the pial surface. We can see that there is a match in terms of a majority of activation regions inferred from SVARO and PMCMC. A closer look reveals that PMCMC tends to make more scattered predictions across the posterior regions of the brain. The number of activation regions from PVB are far greater than the number obtained from the other two approaches, and are more widely dispersed across the brain. From our simulation study results, we suspect that these scattered activated regions are likely false positives.

The main parameters of interest are the regression parameters corresponding to the HRF for the four experimental tasks. We use Geweke (1984) diagnostics along with FDR correction for dependent variables (Benjamini and Yekutieli, 2001) to determine convergence of our model. Although there is a multivariate version of the Geweke diagnostic, our parameter space is too large for this method. Out of the four parameters for each of the 56526 voxels (total of 226104 parameters), 99.89% of the parameters converged to their posterior distributions. In terms of computation time, PVB took about 1 hour, PMCMC took 1 day, while SVARO took about 1 week of computation. Much of our MCMC algorithm is amenable to parallel programming, which is an avenue for further development.

5. Discussion

In this paper, we have developed a Bayesian hierarchical model, SVARO, that allows the AR order to vary spatially across the brain, with the orders themselves displaying a certain level of spatial clustering based on an Ising model. We compared our proposed

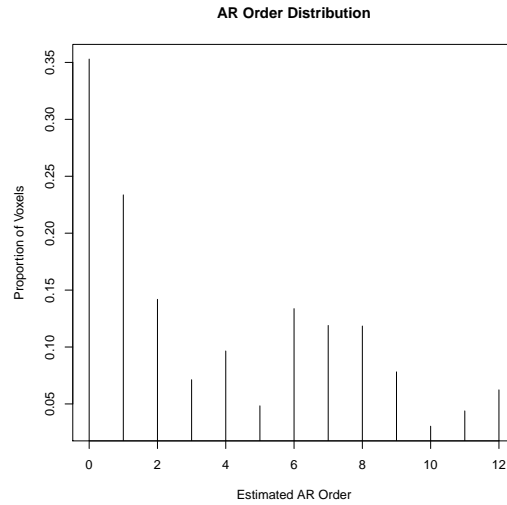


Fig. 6. Histogram of the estimated posterior modal AR order from the SVARO model across all voxels for the face repetition data set.

model with a self-written MCMC sampler for the standard GLM-AR model and the VB implementation for the same model. The results are interesting, under a low SNR ratio, VB seems to suffer from variance overestimation, leading to a larger MSE than the other two methods. It is likely that as temporal noise increases, a more vital role is played by the AR correlation that increases the posterior correlation between different parameters and this makes the mean field assumption underlying the VB approximation less accurate.

We have shown that our model outperforms the GLM-AR model not only in terms of accuracy and sensitivity, but also according to formal model selection based on the LPML criterion. Through an application of our proposed model and through an exploratory analysis, we have shown that the typical constant low-order AR assumption can be violated in real fMRI data. It is very likely that this issue, seen in the face repetition data set, is also present in other fMRI data sets.

There is a computational price to be paid for gaining the flexibility we have proposed in our model. Our model takes a longer time to run than either implementation of GLM-AR. This is mostly due to the estimation of the varying AR orders and the associated greater number of parameters to estimate. However, as previously mentioned, there are elements of our MCMC algorithm that are amenable to parallel programming. This will be investigated in future work.

While we have based our model specification on a set of independent Ising processes, one for each possible order of the AR process, another approach would be to assume a Potts model for the orders of the AR coefficients. A Potts model, combined with a Dirichlet process prior for parameters has been investigated for selecting covariates of interest in brain imaging (Johnson et al., 2013). Here we can also apply it to the selection

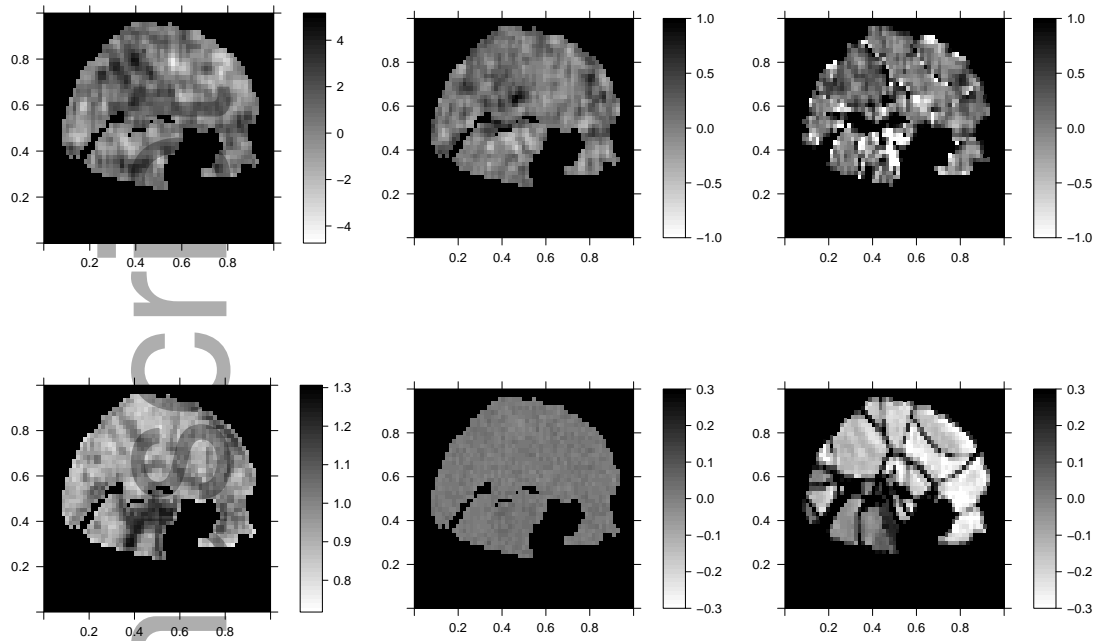


Fig. 7. Posterior estimates from the middle (27/53) slice of the brain, sagittal view. Top row: posterior mean of fame. Bottom row: posterior standard deviation of fame. Left column: results from SVARO. Center column: difference map, SVARO–PMCMC. Right column: difference map, SVARO–PVB.

of autoregressive orders to yield a more parsimonious, yet still flexible, model. Investigation of hyper-parameter estimation in the Ising model and the use of alternative spatial models is also of interest, as is increasing the scope of our comparison of methods to include wavelet approaches that focus on long memory errors, or VAR models (Harrison et al., 2003).

References

- Benjamini, Y. and Yekutieli, D. (2001) The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, **29**, 1165–1188.
- Bezener, M., Eberly, L. E., Hughes, J., Jones, G. and Musgrove, D. R. (2016) Bayesian spatiotemporal modeling for detecting neuronal activation via functional magnetic resonance imaging. *Handbook of Big Data Analytics, Springer Handbooks of Computational Statistics*. Springer.
- Bullmore, E., Brammer, M., Williams, S. C., Rabe-Hesketh, S., Janot, N., David, A., Mellers, J., Howard, R. and Sham, P. (1996) Statistical methods of estimation and




Fig. 8. Activation maps for fame contrast. Top row: anterior view. Bottom row: posterior views. Left column: activation (white colors) from SVARO. Middle column: activation from PMCMC. Right column: activation from PVB.

inference for functional mr image analysis. *Magnetic Resonance in Medicine*, **35**, 261–277.

Bullmore, E., Fadili, J., Maxim, V., Şendur, L., Whitcher, B., Suckling, J., Brammer, M. and Breakspear, M. (2004) Wavelets and functional magnetic resonance imaging of the human brain. *Neuroimage*, **23**, S234–S249.

Castruccio, S., Ombao, H. and Genton, M. G. (2016) A multi-resolution spatio-temporal model for brain activation and connectivity in fmri data. *arXiv preprint arXiv:1602.02435*.

Chang, C. and Glover, G. H. (2010) Time–frequency dynamics of resting-state brain connectivity measured with fmri. *Neuroimage*, **50**, 81–98.

Della-Maggiore, V., Chau, W., Peres-Neto, P. R. and McIntosh, A. R. (2002) An empirical comparison of spm preprocessing parameters to the analysis of fmri data. *Neuroimage*, **17**, 19–28.

Fadili, M. and Bullmore, E. (2002) Wavelet-generalized least squares: a new blu estimator of linear regression models with $1/f$ errors. *NeuroImage*, **15**, 217–232.

- Friston, K., Josephs, O., Zarahn, E., Holmes, A., Rouquette, S. and Poline, J.-B. (2000) To smooth or not to smooth?: Bias and efficiency in fmri time-series analysis. *NeuroImage*, **12**, 196–208.
- Friston, K. J., Holmes, A. P., Poline, J., Grasby, P., Williams, S., Frackowiak, R. S. and Turner, R. (1995) Analysis of fmri time-series revisited. *Neuroimage*, **2**, 45–53.
- Gelfand, A. E. and Dey, D. K. (1994) Bayesian model choice: asymptotics and exact calculations. *Journal of the Royal Statistical Society. Series B (Methodological)*, 501–514.
- Gelman, A. and Meng, X.-L. (1998) Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science*, 163–185.
- George, E. I. and McCulloch, R. E. (1993) Variable selection via gibbs sampling. *Journal of the American Statistical Association*, **88**, 881–889.
- Geweke, J. F. (1984) Measures of conditional linear dependence and feedback between time series. *Journal of the American Statistical Association*, **79**, 907–915.
- Gössl, C., Auer, D. P. and Fahrmeir, L. (2001) Bayesian spatiotemporal inference in functional magnetic resonance imaging. *Biometrics*, **57**, 554–562.
- Harrison, L., Penny, W. D. and Friston, K. (2003) Multivariate autoregressive modeling of fmri time series. *Neuroimage*, **19**, 1477–1491.
- Higdon, D. M. (1998) Auxiliary variable methods for markov chain monte carlo with applications. *Journal of the American Statistical Association*, **93**, 585–595.
- Ising, E. (1925) Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik*, **31**, 253–258.
- Jeong, J., Vannucci, M. and Ko, K. (2013) A wavelet-based bayesian approach to regression models with long memory errors and its application to fmri data. *Biometrics*, **69**, 184–196.
- Johnson, T. D., Liu, Z., Bartsch, A. J. and Nichols, T. E. (2013) A bayesian non-parametric potts model with application to pre-surgical fmri data. *Statistical methods in medical research*, **22**, 364–381.
- Johnson, V. E. and Rossell, D. (2012) Bayesian model selection in high-dimensional settings. *Journal of the American Statistical Association*, **107**, 649–660.
- Kim, S., Smyth, P. and Stern, H. (2010) A bayesian mixture approach to modeling spatial activation patterns in multisite fmri data. *IEEE transactions on medical imaging*, **29**, 1260–1274.
- Lee, K.-J., Jones, G. L., Caffo, B. S. and Bassett, S. S. (2014) Spatial bayesian variable selection models on functional magnetic resonance imaging time-series data. *Bayesian Analysis (Online)*, **9**, 699.

- Li, F., Zhang, T., Wang, Q., Gonzalez, M. Z., Maresh, E. L., Coan, J. A. et al. (2015) Spatial Bayesian variable selection and grouping for high-dimensional scalar-on-image regression. *The Annals of Applied Statistics*, **9**, 687–713.
- Locascio, J. J., Jennings, P. J., Moore, C. I. and Corkin, S. (1997) Time series analysis in the time domain and resampling methods for studies of functional magnetic resonance brain imaging. *Human brain mapping*, **5**, 168–193.
- Makni, S., Ciuciu, P., Idier, J. and Poline, J.-B. (2006) Joint detection-estimation of brain activity in fmri using an autoregressive noise model. In *Biomedical Imaging: Nano to Macro, 2006. 3rd IEEE International Symposium on*, 1048–1051. IEEE.
- Meyer, F. G. (2003) Wavelet-based estimation of a semiparametric generalized linear model of fmri time-series. *IEEE transactions on medical imaging*, **22**, 315–322.
- Mitchell, T. J. and Beauchamp, J. J. (1988) Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, **83**, 1023–1032.
- Musgrove, D. R., Hughes, J. and Eberly, L. E. (2016) Fast, fully bayesian spatiotemporal inference for fmri data. *Biostatistics*, **17**, 291–303.
- O’Hara, R. B., Sillanpää, M. J. et al. (2009) A review of bayesian variable selection methods: what, how and which. *Bayesian analysis*, **4**, 85–117.
- Penny, W., Flandin, G. and Trujillo-Barreto, N. (2007) Bayesian comparison of spatially regularised general linear models. *Human brain mapping*, **28**, 275–293.
- Penny, W., Kiebel, S. and Friston, K. (2003) Variational Bayesian inference for fMRI time series. *NeuroImage*, **19**, 727–741.
- Penny, W. D., Trujillo-Barreto, N. J. and Friston, K. J. (2005) Bayesian fMRI time series analysis with spatial priors. *NeuroImage*, **24**, 350–362.
- Schmidt, D. F. and Makalic, E. (2013) Estimation of stationary autoregressive models with the bayesian lasso. *Journal of Time Series Analysis*, **34**, 517–531.
- Shu, H., Nan, B. and Koeppe, R. (2015) Multiple testing for neuroimaging via hidden markov random field. *Biometrics*, **71**, 741–750.
- Skudlarski, P., Constable, R. T. and Gore, J. C. (1999) Roc analysis of statistical methods used in functional mri: individual subjects. *Neuroimage*, **9**, 311–329.
- Smith, M. and Fahrmeir, L. (2007) Spatial bayesian variable selection with application to functional magnetic resonance imaging. *Journal of the American Statistical Association*, **102**, 417–431.
- Stanley, H. E., Stauffer, D., Kertesz, J. and Herrmann, H. J. (1987) Dynamics of spreading phenomena in two-dimensional ising models. *Physical review letters*, **59**, 2326.
- Swendsen, R. H. and Wang, J.-S. (1987) Nonuniversal critical dynamics in monte carlo simulations. *Physical review letters*, **58**, 86.

- Teng, M., Johnson, T. and Nathoo, F. (2016) A comparison of variational bayes and hamiltonian monte carlo for bayesian fmri time series analysis with spatial priors. *arXiv preprint arXiv:1609.02123*.
- Wang, H., Li, G. and Tsai, C.-L. (2007) Regression coefficient and autoregressive order shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **69**, 63–78.
- Woolrich, M. W., Behrens, T. E. and Smith, S. M. (2004a) Constrained linear basis sets for hrf modelling using variational bayes. *NeuroImage*, **21**, 1748–1761.
- Woolrich, M. W., Jenkinson, M., Brady, J. M. and Smith, S. M. (2004b) Fully bayesian spatio-temporal modeling of fmri data. *IEEE transactions on medical imaging*, **23**, 213–231.
- Woolrich, M. W., Ripley, B. D., Brady, M. and Smith, S. M. (2001) Temporal autocorrelation in univariate linear modeling of fmri data. *Neuroimage*, **14**, 1370–1386.
- Worsley, K. J. and Friston, K. J. (1995) Analysis of fmri time-series revisited again. *Neuroimage*, **2**, 173–181.
- Zarahn, E., Aguirre, G. K. and D’Esposito, M. (1997) Empirical analyses of bold fmri statistics. *NeuroImage*, **5**, 179–197.