University of Michigan School of Public Health
The University of Michigan Department of Biostatistics Working Paper Series

Year 2007 *Paper* 71

Bayesian Spatial Modeling of fMRI data: A Multiple-Subject Analysis

Lei Xu* Timothy Johnson†
Thomas Nichols‡

*lei.xu@vanderbilt.edu
†tdjtdj@umich.edu
‡

Bayesian Spatial Modeling of fMRI data: A Multiple-Subject Analysis

Lei Xu, Timothy Johnson, and Thomas Nichols

Abstract

The aim of this work is to develop a spatial model for multi-subject fMRI data. While there has been much work on univariate modeling of each voxel for single- and multi-subject data, and some work on spatial modeling for single-subject data, there has been no work on spatial models that explicitly account for intersubject variability in activation location. We use a Bayesian hierarchical spatial model to fit the data. At the first level we model "population centers" that mark the centers of regions of activation. For a given population center each subject may have zero or more associated "individual components". While most previous work uses Gaussian mixtures for the activation shape, we instead use Gaussian mixtures for the probability that a voxel belongs to an activated region, assuming homogeneous mean intensity within a region. Our approach incorporates the unknown number of mixture components into the model as a parameter whose posterior distribution is estimated by reversible jump Markov Chain Monte Carlo. We demonstrate our method with a fMRI study of vi- sual working memory and show dramatically better precision of localization with our method relative to the standard mass-univariate method. Although we are motivated by fMRI data, this model could easily be modified to handle other types of imaging data.

Bayesian Spatial Modeling of fMRI data: A Multiple-Subject Analysis

Lei Xu[1*] Timothy D. Johnson[1] Thomas E. Nichols[2,3]

[1]Department of Biostatistics
School of Public Health
University of Michigan
Ann Arbor, MI 48109

[2]GlaxoSmithKline
Clinical Imaging Centre
Imperial College London
Hammersmith Hospital, W12 0NN
United Kingdom

[3]FMRIB Centre
Oxford University
Oxford, OX3 9DU
United Kingdom

Summary. The aim of this work is to develop a spatial model for multi-subject fMRI data. While there has been much work on univariate modeling of each voxel for single- and multi-subject data, and some work on spatial modeling for single-subject data, there has been no work on spatial models that explicitly account for intersubject variability in activation location. We use a Bayesian hierarchical spatial model to fit the data. At the first level we model "population centers" that mark the centers of regions of activation. For a given population center each subject may have zero or more associated "individual components". While most previous work uses Gaussian mixtures for the activation shape, we instead use Gaussian mixtures for the probability that a voxel belongs to an activated region, assuming homogeneous mean intensity within a region. Our approach incorporates the unknown number of mixture components into the model as a parameter whose posterior distribution is estimated by reversible jump Markov Chain Monte Carlo. We demonstrate our method with a fMRI study of visual working memory and show dramatically better precision of localization with our method relative to the standard mass-univariate method. Although we are motivated by fMRI data, this model could easily be modified to handle other types of imaging data.

Key Words: Bayesian hierarchical model; Functional brain mapping; Multi-subject fMRI data analysis; Reversible jump Markov chain Monte Carlo; Spatial mixture model.

*Corresponding Author: email: leix@umich.edu

1 Introduction

Among the methods for mapping brain function, functional magnetic reasonance imaging (fMRI) is one of the most widely used. Conventionally a classical, mass-univariate approach is used, where univariate time series models are fit independently at each voxel. These models are used to create images of parameter estimates and test statistics, which are then assessed for significance (Friston et al., 1995). While this is a computationally efficient method it has several limitations. The approach does not account for the spatial nature of the image data, and does not utilize the prior knowledge that spatially contiguous patterns of activation are expected. Multisubject analyses are particularly problematic since, even after registration of the subjects' brain to a common atlas, there is residual variation in the anatomical landmarks; further, it has been shown that even if sulci and gyri are aligned there is variation in the functional landmarks (Morosan et al., 2001). A mass-univariate model cannot account for any mismatch in activation location and will only detect voxels with consistent change in activation.

In this work we develop a Bayesian hierarchical model that improves the standard methods in several ways. First, we pose an explicit spatial model for activations at the subject level, unlike the mass-univariate method that fits each voxel independently of its neighbors. Second, while some authors have proposed spatial models, they have only considered single subject data; our method is for multi-subject data and accounts for intersubject heterogeneity in activation location about a population location. Lastly, we use a fully Bayesian framework where all sources of uncertainty are considered and quantified.

Our work is motivated by an event-related fMRI study (Park et al., 2003). Park et al. investigated the brain regions involved in a visual working memory task. Briefly, twenty-one adult subjects were scanned during two separate tasks. During the "visual" task a subject was shown a scene for six seconds and asked to memorize it. Immediately thereafter, they were shown a small fragment of a scene and asked whether this fragment was part of the

original scene. During the "maintenance" task, subjects were shown a scene for two seconds and where asked to memorize it for four seconds. They were again shown a small fragment of a scene and asked whether it was part of the original scene. The effect of interest was the difference between neuronal activity during the four second maintenance task and the last four seconds of the visual task. See Park et al. (2003) for full study details..

Inference on mass-univariate models is made by assessing the statistic images, where each voxel measures evidence against the null hypothesis of no effect. A threshold is then applied to define regions of significant activation. Various methods on how to choose the threshold have been proposed (Genovese, Lazar, and Nichols, 2002; Nichols and Hayasaka, 2003). While some methods find a threshold that accounts for correlation in the null hypothesis statistic images (Worsley et al., 1996), the spatial properties of the signal are not modeled. Recently, however, several authors have developed spatial models for the signal in fMRI data using Bayesian methods.

Woolrich et al. (2005) develop a spatial mixture model using a discrete Markov random field (MRF) prior on a spatial map of classification labels. For each voxel there is a binary latent class variable denoting whether this voxel is active or inactive. Within this framework one can calculate the posterior probability of a voxel being activated. Spatial regularization is controlled by a parameter in the discrete MRF prior that penalizes neighboring voxels of different classes. The spatial regularization parameter is typically assumed to be known. However, results can be highly dependent on this parameter. Estimating this parameter is problematic because the computational evaluation of the normalizing constant of the discrete MRF prior is difficult, yet necessary. Hence, Woolrich et al. (2005) cleverly approximate the discrete MRF prior with a continuous MRF prior whose normalizing constant is easily calculated. Hartvig and Jensen (2000) also use a spatial mixture model and achieve computational feasibility by formulating the model through the marginal distribution on a small grid of voxels. By carefully choosing prior distributions they are able to calculate the pos-

terior distribution analytically. This provides inference that is much faster than simulation based methods. In another paper, Hartvig (2002) proposes a regression based spatial model using the idea of "activation centers". The activation surface is modeled as a sum of Gaussian functions with parameters describing the shape and the magnitude of the center. They design a reversible jump algorithm to insert, delete and change an activation center given the variance parameters are known. They estimate the variance parameters separately in a different procedure using the distribution of some sufficient statistics.

While these other methods model spatial dependence in the signal, they focus solely on single subject analyses. We are aware of only one method for multi-subject data. Miglioretti, McCulloch, and Zeger (2002) develop a multi-subject model for surface electrode data. Their model is developed for binary observations that is not applicable in the current setting.

In this paper we propose a Bayesian spatial model for multi-subject fMRI data. Our model can be used to make inference on all facets of the data, but specifically, on the population location of activations, population prevalance of an activation at a location, and intersubject variability in a location of activation. Throughout we consider the effect magnitude image as the observations (or the data), one per subject. These magnitudes can be interpreted, roughly, as map of the percent change in blood flow—an indirect measure of neuronal activity. Although our model can be easily extended to analyze 3D data, in this paper we will only focus on one slice of the 3D image. The top row of Figure 3 displays data from 4 of the 21 subjects in the data set.

The paper is organized as follows. In Section 2, we introduce the spatial mixture model in a fully Bayesian framework. In Section 3, we describe the algorithm used for posterior inference and discuss how we summarize posterior inference in Section 4. In Section 5, we present simulation results as well as results from our motivating example and a sensitivity analysis. We conclude the manuscript with a discussion and ideas for future work.

2 Model

We begin with an overview of the model and notation (with all subscripts initially suppressed), after which we present the likelihood and priors in detail.

Our model is specified hierarchically. At the first level we have an unknown number of "population centers", $\boldsymbol{\mu}$, that are uniformly distributed within the confines of the brain. At the second level, an unknown number of "individual component" means, $\boldsymbol{\eta}$, are distributed as isotropic Gaussian mixtures whose mixture component means are the population centers with variances τ^2. We assume that each subject has been fit with an intrasubject fMRI signal model, producing scalar images of the fMRI blood oxygenation level dependent (BOLD) effect magnitude; we refer to these intrasubject summary measures as "the data". At the third and final level, we assume the data, y, for each subject are distributed as a Gaussian mixture with an unknown number of mixing components whose means are θ with variances σ^2. The mixing weights for the datum at pixel v are proportional to the statistical distance from v to $\boldsymbol{\eta}$ with a variance r^2. There is one special component representing the constant background intensity for each subject. We now present the details of our model.

2.1 Likelihood

We begin with the third level of the hierarchy. Let y_{jv} be the observed effect magnitude at pixel v, $v = 1, \ldots, V$, for subject j, $j = 1, \ldots, J$. We use a single index v to reference a 2D pixel, $\mathbf{x}_v = (x_{1v}, x_{2v})$. We assume that observations are distributed as a mixture of Gaussian components:

$$f(y_{jv} \mid p_{jvl}, \theta_{jl}, \sigma^2_{jl}, c_j) = p_{jv0}\, \phi(y_{jv}; \theta_0, \sigma^2_0) + \sum_{l=1}^{c_j} p_{jvl}\, \phi(y_{jv}; \theta_{jl}, \sigma^2_{jl}). \tag{1}$$

Here, c_j is the number of mixture components for subject j (not including the background component) and $\phi(a; b, c^2)$ is the density at a of a normal distribution with mean b and variance c^2. θ_0 and σ^2_0 represent the mean and variance of the background component. By

introducing a latent variable ω_{jv} with $\Pr(\omega_{jv} = l) = p_{jvl}$, the likelihood can be written conditionally as

$$\begin{aligned} f(y_{jv} \mid \theta_0, \sigma_0^2, \omega_{jv} = 0) &= \phi(y_{jv}; \theta_0, \sigma_0^2) \quad \text{for background,} \\ f(y_{jv} \mid \theta_{jl}, \sigma_{jl}^2, \omega_{jv} = l) &= \phi(y_{jv}; \theta_{jl}, \sigma_{jl}^2) \quad \text{otherwise.} \end{aligned}$$

The latent variable ω_{jv} can be interpreted as an allocation variable in the sense that observation y_{jv} is assigned to one of the mixture components with probability p_{jvl}. A pixel is considered active if that pixel belongs to a component with index $l > 0$ and is inactive if that pixel belongs to the background component ($l = 0$).

Let $\phi_2(\cdot; \mathbf{a}, \mathbf{B})$ denote the density of a bivariate normal distribution with mean $\mathbf{a}$ and covariance matrix $\mathbf{B}$. The mixing weights, p_{jvl}, take the form

$$p_{jvl} \propto \begin{cases} m & l = 0 \\ \phi_2(\mathbf{x}_v; \boldsymbol{\eta}_{jl}, \mathbf{R}_{jl}) & l = 1, \ldots, c_j \end{cases}$$

with $\sum_{l=0}^{c_j} p_{jvl} = 1$. Here $\mathbf{x}_v$ is the spatial location of pixel v and $\boldsymbol{\eta}_{jl}^{\mathrm{T}} = (\eta_{1jl}, \eta_{2jl})$ is the mean of individual component l, $l = 1, ..., c_j$ (we refer to the mean of the individual component as the "component center"). $\mathbf{R}_{jl} = r_{jl}^2 \mathbf{I}_2$ where $\mathbf{I}_2$ is the 2×2 identity matrix. The spatial dependence of the data is captured by the weights, p_{jvl}. Given m, the weights largely depend on the distance from the pixel to each of the component centers. This implies that observations that are spatially close to one another are more likely to have similar weights. Hence, spatially close observations will be more correlated than distant observations. Note that an observation distant from all components centers, a priori, will have probability that it belong to the background center close to 1. Furthermore, if a pixel and component center are coincident, then the a priori probability that this pixel belongs to the background is approximately $m/(r_{jl}^{-2}/2\pi + m)$. fMRI experiments are usually designed such that only a small percentage of the brain (roughly 1%–5%) actually shows any activity. Thus, a priori, we set $m = 19$ and $E(r_{jl}^{-2}) = (2\pi)$. This gives $m/(E(r_{jl}^{-2})/2\pi + m) = 0.95$, reflecting this a priori belief.

The number of mixing components, c_j, for each subject is not known and we estimate it along with all other parameters. A priori we assume c_j is a Poisson random variate with mean 5, for $j = 1, \ldots, J$.

Now we move on to the second and first levels of the hierarchy.

2.2 Priors

The joint prior distribution is factored hierarchically as given in the Appendix.

The priors of the component means make up the second level of the hierarchy. The prior of each component mean, $\boldsymbol{\eta}_{jl}$, $l = 1, \ldots, c_j$ for subject j, is taken to be a mixture of c_p bivariate normals. Each component mean is associated with a particular population center, $\boldsymbol{\mu}_i$, $i = 1, \ldots, c_p$:

$$\pi(\boldsymbol{\eta}_{jl} \mid \psi_i, \boldsymbol{\mu}_i, \tau_i^2, c_p) = \sum_{i=1}^{c_p} \psi_i \, \phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}_i, \tau_i^2 \mathbf{I}_2). \tag{2}$$

Here, the ψ_i are mixing weights. We also introduce another latent variable, z_{jl}, such that $\Pr(z_{jl} = i) = \psi_i$. Thus, conditional on $z_{jl} = i$, $\boldsymbol{\eta}_{jl}$ is bivariate normal with mean $\boldsymbol{\mu}_i$ and covariance matrix $\tau_i^2 \mathbf{I}_2$. A natural choice for the prior on ψ_i is a symmetric Dirichlet distribution: $\psi_i \mid c_p \sim D(1, 1, \ldots, 1)$ where the parameter dimension is $c_p - 1$. The number of population centers, c_p, is a priori unspecified and is to be estimated along with all other model parameters.

At the first level of the hierarchy, the parameter $\boldsymbol{\mu}_i^{\mathrm{T}} = (\mu_{1i}, \mu_{2i})$ is the location of population activation center i, $i = 1, ..., c_p$. Let A_j denote the cross-sectional area of the given MRI slice of the brain of subject j, $j = 1, \ldots, J$. Set $A = \cup_{j=1}^{J} A_j$. (We note here that, although all subject's data have been mapped onto a common brain atlas, due to motion artifacts and field inhomogeneities, there are missing data. Typically, fMRI analyses are performed on the intersection of the A_j. By taking the union, we allow for the possibility that a population center is in a region where some subjects may have missing data.) A priori, we assume these population activation centers follow a homogenous spatial Poisson process with rate

λ_p defined over A. Thus, we can write

$$c_p \sim P(\lambda_p A) \quad \text{and} \quad \pi(\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_{c_p} \mid c_p) = A^{-c_p}.$$

We take $\lambda_p = 5A^{-1}$ reflecting our belief that the number of population centers should be small. Note that although the prior distribution of the population centers is a homogenous point process, it's posterior is not necessarily. The posterior will depend, in large part, on the posterior distribution of the $\boldsymbol{\eta}_{jl}$.

We place inverse gamma (parametrized such that if $x \sim IG(\alpha, \beta)$, $E(x) = \beta/(\alpha - 1)$) priors on all variance parameters for mathematical convenience:

$$r_{jl}^2 \sim IG(2\pi, \beta_r),\ \tau_i^2 \sim IG(3, \beta_\tau),\ \sigma_{jl}^2 \sim IG(3, \beta_\sigma),\ \sigma_0^2 \sim IG(.001, .001).$$

Hyperprior distributions are then placed on β_r, β_τ and β_σ: $\beta_x \sim G(.01, .01)$ where x can be any one of r, τ, or σ. These choices result in informative priors. However, placing rather vague hyperprior distributions on the scale parameters reduces the influence of the prior on the posterior estimates of r_{jl}^2, τ_i^2 and σ_{jl}^2.

It remains to specify the priors of the intensity parameters, θ_{jl}, in the likelihood specified in equation (1). Recall that θ_{jl}, $l > 0$, are the mean intensity levels of the individual mixture components for active pixels and θ_0 is the mean intensity level of inactivated pixels. We choose the following priors:

$$\theta_0 \sim N(0, 1) \quad \text{and} \quad \theta_{jl} \sim \text{trunc}_{(0,\infty)} N(\lambda_\theta, \sigma_\theta^2),$$

where

$$\lambda_\theta \sim N(35, 1 \times 10^8) \quad \text{and} \quad \sigma_\theta^2 \sim IG(.01, .01).$$

Furthermore, $\text{trunc}_{(0,\infty)} N(\lambda_\theta, \sigma_\theta^2)$ denotes the normal distribution with mean λ_θ and variance σ_θ^2 truncated to $(0, \infty)$. Results from a sensitivity analysis are given in Section 5.3.

3 Posterior estimation

The full posterior distribution does not have an analytic solution. Thus, the posterior distribution is simulated via Markov chain Monte Carlo (MCMC) techniques, and, in particular reversible jump MCMC (RJMCMC). RJMCMC was first introduced by Green (1995) and can be viewed as a Metropolis-Hastings (MH) (Hastings (1970)) method adapted to varying dimensional parameter spaces. In our example this corresponds to the addition and deletion of a population center or an individual component.

Briefly, we propose to add a new population center or an individual component, each with probability 0.5 and propose to delete a population center or individual component, each with probability 0.5 at each iteration of the algorithm. We over-sample the RJMCMC moves three times per iteration which results in better mixing. When we propose to add an individual component the parameters defining the new component are drawn from their prior distributions. The use of the prior distributions in proposing new values leads to a simplification of the acceptance probability. When we propose to add a population center, a new location $\boldsymbol{\mu}_*$ and variance τ_*^2 are drawn from their respective priors. A new mixing weight ψ_* is drawn from Beta$(1, c_p)$ and its kernel cancels the Jacobian of the transformation. We re-scale the old weights $\psi' s$ according to $\psi' = \psi(1 - \psi_*)$ such that all weights sum up to 1. The deletion move is the inverse of this construction. Conditional on the number of population centers and individual components other parameters are updated using a Gibbs or a random walk MH step. The variances in the proposal distribution for the MH steps were chosen to obtain acceptance rates of approximately 35%. Following Fernandez and Green (2002) we use the marginal expression for the likelihood and for the priors of the $\boldsymbol{\eta}_{jl}$, as specified in equations (1) and (2) to obtain better mixing. Details are given in the Appendix.

4 Summarizing posterior inference

Simultaneous visualization of the joint posterior distribution of all parameters is infeasible. Instead we view the distributions of certain univariate parameters and create images summarizing the posteriors of the various spatial parameters. In this section we review the approaches we have used to understand our posterior and assess model fit.

We create a "Posterior Probability of Activation" image for each subject, the pixel-wise posterior probability that a pixel is activatived. Precisely, we estimate the marginal posterior probability that subject j has pixel v activated, $\Pr(\omega_{jv} > 0 \mid \mathbf{y})$. The Bayesian estimate of this quantity is the marginal posterior mean of $I_{(\omega_{jv}>0)}$, where $I_{(\omega_{jv}>0)} = 0$ if $\omega_{jv} = 0$ and 1 otherwise.

We create an image of the average "Individual Component Posterior", a pixelization of the (2D) posterior predictive density of the location of a new individual component for a new subject. Precisely, we seek to estimate the posterior distribution of $\tilde{\boldsymbol{\eta}}$, the location of a single individual component for a randomly selected subject. At each sweep we compute the marginal distribution (over population centers and subjects) of the individual component locations, $\sum_{i=1}^{c_p} \psi_i \, \phi_2(\mathbf{x}_v; \boldsymbol{\mu}_i, \tau_i^2 \mathbf{I}_2)$ for each v. Averaging this over sweeps creates an estimate of the density of $\tilde{\boldsymbol{\eta}} \mid \boldsymbol{y}$.

The Individual Component Posterior shows the most probable locations of individual components $\boldsymbol{\eta}$ and is most valuable for visualizing the spread of individual component centers about the population centers, as parameterized by τ_i^2. Care must be taken not to over-interpret the relative height of this image, however. As can be seen by the prior for $\boldsymbol{\eta}$ (equation (2)), the mode height in the Individual Component Posterior are affected by three things. First, all things equal, a smaller τ_i^2 means a more concentrated density which will result in a higher mode. Second, a mode's height will grow as more subjects have components associated with that population center; e.g. if component l is associated with center i, then as $\sum_j I_{(z_{jl}=i)}$ grows so will ψ_i, the mixing weight for center i. And lastly, a subject that

requires more components to fit their profile of activation ($\{y_{jv}\}_v$) associated with center i will also raise ψ_i.

We also create three images to characterize the population centers. The "Population Center Location" image is a pixelization of the (2D) posterior rate function (counts per pixel) for $\boldsymbol{\mu}$. We estimate this by computing the 2D histogram of $\{\boldsymbol{\mu}_i\}$ for each iteration and then average this over iterations. Note that this image does not sum to one, but rather sums to the posterior mean number of population centers, $E(c_p \mid \boldsymbol{y})$.

The "Population Center Scale" image shows the standard deviation of individual component locations about population centers. Exactly, this the posterior average of τ conditional on a population center being in pixel v. Note here we have intentionally suppressed the population center index i, as we marginalize over population center.

The "Population Center Prevalence" image shows the fraction of subjects that possess a population center. To precisely define this quantity, let $c_{ij} = \sum_l I_{(z_{jl}=i)}$ be the count of individual components that subject j has for population center i. The count of subjects with center i is then $c_i = \sum_{j=1}^{J} I_{(c_{ij}>0)}$, and the population prevalance is c_i/J. Conditional on a population center being in pixel v, this image is the posterior average of c_i/J after marginalizing over population centers $i = 1, ..., c_p$.

For comparison with other methods we create two other images. A "Classical t-image" is the one-sample t-test on the BOLD effect magnitudes ($\boldsymbol{y_v} = \{y_{jv}\}_j$) at each voxel. To make inference on this result, we compute an image of $-\log_{10}$P-values for a one-sided, one sample t-test. The $-\log_{10}$ transformation makes for easier visualization, creating an image that should be "bright" in pixels with evidence for an non-null effect magnitude.

5 Results

5.1 Simulation results

To assess the performance of the proposed model, we simulated data for 10 subjects. The size of each image was 40×40 pixels. The parameters in the simulation were chosen to include different numbers of individual activation components and to include centers with different mean intensity levels and different sizes. The top row of Figure 1 shows four of the simulated data sets. The variance of inactivated pixels is 10 for the first 6 subjects and 8 for the last 4 subjects. There are three population activation centers. Simulated subjects 1 and 2 possess activations corresponding to all three centers. Subjects 3 through 7 possess activations corresponding to only two of the three population level centers. Subjects 8 through 10 possess only one activation region each. The mean intensities of these three population level centers are 10, 20 and 25. The variance of the relative intensity for activated pixels are, respectively, 1, 4 and 9. We set $r_{jl}^2 = 6$, 1 and 4 for all j for the three different individual activation components $l = 1, 2, 3$ and $\tau_i^2 = 4$, 9 and 1 to obtain different sizes and different spreads for each individual component. The actual locations of the population centers and the individual components are summarized in Figure 2(a). We ran the chain for 10,000 iterations and discarded the first 5,000 as burn-in. The hyperparameter controlling the background probability was set to $m = 4$, corresponding to an a priori probability of 0.8 of belonging to the background (see page 6, Section 2.1). Other hyperparameter values are as given in Section 2.2.

To assess convergence we started at different random initial values and found no appreciable differences in the results. The bottom row of Figure 1 shows the Posterior Probability of Activation images. Figure 2(b) shows the Individual Component Posterior. The locations of the three major modes are close to the true population values. Also, note that the lower left activation has the greatest intensity, reflecting that more individuals have this activation (as can be seen in Figure 2(a)) and the variation is larger compare to the other two cen-

ters. Figure 2(c) shows the Population Center Location image and gives similar information about the location of the 3 population centers as does Figure 2(b). However, note that the variability of the population locations is much smaller than the variability across subjects. In fact, the population level information is quite precise. The posterior distribution of the number of population activation centers, c_p, has a single mode at 3—in agreement with the true number of population centers. For comparison we display in Figure 2(d) the classical t-image and in Figure 2(e) the $-\log_{10}$P-value image.

Our result is both more informative and less noisy. First, our method separately provides inference on (1) location of population activations, (2) the consistency with which subjects express such activations, and (3) the intrasubject spread of activations about the population centers. The classical method can only assess consistency of signal change at each voxel. With no spatial model, it cannot effect the smoothing of our model, nor can it account for intersubject variability in activation location.

5.2 Real data results

We now summarize results of our model on the working memory data. We run the algorithm for 10,000 iterations with a burn-in of 5,000, saving every 5th iteration to summarize results. The algorithm takes approximately 105 minutes to run on a Mac 2.7 GHz PowerPC G5 processor. The algorithm was written in C++. The acceptance rate for the population level birth/death RJMCMC is about 5%. Figure 3 shows source data (y) and the Posterior Probability of Activation for a representative sample of subjects (1, 3, 9 and 20). The results demonstrate our model capturing the focal signals in subjects 1, 3 and 9, but appropriately finding no signal in subject 20.

The main focus of our work is the population level, as shown in Figure 4(a,b,c,d). Figure 4(a) displays the Population Center Location posterior; this image shows evidence of 5 to 7 population centers, localized to a a few pixels. Figure 4(b) shows the Population Center Prevalence image; of the 5 to 7 centers, only 3 are present in more than 50% of the subjects:

The bottom left and right bilateral pair of activations are in the dorsal occiptal lobe, an area involved in visual processing; the upper left activation is in the dorsolateral prefrontal cortex (DLPFC), an area important maintaining "on-line", short-term representations of information. Figure 4(c) shows the Population Center Scale image, which quantifies the spread of individual components in terms of τ. Figure 4(d) shows the Individual Component Posterior, which qualitatively depicts the spread of individual components about population centers. These two previous images show that the dispers! al of individual components is mostly homogeneous ($\tau \approx 2$), with only the relatively rare right frontal center having larger spread.

Figure 4(e,f) display the classical t-image and the $-\log_{10}(p)$ image. While these images similarly illustrate the spread of activation between subjects, they cannot quantify the spatial precision of the results as our method does.

5.3 Sensitivity analysis

We have studied the sensitivity of our approach to the choice of hyperprior specifications. Our main interest is in how the population level results are related to various prior specifications. We studied the sensitivity of the posterior to the five following scenarios:

(a) $m = 19$, $c_j \sim P(5)$, $j = 1, \ldots, J$, $\lambda_p = 5A^{-1}$, $\alpha_\tau = \alpha_\sigma = 3$, $\lambda_\theta \sim N(35, 10^8)$, the default setting.

(b) same as (a) but with different random seeds (different initial values).

(c) same as (a) but with $m = 5\frac{2}{3}$, which corresponds to an a priori probability that a pixel belongs to the background center of approximately 0.85.

(d) same as (a) but with $\lambda_p = 10A^{-1}$ and $c_j \sim P(10)$, $j = 1, \ldots, J$.

(e) same as (a) but with $\alpha_\tau = \alpha_\sigma = 2$, $\lambda_\theta \sim N(70, 10^8)$.

Overall we found no appreciable differences in the population level results (i.e. Figures 4(a) through 4(d)), under all five scenarios. As expected, scenarios (a) and (b) give almost iden-

tical results for both the individual and the population level. As m decreases the probability that a pixel belongs to the background decreases. This encourages more individual components at low intensity areas. Also, the posterior distribution of the population centers is slightly influenced. We observe a shift in the posterior distribution of the number of individual components, c_j, $j = 1, \ldots, J$, and in the distribution of the number of population centers c_p. For example under (a) the mode of the estimated marginal posterior distribution of the number of mixture components for subject 1, c_1, is 6 (32% of iterations). The estimated marginal posterior distribution of the number of population centers, c_p, is highly concentrated about its mode of 8 (79%). However, under (d)! the mode of the estimated marginal posterior distribution of c_1 becomes 7 (29%); the mode of the estimated marginal posterior of c_p remains 8, but is more diffuse (53%). This influence of the Poisson mean on the posterior distribution of the number of components is expected, see, e.g., Green (1995).

6 Future work and discussion

We have described a Bayesian mixture model for fMRI data analysis. The method considers the spatial structure of the signal, which is often ignored in frequentist approaches. Moreover this method extends the current Bayesian spatial modeling literature in two key ways. First, we consider multi-subject data, explictly population centers and the dispersion of individvual's response about those centers. Second, instead of assuming a normal shape model for activation magnitude, we assume a normal shape model for probability of activation and assume homogeneous magnitude with component. We argue that this leads to a more flexible yet still interpretable parameterization.

One of the limitations of our work is that we assumes spherical Gaussians in our spatial mixture models. For future work, we can use elliptical shape activations using general variance-covariance structures for some of the variance parameters. Another limitation is the computational intensity. The present method took about 105 minutes on a MAC 2.7 GHz

PowerPC G5. Although this is not excessive, with larger 3D images the compute time my be excessive. One avenue to explore would be to parallelize the code. With the proliferation of compute clusters, it would be a simple matter to schedule individuals to separate nodes. All population parameters would then be processed at each iteration after the scheduler returns results from the individuals. Another possible avenue would a Variational Bayes approach (Attias, 2000), which would provide approximate posterior means and variances.

The flexibility of our model does make for some interpretive limitations. While most users of fMRI conceive of activation "loci" (x,y,z locations), in our model the population centers are random variables $\boldsymbol{\mu}$ whose distribution can't easily be summarized. We visualize this distribution with the Population Center Location posterior rate function image, inspecting for modes and assessing the spread about modes.

Likewise, at the subject level, the mixture components are not summarized by loci, but by the Individual Component Posterior density image. While this image is valuable for visualizing the intrasubject spread of activation, the intensity of a peak confounds three things: The scale of intersubject spread, the population prevalance of activation, and the rate of individual components per subject. This can be observed by inspecting the prior for $\boldsymbol{\eta}$ (Eqn. (2)). All things equal, a smaller τ_i^2 results in a more concentrated density which will result in a higher mode. Second, the intensity of a mode will grow as more subjects have evidence for that population center; e.g. if component l is associated with center i, then as $\sum_j I_{(z_{jl}=i)}$ grows so will ψ_i, the mixing weight for center i. And lastly, a subject that requires more components to fit their profile of activation ($\{y_{jv}\}_v$) associated with center i will also raise ψ_i.

Finally, a limitation of our model is that we do not have explicit linking of population centers to individual components. A fMRI practitioner would ideally like to point to population loci $\mathcal{X}$ and ask "which subjects have evidence for that loci, and what is the pattern of activation in each subject corresponding to that loci." For any one MCMC sweep this

connection is known through latent variables z_{jl} and ω_{jv}, but cannot be easily summarized over different sweeps as the number and location of population centers may change.

In future work we would like to address these short commings, perhaps by introducing new latent variables that specifically correspond to investigators notions of activation loci. For example, we could define a loci as "nearest center $\boldsymbol{\mu}_i$ to anatomical landmark X". Alternatively, local maxima on the classical t-statistic image could be used to label population centers; once population centers are identifiable, they can be used to track summary measures of interest.

Acknowledgments

This work is part of my Ph.D. dissertation at the University of Michigan. I would like to express deep gratitude to my advisors whose guidance and support were crucial for the successful completion of this project. Lei Xu and Tim Johnson were partially funded by U.S. National Institute of Health contract/grant number PO1 CA087684. Tim Johnson and Tom Nichols were partially funded by U.S. National Institute of Health contract/grant number RO1 MH069326.

References

Attias, H. (2000). *A variational Bayesian framework for graphical models. Advances in Neural Information Processing Systems*. Cambridge, MA: MIT Press.

Fernandez, C. and Green, P. (2002). Modelling spatially correlated data via mixtures:a bayesian approach. *Journal of the Royal Statistical Society. B* **64**, 805–826.

Friston, K. J. and Holmes, A. P. and Poline, J. B. and Grasby, P. J. and Williams, S. C. R. and Frackowiak, R. S. J. and Turner, R. (1995). Analysis of fmri time-series revisited. *NeuroImage* **2**, 45–53.

Genovese, C., Lazar, N., and Nichols, T. (2002). Thresholding of statistical maps in

functional neuroimaging using the false discovery rate. *NeuroImage* **15**, 870–878.

Green, P. J. (1995). Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika* **82**, 711–732.

Hartvig, N. V. (2002). A stochastic geometry model for functional magnetic resonance images. *Board of the Foundation of the Scandinavian Journal of Statistics* **29**, 333–353.

Hartvig, N. V. and Jensen, J. L. (2000). Spatial mixture mixture modelling of fmri data. *Human Brain Mapping* **11**, 233–248.

Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika* **57**, 97–109.

Miglioretti, D. L., McCulloch, C., and Zeger, S. L. (2002). Combining images across multiple subjects: a study of direct cortical electrical interference. *Journal of the American Statistical Association* **97**, 125–135.

Morosan, P. and Rademacher, J. and Schleicher, A. and Amunts, K. and Schormann, T. and Zilles, K. (2001). Human primary auditory cortex: Cytoarchitectonic subdivisions and mapping into a spatial reference system. *NeuroImage* **13**, 684–701.

Nichols, T. and Hayasaka (2003). Controlling the familywise error rate in functional neuroimaging: A comparative review. *Statistical Methods in Medical Research* **12**, 419–446.

Park, D. C. and Welsh, R. C. and Marshuetz, C. and Gutchess, A. H. and Mikels, J. and Polk, T. and Noll, D. C. and Taylor, S. F. (2003). Working memory for complex scenes: Age differences in frontal and hippocampal activations. *Journal of Cognitive Neuroscience* **15**, 1122–1134.

Woolrich, M. W. and Behrens, T. E. J. and Beckmann, C. F. and Smith, S. M. (2005). Mixture models with adaptive spatial regularisation for segmentation with an appli-

cation to fmri data. *IEEE transactions on medical imaging* **24**, 1–11.

Worsley, K.J. and Marrett, S. and Neelin, P. and Vandal, A.C. and Friston, K.J. and Evans, A.C. (1996). A unified statistical approach for determining significant signals in images of cerebral activation. *Human Brain Mapping* **4**, 58–73.

Appendix

Posterior Distribution

By combining the likelihood function and the prior distributions all together the posterior distribution can be written as

$$
\begin{aligned}
& \pi\left(\{\theta_{jl}\},\, \{\boldsymbol{\eta}_{jl}\},\, \{\boldsymbol{\mu}_i\},\, \{\psi_i\},\, \{\sigma^2_{jl}\},\, \{r^2_{jl}\},\, \{\tau^2_i\},\, \{c_j\},\, c_p,\, \beta_r,\, \beta_\tau,\, \beta_\sigma,\, \lambda_\theta,\, \sigma^2_\theta \mid \mathbf{y}\right) \\
\propto\ & f\left(\mathbf{y} \mid \{\theta_{jl}\},\, \{\boldsymbol{\eta}_{jl}\},\, \{\sigma^2_{jl}\},\, \{r^2_{jl}\},\, \{c_j\}\right)\, \pi\left(\{\theta_{jl}\} \mid \{c_j\}\right) \\
& \pi\left(\{\boldsymbol{\eta}_{jl}\} \mid \{\psi_i\},\, \{\boldsymbol{\mu}_i\},\, \{\tau^2_i\},\, \{c_j\}\right) \times \\
& \pi\left(\{\psi_i\} \mid \lambda_\psi,\, c_p\right)\, \pi\left(\{\boldsymbol{\mu}_i\} \mid c_p\right)\, \pi\left(\{\sigma^2_{jl}\} \mid \beta_\sigma,\, \{c_j\}\right)\, \pi\left(\sigma^2_0 \mid \beta_{\sigma_0}\right)\, \pi\left(\{\tau^2_i\} \mid \beta_\tau,\, c_p\right) \times \\
& \pi\left(\{r^2_{jl}\} \mid \alpha_r,\, \beta_r,\, c_p\right)\, \pi\left(\{c_j\}\right)\, \pi\left(c_p\right)\, \pi\left(\beta_r\right) \pi\left(\beta_\tau\right)\, \pi\left(\beta_\sigma\right)\, \pi\left(\lambda_\theta\right)\, \pi\left(\sigma^2_\theta\right) \\
= \ & \prod_{j=1}^{n} \prod_{v=1}^{V} \Bigg(p_{jv0} (2\pi\sigma^2_0)^{-1/2} \exp\left[-0.5(y_{jv} - \theta_0)^2/\sigma^2_0\right] + \\
& \sum_{l=1}^{c_j} p_{jvl} (2\pi\sigma^2_{jl})^{-1/2} \exp\left[-0.5(y_{jv} - \theta_{jl})^2/\sigma^2_{jl}\right] \Bigg) \times \\
& (2\pi)^{-1/2} \exp\left(-0.5\theta^2_0\right) \prod_{j=1}^{n} \prod_{l=1}^{c_j} (2\pi\sigma^2_\theta)^{-1/2} \exp\left[-0.5(\theta_{jl} - \lambda_\theta)^2/\sigma^2_\theta\right]\ I(\theta_{jl} > 0) \times \\
& \prod_{j=1}^{n} \prod_{l=1}^{c_j} \sum_{i=1}^{c_p} \psi_i (2\pi\tau^2_i)^{-1} \exp\left[-0.5(\boldsymbol{\eta}_{jl} - \boldsymbol{\mu}_i)^T (\boldsymbol{\eta}_{jl} - \boldsymbol{\mu}_i)/\tau^2_i\right]\ \times \\
& \frac{\Gamma(c_p \lambda_\psi)}{\Gamma(\lambda_\psi)^{c_p}} \prod_{i=1}^{c_p} \psi_i^{\lambda_\psi - 1} \prod_{i=1}^{c_p} \left[\frac{\beta^3_\tau}{\Gamma(3)} (\tau^2_i)^{-3-1} \exp\left(-\beta_\tau/\tau^2_i\right) \right] \times \\
& \prod_{j=1}^{n} \prod_{l=1}^{c_j} \left[\frac{\beta^{2\pi}_r}{\Gamma(2\pi)} (r^2_{jl})^{-2\pi-1} \exp\left(-\beta_r/r^2_{jl}\right)\ \frac{\beta^3_\sigma}{\Gamma(3)} (\sigma^2_{jl})^{-3-1} \exp\left(-\beta_\sigma/\sigma^2_{jl}\right) \right] \times \\
& \frac{10^{-6}}{\Gamma(10^{-3})} (\sigma^2_0)^{-10^{-3}-1} \exp\left(-10^{-3}/\sigma^2_0\right)\ \frac{10^{-4}}{\Gamma(10^{-2})} \beta_r^{10^{-2}-1} \exp\left(-10^{-2}\beta_r\right) \times \\
& \frac{10^{-4}}{\Gamma(10^{-2})} \beta_\tau^{10^{-2}-1} \exp\left(-10^{-2}\beta_\tau\right)\ \frac{10^{-4}}{\Gamma(10^{-2})} \beta_\sigma^{10^{-2}-1} \exp\left(-10^{-2}\beta_\sigma\right) \times \\
& (2\pi 10^8)^{-1/2} \exp\left(-0.5(\lambda_\theta - 35)^2/10^8\right)\ \frac{10^{-4}}{\Gamma(10^{-2})} (\sigma^2_\theta)^{-10^{-2}-1} \exp\left(-10^{-2}/\sigma^2_\theta\right) \times \\
& [5^{c_p} \exp(-5)/c_p!] \prod_{i=1}^{c_p} I(\boldsymbol{\mu}_i \in \cup_{j=1}^{n} A_j) \prod_{j=1}^{n} \left[5^{c_j} \exp(-5^{c_j})/c_j!\right].
\end{aligned}
$$

Details of the MCMC algorithm

Current parameter values will be referred to by name and proposed parameters will be referred to by name superscripted with a prime ($'$). Details of the MCMC steps are given below. Each step is classified as a reversible jump step (RJ), a Metropolis-Hastings step (MH) or a Gibbs (Gibbs) step.

1. Birth of a population center: RJ.

   A population center birth is proposed with probability 0.5. Suppose there are c_p centers. Generate $\boldsymbol{\mu}'_{c_p+1}$, and $\tau^{2'}_{c_p+1}$ from their prior distributions:

   $$\boldsymbol{\mu}'_{c_p+1} \sim \text{uniform on} \cup_{j=1}^{J} A_j \quad \text{and} \quad \tau^{2'}_{c_p+1} \sim IG(3, \beta_\tau)$$

   Draw $\psi'_{c_p+1} \sim \text{Beta}(1, c_p)$ and re-scale the existing weights using $\psi'_i = \psi_i(1 - \psi'_{c_p+1})$ for $i = 1, \ldots, c_p$. Set $\boldsymbol{\mu}'_i = \boldsymbol{\mu}_i$ and $\tau^{2'}_i = \tau^2_i$ for $i = 1, \ldots, c_p$. Accept the new population center with probability

   $$\min\left\{1, \frac{5}{(c_p+1)^2} \prod_{j=1}^{J} \prod_{l=1}^{c_j} \frac{\sum_{i=1}^{c_p+1} \psi'_i \ \phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}'_i, \tau^{2'}_i \mathbf{I}_2)}{\sum_{i=1}^{c_p} \psi_i \ \phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}_i, \tau^2_i \mathbf{I}_2)}\right\}.$$

2. Death of a population center: RJ.

   The death of a population center is proposed with probability 0.5. Randomly choose a center from the set {1,...,c_p}. Suppose it is j. Set $\psi'_i = \psi_i/(1 - \psi_j)$ for all $i \in \{1, \ldots, c_p\} \setminus \{j\}$ and relabel the indices $1, \ldots, c_p - 1$. Set $\boldsymbol{\mu}'_i = \boldsymbol{\mu}_i$ and $\tau^{2'}_i = \tau^2_i$ for $i = 1, \ldots, c_p - 1$. Accept the death of this population center with probability

   $$\min\left\{1, \frac{{c_p}^2}{5} \prod_{j=1}^{J} \prod_{l=1}^{c_j} \frac{\sum_{i=1}^{c_p-1} \psi'_i \phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}'_i, \tau^{2'}_i \mathbf{I}_2)}{\sum_{i=1}^{c_p} \psi_i \phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}_i, \tau^2_i \mathbf{I}_2)}\right\}.$$

3. Birth of an individual component: RJ.

   For each subject j, a birth of a component is proposed with probability 0.5. Suppose there are c_j components for subject j. Proposed parameters are drawn from their prior distributions:

   $$r^{2'}_{j,c_j+1} \sim IG(2\pi, \beta_r), \quad \sigma^{2'}_{j,c_j+1} \sim IG(3, \beta_\sigma)$$

$$Pr(z'_{j,\,c_j+1} = i) = \psi_i, \quad \left[\boldsymbol{\eta}'_{j,\,c_j+1} \mid z'_{j,\,c_j+1} = i\right] \sim N(\boldsymbol{\mu}_i, \tau_i^2), \quad \theta'_{j,\,c_j+1} \sim trunc_{(0,\infty)}N(\lambda_\theta, \sigma_\theta^2).$$

Calculate p'_{jvl} for all v, l and $j = 1, \ldots, c_j + 1$. The acceptance probability is

$$\min\left\{1, 5(c_j+1)^{-2}(\text{likelihood ratio})\right\}.$$

4. Death of an individual component: RJ.

   For each subject j, a death of a component is proposed with probability 0.5. Suppose there are c_j components for subject j. Randomly choose a component, j, from the set $\{1, ..., c_j\}$. Calculate p'_{jvl} for all v, l and $j = 1, \ldots, c_j - 1$ and relabel the components $1, \ldots, c_j - 1$. The acceptance probability is

$$\min\left\{1, {c_j}^2 5^{-1}(\text{likelihood ratio})\right\}.$$

5. Update $\boldsymbol{\mu}_i$ for all $i = 1, \ldots, c_p$: MH.

   Propose $\boldsymbol{\mu}'_{\boldsymbol{i}}$ from $N(\boldsymbol{\mu}_i, \sigma_\mu^2\mathbf{I}_2)$. Set $\boldsymbol{\mu}'_j = \boldsymbol{\mu}_j$ for all $j \neq i$. Accept with probability

$$\min\left\{1, \frac{I(\boldsymbol{\mu}'_i \in \cap_{j=1}^J A_j)}{I(\boldsymbol{\mu}_i \in \cap_{j=1}^J A_j)} \prod_{j=1}^{J}\prod_{l=1}^{c_j} \frac{\sum_{i=1}^{c_p}\psi_i\phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}'_i, \tau_i^2\mathbf{I}_2)}{\sum_{i=1}^{c_p}\psi_i\phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}_i, \tau_i^2\mathbf{I}_2)}\right\}.$$

6. Update $\boldsymbol{\eta}_{jl}$ for all $j = 1, \ldots, J$ and $l = 1, \ldots, c_j$: MH.

   For subject j, component l, propose $\boldsymbol{\eta}'_{jl} \sim N(\boldsymbol{\eta}_{jl}, {\sigma_\eta}^2\mathbf{I}_2)$. Calculate p'_{jvl}. Set $\boldsymbol{\eta}'_{kl} = \boldsymbol{\eta}_{kl}$ for all $k \neq j$ and all $l = 1, \ldots, c_j$. Accept with probability

$$\min\left\{1, (\text{likelihood ratio}) \times \frac{\sum_{i=1}^{c_p}\psi_i\phi_2(\boldsymbol{\eta}'_{jl}; \boldsymbol{\mu}_i, \tau_i^2\mathbf{I}_2)}{\sum_{i=1}^{c_p}\psi_i\phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}_i, \tau_i^2\mathbf{I}_2)}\right\}.$$

7. Update τ_i^2 for all $i = 1, \ldots, c_p$: Gibbs.

   Let D_i be the set of individual components with $z_{jl} = i$ and let N_{D_i} be the number of individual components in D_i.

   Draw $\tau_i^2 \sim IG(N_{D_i} + 3, 0.5\sum_{jl\in D_i}(\boldsymbol{\eta}_{jl} - \boldsymbol{\mu}_i)^T(\boldsymbol{\eta}_{jl} - \boldsymbol{\mu}_i) + \beta_\tau)$.

8. Update r_{jl}^2 for all $j = 1, \ldots, J$ and $l = 1, \ldots, c_j$: MH.

   For subject j, component l, propose $\log r_{jl}^{2\prime} \sim N(\log r_{jl}^2, \sigma_r^2)$. Calculate p'_{jvl} for all v. Accept with probability

$$\min\left\{1, (r_{jl}^2/r_{jl}^{2\prime})^{2\pi} \exp[\beta_r(r_{jl}^{-2} - r_{jl}^{-2\prime})](\text{likelihood ratio})\right\}.$$

9. Update $\psi_1, \ldots, \psi_{c_p}$: Gibbs.

   Draw ψ_1, ψ_2, ... ,ψ_{c_p} from $D(1 + N_{D_1}, ..., 1 + N_{D_{c_p}})$.

10. Update z_{jl} for all $j = 1, \ldots, J$ and $l = 1, \ldots, c_j$: Gibbs.

    Draw z_{jl} according to $\Pr(z_{jl} = i \mid .) \propto \psi_i \phi_2(\boldsymbol{\eta}_{jl}; \boldsymbol{\mu}_i, \tau_i^2 \mathbf{I}_2)$.

11. Update ω_{jv} for all $j = 1, \ldots, J$ and all pixels v: Gibbs.

    Draw ω_{jv} from $\Pr(\omega_{jv} = l \mid .) \propto p_{jvl} \phi(y_{jv}; \theta_{jl}, \sigma_{jl}^2)$.

12. Update σ_{jl}^2 for all $j = 1, \ldots, J$ and $l = 1, \ldots, c_j$: Gibbs.

    Let G_{jl} denote the set of pixels with $\omega_{jv} = l$ for subject j. Let $N_{G_{jl}}$ be the number of voxels in G_{jl}. Draw $\sigma_{jl}^2 \sim IG(0.5 N_{G_{jl}} + 3, 0.5 \sum_{v \in G_{jl}} (y_{jv} - \theta_{jl})^2 + \beta_\sigma)$.

13. Update σ_0^2: Gibbs.

    Let G_0 denote the set of pixels with $\omega_{jv} = 0$ for subject j.

    Draw $\sigma_0^2 \sim IG\left(0.5 \sum_{j=1}^J N_{G_{j0}} + 10^{-3}, 0.5 \sum_{j=1}^J \sum_{v \in G_{j0}} (y_{jv} - \theta_0)^2 + 10^{-3}\right)$.

14. Update θ_{jl} for all $j = 1, \ldots, J$ and $l = 1, \ldots, c_j$: Gibbs.

    Draw $\theta_{jl} \sim N(mv, v)$, where

$$m = \sigma_\theta^{-2} \lambda_\theta + \sigma_{jl}^{-2} \sum_{v \in G_{jl}} y_{jv} \quad \text{and} \quad v = \left(\sigma_{jl}^{-2} N_{G_{jl}} + \sigma_\theta^{-2}\right)^{-1}.$$

    Accept if $\theta'_{jl} > 0$.

15. Update θ_0: Gibbs.

Draw $\theta_0 \sim N\left(mv, v\right)$, where

$$m = \sigma_0^{-2} \sum_{j=1}^{J} \sum_{v \in G_{j0}} y_{jv} \quad \text{and} \quad v = \left(\sigma_0^{-2} \sum_{j=1}^{J} N_{G_{j0}} + 1 \right)^{-1}.$$

16. Update β_σ: Gibbs.

    Draw $\beta_\sigma \sim G(3 \sum_{j=1}^{J} c_j + 0.01, \sum_{j=1}^{J} \sum_{l=1}^{c_j} \sigma_{jl}^{-2} + 0.01)$.

17. Update β_τ: Gibbs.

    Draw $\beta_\tau \sim G(3c_p + 0.01, \sum_{i=1}^{c_p} \tau_i^{-2} + 0.01)$.

18. Update β_r: Gibbs.

    Draw $\beta_r \sim G(2\pi \sum_{j=1}^{J} c_j + 0.01, \sum_{j=1}^{J} \sum_{l=1}^{c_j} r_{jl}^{-2} + 0.01)$.

19. Update λ_θ: Gibbs.

    Draw $\lambda_\theta \sim N(mv, v)$, where

$$m = \sigma_\theta^{-2} \sum_{j=1}^{J} \sum_{l=1}^{c_j} \theta_{jl} + (35)1e^{-8} \quad \text{and} \quad v = (\sigma_\theta^{-2} \sum_{j=1}^{J} c_j + 1e^{-8})^{-1}.$$

20. Update σ_θ^2: Gibbs.

    Draw $\sigma_\theta^2 \sim IG(0.5 \sum_{j=1}^{J} c_j + 0.01, 0.5 \sum_{j=1}^{J} \sum_{l=1}^{c_j} (\theta_{jl} - \lambda_\theta)^2 + 0.01)$.

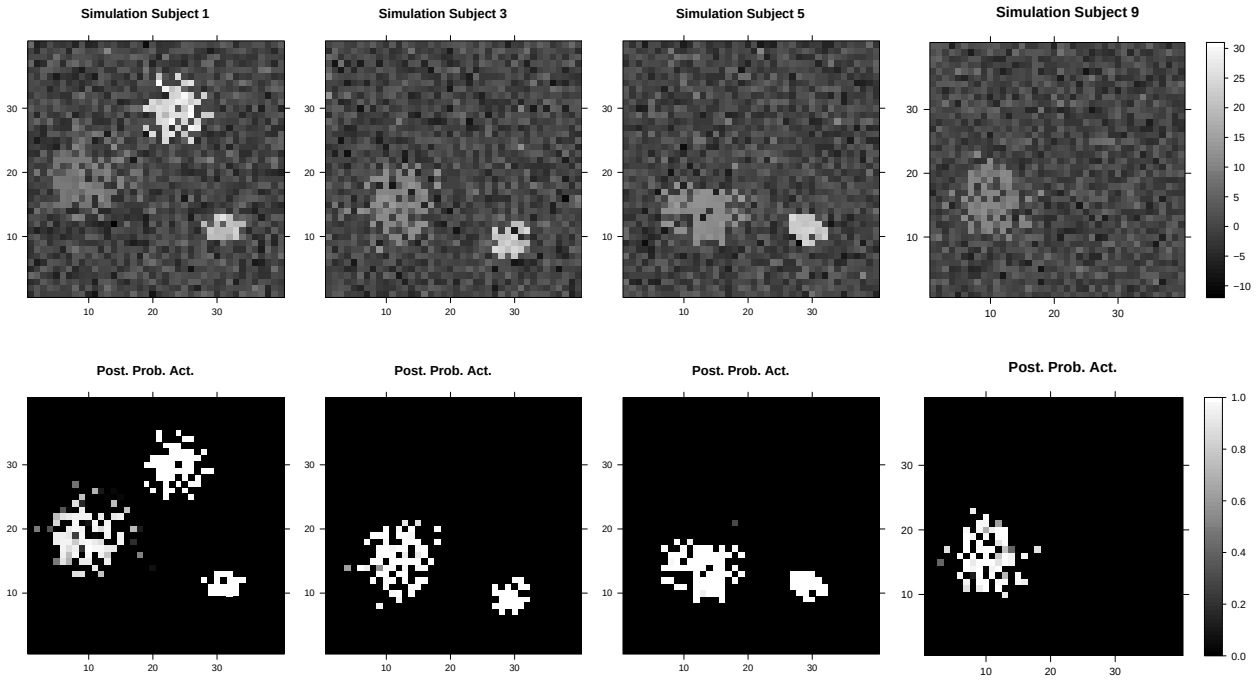

Figure 1: Top row: The intensity data from 4 simulated subjects. Bottom row: The marginal posterior probability of activation: $\Pr(\omega_{jv} > 0 \mid \mathbf{y})$.

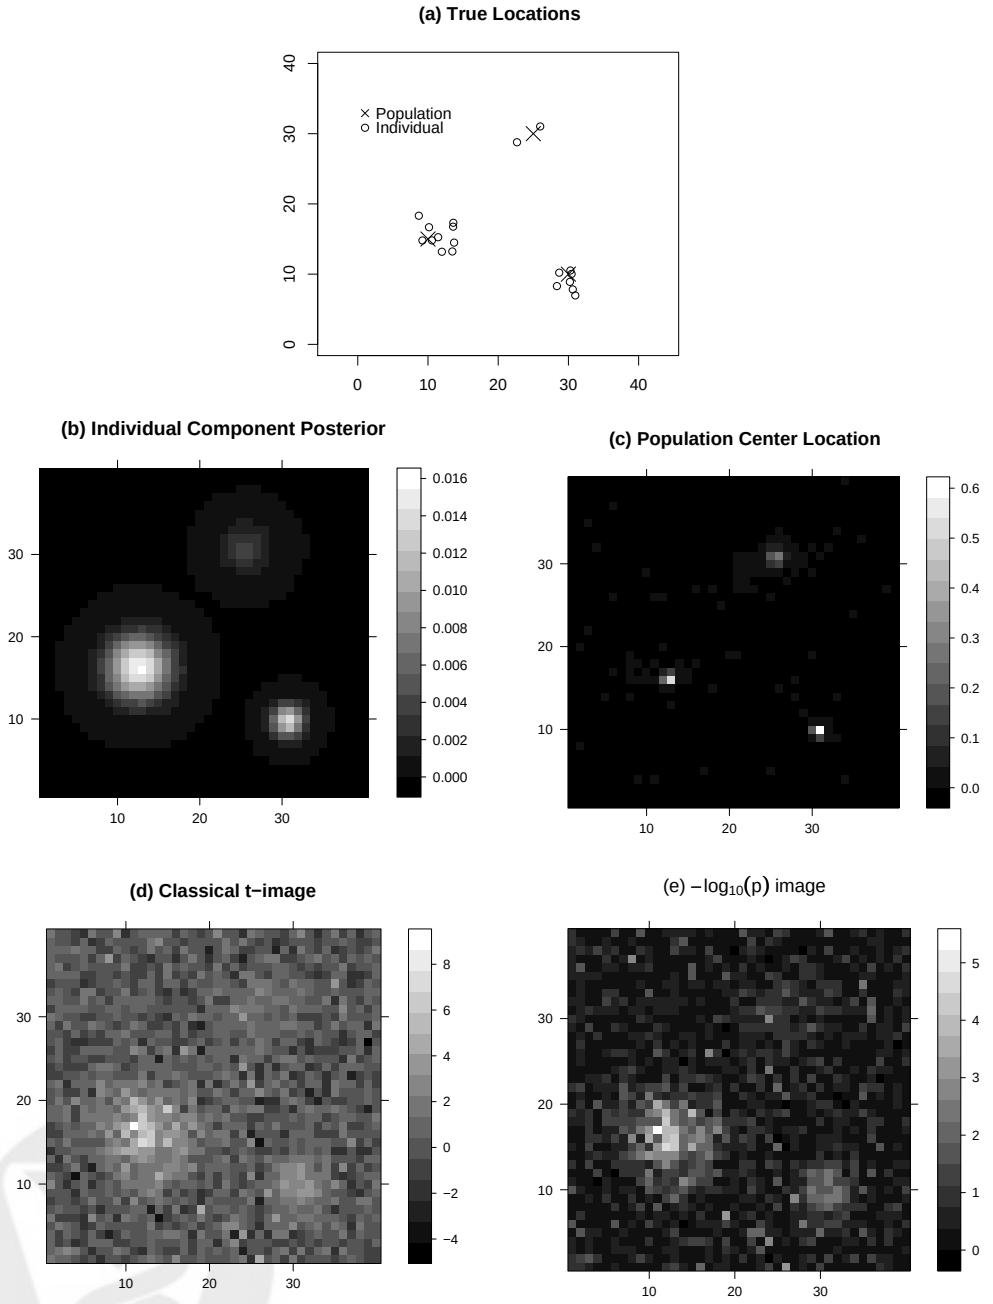

Figure 2: Simulation results. (a) True locations of population and individual locations for simulation. Note that not all subject have centers for each population center. (b) Individual Component Posterior. Note how each component has varying scale, depicting the variation in intrasubject spread of activation about population centers. (c) Population Center Location image. Note the precision of the location of population centers, especially relative to the spread of Individual Component Posterior. (d) Classical one-sample t-test and (e) $-\log_{10}$P-values both show greater noise and poorer localization than our method.

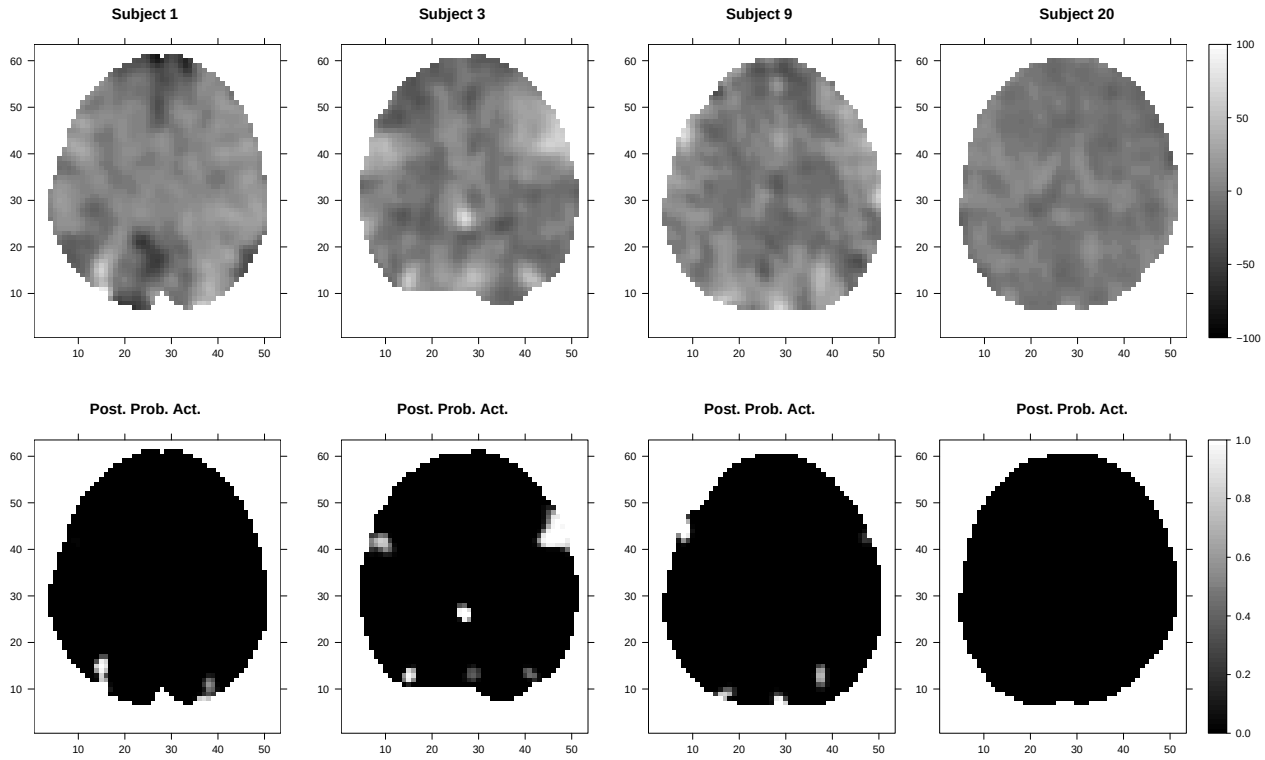

Figure 3: Top row: The intensity data from slice 21 for 4 subjects. Bottom row: The marginal posterior probability of activation: $\Pr(\omega_{jv} > 0 \mid \mathbf{y})$.

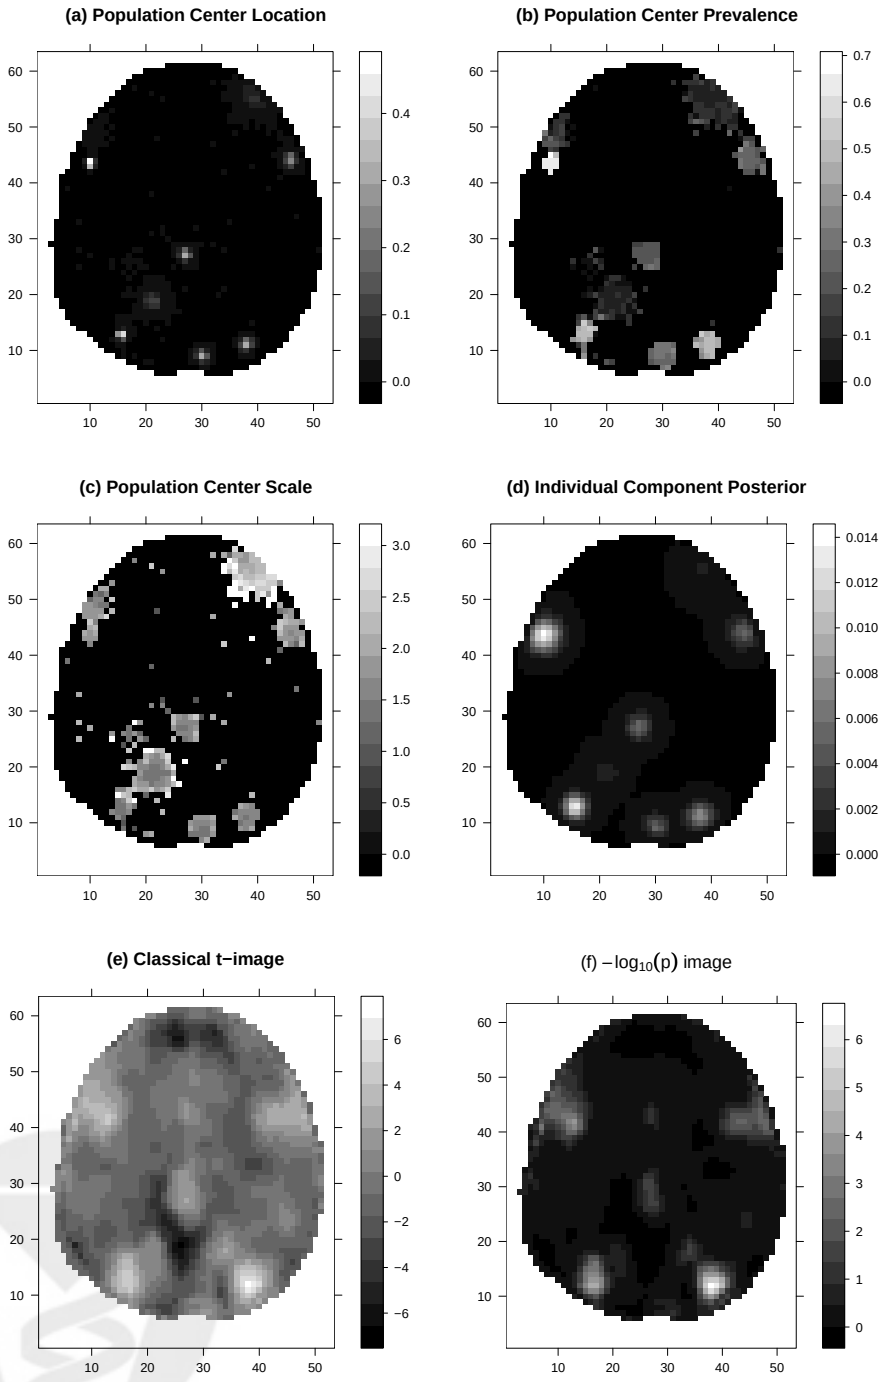

Figure 4: (a) posterior distribution of population center locations; (b) proportion of subjects containing evidence for a population center; (c) standard deviation of individual components about population centers; (d) posterior-predictive density of individual center locations $p(\tilde{\boldsymbol{\eta}} \mid \mathbf{y})$; (e) classical t-image; (f) minus log base 10 p-values from the t-image ($-\log_{10}(p)$).