



Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Combining spatial priors and anatomical information for fMRI detection

Wanmei Ou^{a,*}, William M. Wells III^{a,b}, Polina Golland^a

^a Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, United States

^b Surgical Planning Laboratory, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, United States

ARTICLE INFO

Article history:

Received 26 September 2008

Received in revised form 7 February 2010

Accepted 12 February 2010

Available online 6 March 2010

Keywords:

fMRI detection

Markov Random Field

Mean Field

Variational approximation method

Spatial prior

Anatomical information

GLM

ABSTRACT

In this paper, we analyze Markov Random Field (MRF) as a spatial regularizer in fMRI detection. The low signal-to-noise ratio (SNR) in fMRI images presents a serious challenge for detection algorithms, making regularization necessary to achieve good detection accuracy. Gaussian smoothing, traditionally employed to boost SNR, often produces over-smoothed activation maps. Recently, the use of MRF priors has been suggested as an alternative regularization approach. However, solving for an optimal configuration of the MRF is NP-hard in general. In this work, we investigate fast inference algorithms based on the Mean Field approximation in application to MRF priors for fMRI detection. Furthermore, we propose a novel way to incorporate anatomical information into the MRF-based detection framework and into the traditional smoothing methods. Intuitively speaking, the anatomical evidence increases the likelihood of activation in the gray matter and improves spatial coherency of the resulting activation maps within each tissue type. Validation using the receiver operating characteristic (ROC) analysis and the confusion matrix analysis on simulated data illustrates substantial improvement in detection accuracy using the anatomically guided MRF spatial regularizer. We further demonstrate the potential benefits of the proposed method in real fMRI signals of reduced length. The anatomically guided MRF regularizer enables significant reduction of the scan length while maintaining the quality of the resulting activation maps.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Functional magnetic resonance imaging (fMRI) provides a non-invasive dynamic method for studying brain activation by capturing changes in the blood oxygenation level. In this paper, we focus on intra-subject fMRI activation detection and aim to improve the accuracy of the activation localization in individual subjects. Most fMRI detection methods operate by comparing the time course of each voxel with the experimental protocol, labelling as “active” those voxels whose time courses correlate significantly with the protocol. The commonly used general linear model (GLM) (Friston et al., 1995; Worsley et al., 2002) further models the fMRI signal as an output of a linear system driven by the stimulus. Application of GLM to an fMRI time series results in the so-called *statistical parametric map* (SPM), which is often thresholded to produce a binary map of active areas or, more generally, areas that show different activation levels under different conditions during the experiment. However, because of low signal-to-noise ratio (SNR), the binary maps typically contain many small false positive islands.

One approach to reducing such false positives is to enforce spatial continuity of the data or the estimated activations as a

pre-processing or post-processing step, respectively. Among the pre-processing approaches, smoothing with a Gaussian filter is the most popular method. It effectively increases the SNR of the signal. Unfortunately, Gaussian smoothing can produce overly smoothed SPMs, leading to a loss of details in the resulting binary activation maps. Alternatively, Descombes et al. (1998) employed the Markov Random Field (MRF) model to smooth the data, but their model required large amount of computation due to the application of MRF to the entire 4D fMRI volume. In contrast to smoothing in the spatial domain, Van De Ville et al. (2007) proposed a wavelet-based method to capture the underlying smooth activation pattern. As a post-processing step, Friston et al. (1993) proposed an adaptive thresholding method that assesses the statistical significance of active regions according to their size based on the Gaussian Random Field theory.

An alternative is to integrate a model for spatial coherency of activation into estimation to avoid the above two-step procedure, in which the smoothing step in general may not adequately compensate for the estimation errors. MRFs and Laplacian models are widely used for encouraging spatial continuity. Penny et al. (2005) employed a hierarchical model that imposed a Laplacian spatial prior on the regression coefficients of the GLM detector. Without the hyper prior, their approach effectively smooths the regression coefficients with a fixed spatial kernel that penalizes large magnitudes of the first order spatial derivatives of the

* Corresponding author.

E-mail addresses: wanmei@csail.mit.edu (W. Ou), sw@bwh.harvard.edu (W.M. Wells III), polina@csail.mit.edu (P. Golland).

regression coefficients. The hyper prior helps to obtain appropriate weights for the smoothing. Flandin and Penny (2007) extended this approach to the wavelets domain. Svensen et al. (2000) applied the MRF prior to the parameters of the hemodynamic response functions (HRF) in order to recover the missing activation regions in a fixed-HRF detector. Cosman et al. (2004), Rajapakse and Piyaratna (2001), and Woolrich et al. (2005) employed the MRF model to encourage adjacent locations to share similar activation states, for binary (Cosman et al., 2004; Rajapakse and Piyaratna, 2001) and trinary (Woolrich et al., 2005) activation models.

Motivated by the MRF model in (Cosman et al., 2004), we focus on MRFs for modeling spatial coherence of the activation maps and study their performance in conjunction with a refined version of the MRF model. Despite recent algorithmic development in capturing “continuous” activations, many fundamental neuroscience and surgical applications still employ binary activation maps to make inference about functional areas in the brain. In particular, neuroscientists and neurosurgeons are interested in identifying voxels that are consistently activated in a certain task (Amedi et al., 2001; Ash-tari et al., 2005). Since the fMRI voxels are relatively large, a one-to-two voxel shift often leads to activation localization in a different anatomical region. Therefore, there is still a real need to improve detector accuracy in delineating active/not active regions in an fMRI volume. Similar to previous work (Cosman et al., 2004; Rajapakse and Piyaratna, 2001; Woolrich et al., 2005), we assume that given the activation state of each voxel, the time courses of different voxels are conditionally independent. From the detection point of view, our model directly regularizes the detection results rather than a set of nonlinear HRF parameters implicitly connected to the activation maps of interest (Svensen et al., 2000). More complex temporal models commonly used in fMRI detection, e.g., the autoregressive (AR) models (Burock and Dale, 2000), can be easily integrated into our framework by applying a pre-processing whitening step.

MRFs are known to be much more computationally intensive than linear filters, including the commonly used Gaussian filter. For an MRF with binary states, an exact solution can be obtained in polynomial time (Greig et al., 1989), but it is still too computationally demanding for fMRI detection applications due to a large number of voxels. An fMRI detection algorithm based on the GLM statistic and the binary activation states was demonstrated in (Cosman et al., 2004). If one wants to go beyond binary states (e.g., treating positively and negatively activated voxels differently or including states corresponding to anatomical segmentation), estimation of the optimal activation configuration becomes intractable and approximation algorithms must be used. Prior work in solving the MRF model employed simulated annealing (Descombes et al., 1998; Rajapakse and Piyaratna, 2001), the iterated conditional mode algorithm (Besag, 1974; Salli et al., 2001), and the graph cuts algorithm (Boykov et al., 2001). Recent work by Woolrich et al. (2005) proposed to approximate the MRF using a mixture model. It effectively converts discrete inference into inference in Gaussian Random Fields in order to approximate the partition function. Although the method is computationally efficient, the results depend on the nonlinear logistic transform in the approximation, and the accuracy of the approximated partition function is not quantified. In this work, we adopt the Mean Field approach, developed in statistical physics (Parisi, 1998), and widely employed in MRF-based image processing applications (Kapur et al., 1998; Langan et al., 1992; Pohl et al., 2002). We note that Svensen et al. (2000) also used Mean Field in application of MRFs to spatial regularization of the HRF shape.

We further refine the activation prior by incorporating anatomical information. Similar to the atlas-based segmentation that employs spatially varying priors on tissue types, the anatomy can provide prior information on the activation map. Intuitively speaking, we want the prior to reflect the fact that activation is more likely to occur in gray matter than in white matter, and is not at all likely in

cerebrospinal fluid or bone. In addition, we expect the activation to be spatially coherent within each tissue type but not across tissue boundaries. We use a hidden state variable to encode both the tissue type and the activation state. Segmentation provides an additional, potentially noisy, observation at each node. We derive the detection algorithm for this model and evaluate it on simulated and real data, achieving high detection accuracy with significantly shorter time courses when compared to the standard GLM detector.

Anatomical scans have certainly been used in fMRI analysis and visualization before. Hartvig (2002) used anatomical information in his marked point process spatial prior. Penny et al. (2007) applied a tissue-type prior to explain the observed spatial variability in the temporal AR model. Mapping fMRI data to the cortical surface extracted from an anatomical scan, followed by activation detection on the surface, has been shown to achieve robust activation detection (Andrade et al., 2001; Kiebel et al., 2000). Moreover, some systems, such as BrainVoyager and FSL, employ topologically correct representation of the neocortex for fMRI analysis and visualization (Dale et al., 1999; Goebel et al., 2006; Woolrich et al., 2001). In contrast, our approach does not require a surface extraction algorithm but instead utilizes anatomical information to inject coherency bias into the detection algorithm while performing the computation directly on the volumetric data.

Similar to previous methods that rely on anatomical scans to provide guidance and visualization for fMRI detection, our technique depends on registration of fMRI data to the anatomical scan of the same subject. T2*-weighted echo-planar imaging (EPI) used in fMRI typically suffers from signal distortion or drop-out in certain areas due to magnetic field inhomogeneities. Recent work in parallel imaging acquisition promises to reduce such artifacts (Lin et al., 2005). Moreover, artifact correction and EPI unwarping based on field maps (Jezzard and Balaban, 1995) improve the quality of the fMRI alignment with anatomical MRI. In our real fMRI experiments, field map based unwarping was included.

Lack of ground truth activation in real fMRI data presents a serious challenge for validation of detection methods. Previously proposed methods for evaluating fMRI detectors (Genovese et al., 1997; Liou et al., 2006; Strother et al., 2002) focus on quantifying repeatability of the results on a large set of repeated trials. However, our model explicitly violates the voxel independence assumption required by these evaluation methods. Instead, we evaluate the detection results by comparing activation maps based on reduced-length time courses with *pseudo ground truth* activation maps created from full length time courses over multiple runs. This effectively evaluates how the reduction in observations degrades the ability of the method to detect activations.

This paper extends our preliminary results reported in (Ou et al., 2005). Here we present an alternative formulation of the observation model in order to make the MRF model more transparent and refine the procedure for estimating the model parameters, leading to higher detection accuracy. In addition, we include an in-depth discussion on estimating the model parameters.

In the next section, we briefly review GLM, MRF, and the Mean Field algorithm. In Section 3, we demonstrate how to construct the likelihood term, closely related to the GLM framework, for the MRF model. Section 4 describes our MRF prior model for fMRI detection, and Section 5 extends the MRF model to include anatomical information. Section 6 reports experimental results on synthetic and real fMRI data sets, followed by conclusions and the discussion of future work in Section 7.

2. Background and notation

In this section, we introduce notation and briefly review the necessary background on GLM detection, the MRF models,

and the Mean Field approach to approximate inference on MRFs.

Throughout the paper, we use bold face to denote vectors in time, and “ \rightarrow ” to denote vectors in space. We let random variable $\bar{X} = [X_1, \dots, X_N]$ represent an activation configuration of all N voxels in the volume, and let vector $\vec{x} = [x_1, \dots, x_N]$ be one possible configuration i.e., an activation map. In binary detection, the random variable X_i , which represents the activation state of voxel i , is also binary. We use random variable \mathbf{Y}_i to denote the time course of voxel i ($i = 1, \dots, N$). We let vector $\mathbf{y}_i \in \mathbb{R}^T$ be the time course observed for voxel i in the fMRI scan, where T is the number of time samples in the scan. Given the fMRI scan $\vec{\mathbf{y}} = [\mathbf{y}_1, \dots, \mathbf{y}_N]$, our goal is to produce an activation map \vec{x}^* that represents the best estimate of the true activation map.

2.1. General linear model (GLM)

The GLM detector represents a time course as a linear combination of the protocol-dependent component B , and the protocol-independent component A , such as the cardiopulmonary contributions to the fMRI signal. For our purposes, it is convenient to represent activation detection as a binary hypothesis test for the presence of the protocol-dependent component B :

$$H_0: \mathbf{Y}_i = A\boldsymbol{\alpha}_i + \boldsymbol{\epsilon}_i \quad (1)$$

$$H_1: \mathbf{Y}_i = A\boldsymbol{\alpha}_i + B\boldsymbol{\beta}_i + \boldsymbol{\epsilon}_i$$

for $i = 1, \dots, N$, and $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}, \sigma_i^2 I)$. Appendix A shows that this formulation is equivalent to the more commonly used formulation of GLM that employs a contrast matrix to construct the test statistic. Matrix $C = [AB]$ is referred to as the design matrix. The least squares estimates of the activation response $\boldsymbol{\beta}_i$ and the protocol-independent factor $\boldsymbol{\alpha}_i$

$$\begin{bmatrix} \hat{\boldsymbol{\alpha}}_i \\ \hat{\boldsymbol{\beta}}_i \end{bmatrix} = (C^T C)^{-1} C^T \mathbf{y}_i \quad (2)$$

and the covariance of the estimates $\hat{\boldsymbol{\beta}}_i$ are used to form the corresponding F -statistic for rejecting the null hypothesis, $F_i = \frac{1}{N_\beta} \hat{\boldsymbol{\beta}}_i^T \hat{\Sigma}_i^{-1} \hat{\boldsymbol{\beta}}_i$ with $(N_\beta, T - \text{rank}(C))$ degrees of freedom, where N_β is the number of regression coefficients in $\boldsymbol{\beta}_i$. A detailed discussion on the GLM framework can be found in (Friston et al., 1995; Worsley et al., 2002).

Given an fMRI scan $[\mathbf{y}_1, \dots, \mathbf{y}_N]$, the GLM detector produces an activation map \vec{x}^* by thresholding the statistic value for each voxel at a user-specified level. Some variants of the GLM framework employ the Z -scores or the t -statistic instead. While the magnitude of the statistic could provide additional information about the nature of activation, the statistical framework above leads to a binary answer that is a map of all voxels whose activation scores passed the threshold for statistical significance. In this work, we adopt this commonly used convention of modeling activation maps as binary labels that contain voxels that are modulated significantly by the experimental protocol.

2.2. Markov Random Fields (MRFs)

MRFs are widely used in computer vision (Boykov et al., 2001; Freeman et al., 2000) as a prior for coherency of the underlying structure. For illustration purposes, we couple the review of MRFs with the concrete example of fMRI detection.

In this work, we impose the MRF prior on the hidden activation configuration \bar{X} :

$$P_{\bar{X}}(\vec{x}) = \frac{1}{Z} \prod_{(i,j)} \Psi_{ij}(x_i, x_j) \prod_i \Phi_i(x_i), \quad (3)$$




Fig. 1. Graphical model for $P_{\bar{X}|\vec{\mathbf{y}}}$. In fMRI detection, X_i is the hidden activation label, and \mathbf{Y}_i is the observed fMRI time course.

where the singleton potential $\Phi_i(x_i)$ provides bias over state values x_i for voxel i , and the pairwise potential $\Psi_{ij}(x_i, x_j)$ models the compatibility of a pair of neighboring voxels. Z is the partition function. Here we consider two adjacent voxels as a clique of two, i.e., the first product in Eq. (3) is restricted to adjacent voxels (i, j) . The Ising model is an MRF with binary state variables X_i 's; the Potts model is an MRF with discrete state variables distributed over a finite set of values (Wu, 1982).

Given the time courses of all voxels $\vec{\mathbf{y}} = [\mathbf{y}_1, \dots, \mathbf{y}_N]$, we seek the maximum a posteriori (MAP) estimate of the activation configuration:

$$\begin{aligned} \vec{x}^* &= \arg \max_{\vec{x}} P_{\bar{X}|\vec{\mathbf{y}}}(\vec{x}|\vec{\mathbf{y}}) = \arg \max_{\vec{x}} P_{\bar{X}}(\vec{x}) P_{\vec{\mathbf{y}}|\bar{X}}(\vec{\mathbf{y}}|\vec{x}) \\ &= \arg \max_{\vec{x}} \frac{1}{Z} \prod_{(i,j)} \Psi_{ij}(x_i, x_j) \prod_i \Phi_i(x_i) P_{\mathbf{Y}_i|X_i}(\mathbf{y}_i|x_i), \end{aligned} \quad (4)$$

where $P_{\mathbf{Y}_i|X_i}$ is the probability of the fMRI signal given the activation state of the voxel. The last equality is based on the assumption that the observations at different voxels are independent given the activation state of each voxel. Fig. 1 shows the corresponding graphical model, using a 2D grid for illustration purposes only. In all experiments reported in this paper, we perform the estimation fully in 3D, with six immediately adjacent voxels as neighbors. The observation model $P_{\mathbf{Y}_i|X_i}$ and the potentials Φ_i and Ψ_{ij} fully define the MRF model. The specifics of the model depend on the application of interest. This model explicitly introduces dependencies around the time courses of different voxels \mathbf{Y}_i 's through the hidden activation state \bar{X} . We describe the details of MRF construction for fMRI detection in Sections 3 and 4.

Direct search for the optimal activation configuration is intractable in general. A polynomial-time algorithm for exact MAP estimation exists for binary MRFs (Greig et al., 1989), based on a reduction to the Minimum-Cut-Maximum-Flow problem. We refer to this exact solver as Min-Max. Min-Max is still computationally intensive when applied to the volumetric data. Our experimental results in Section 6.2 show that the Mean Field solver provides a close approximation to the Min-Max solutions but requires much less computation.

2.3. The Mean Field solver

The Mean Field algorithm can be derived in a number of ways. In Appendix B, we provide a detailed development based on the variational approach of approximating the posterior probability of configuration \vec{x} given observation $\vec{\mathbf{y}}$, $P_{\bar{X}|\vec{\mathbf{y}}}(\vec{x}|\vec{\mathbf{y}})$, by a product distribution $Q(\vec{x}) = \prod_{i=1}^N b_i(x_i)$ through minimization of the KL-divergence between the two distributions (Jaakkola, 2000). The optimization leads to a fixed-point iterative update rule

$$b_i^{t+1}(x_i) \leftarrow \gamma P_{\mathbf{Y}_i|X_i}(\mathbf{y}_i|x_i) \Phi_i(x_i) e^{2 \sum_{j \in \mathcal{N}(i)} \sum_{x_j=0}^{M-1} b_j^t(x_j) \log \Psi_{ij}(x_i, x_j)} \quad (5)$$

where $\mathcal{N}(i)$ is the set of voxel i 's neighbors, and M is the number of possible values the state variables X_i can take. The normalization constant γ ensures the solution is a valid probability distribution (it sums to 1). At iteration $t + 1$, the voxel's belief is updated according to a linear combination of its neighbors' beliefs at iteration t . The probability model (i.e., $P_{\mathbf{Y}_i|X_i}$, Φ_i , and Ψ_{ij}) determines the exact form of the update rule. We stop the iterations when the belief for each voxel changes less than a pre-selected threshold.

The resulting approximated posterior probability of activation has two interpretations. First, if we are interested in the MAP solution, we can assign each voxel the state value with the highest belief (for binary MRFs, voxel i is set active if $b_i(1) > b_i(0)$). Alternatively, we can use the belief b_i as a detection statistic. We follow the latter strategy in the ROC analysis in Section 6.

3. Observation model

In this section, we construct the observation model $P_{\mathbf{Y}_i|X_i}$ that is consistent with the GLM framework for fMRI detection.

As demonstrated in (Cosman et al., 2004), the GLM detection can be equivalently formulated as a Generalized Likelihood Ratio (GLR) test:

$$\ell = \frac{\max_{\boldsymbol{\alpha}_i, \beta_i, \sigma_i^2} \mathcal{N}(\mathbf{y}_i; B\boldsymbol{\beta}_i + A\boldsymbol{\alpha}_i, \sigma_i^2 I)}{\max_{\boldsymbol{\alpha}_i, \sigma_i^2} \mathcal{N}(\mathbf{y}_i; A\boldsymbol{\alpha}_i, \sigma_i^2 I)}. \quad (6)$$

It is easy to show that under the null hypothesis ℓ is a monotonic function of the F -statistic under the white noise assumption (Rencher, 2002). If the noise has an unknown covariance, there is no longer an equivalence between the F -statistic and the GLR test. Moreover, neither the numerator nor the denominator in Eq. (6) is a proper probability distribution function since they violate the normalization axiom. But since the maximum likelihood (ML) estimators asymptotically converge to the true values of the parameters under the correct hypothesis, $\max_{\boldsymbol{\alpha}_i, \beta_i, \sigma_i^2} \mathcal{N}(\mathbf{y}_i; B\boldsymbol{\beta}_i + A\boldsymbol{\alpha}_i, \sigma_i^2 I)$ and $\max_{\boldsymbol{\alpha}_i, \sigma_i^2} \mathcal{N}(\mathbf{y}_i; A\boldsymbol{\alpha}_i, \sigma_i^2 I)$ are asymptotically equivalent to the likelihood terms $P_{\mathbf{Y}_i|X_i}(\mathbf{y}_i|X_i = 1)$ and $P_{\mathbf{Y}_i|X_i}(\mathbf{y}_i|X_i = 0)$, respectively.

To augment the GLM detector with an MRF prior, we choose to approximate the observation model with the ML estimates of the corresponding parameters. We set

$$P_{\mathbf{Y}_i|X_i}(\mathbf{y}|X = 1) = \mathcal{N}(\mathbf{y}; B\hat{\boldsymbol{\beta}}_i + A\hat{\boldsymbol{\alpha}}_i, \hat{\sigma}_i^2 I), \quad (7)$$

$$P_{\mathbf{Y}_i|X_i}(\mathbf{y}|X = 0) = \mathcal{N}(\mathbf{y}; A\hat{\boldsymbol{\alpha}}_i, \hat{\sigma}_i^2 I), \quad (8)$$

where $\hat{\boldsymbol{\alpha}}_i$, $\hat{\boldsymbol{\beta}}_i$, and $\hat{\sigma}_i$ are the ML estimates of the model parameters for voxel i . Since $P_{\mathbf{Y}_i|X_i}(\mathbf{y}|X = 1)$ is greater than $P_{\mathbf{Y}_i|X_i}(\mathbf{y}|X = 0)$ in the current likelihood setting, we design the prior on X_i to counter this bias as shown in the next section. In Appendix C, we also show an approach to constructing an MAP estimation procedure based on a voxel-independent prior, instead of the MRF prior, that yields activation maps identical to those obtained through the standard GLM detection.

4. Prior model

We aim to set the prior model parameters to represent our knowledge about the statistical properties of the *true* activation maps. In this work, we employ a spatially stationary prior by using the same potential functions Φ_i and Ψ_{ij} for all voxels in the volume. The observations (fMRI signals and anatomical information, if available) move the MAP estimate away from the configurations that are favored by the spatially stationary distribution. In this stationary model, the subscripts of the potential functions can be removed.

Setting the prior model parameters is a major challenge in MRF-based image analysis applications. Since no ground truth for fMRI

detection is available, we choose to learn the model parameters from the data itself. While this approach represents an approximation for the optimal parameter setting, we found that it works reasonably well in practice. We use the results of the voxel-wise detector at a pre-specified threshold on activation values to estimate the expected number of active voxels in the volume. More specifically, we compute the GLR statistic in Eq. (6) and threshold it to obtain an initial estimate of the activation map. We then use the appropriate voxel counts to estimate the model parameters for the MRF prior as described below.

We set the singleton potential $\Phi(a)$ to the expected percentage of voxels in state a . For example, in the binary activation scenario

$$\Phi(a) = \frac{1}{N} \sum_{j=1}^N \delta(\tilde{x}_j - a), \quad (9)$$

where $\tilde{x}_j = 1$ if the GLR statistic of Eq. (6) at voxel j exceeds the detection threshold and $\tilde{x}_j = 0$ otherwise. This effectively incorporates the detection threshold into the prior on the activation frequency in the volume.

The pairwise potential function $\Psi(a, b)$ should encourage the neighboring voxels to share the activation state. Ideally, we should choose a value for the pairwise potential function in such a way that samples drawn from the resulting prior model agree with our notion of the spatial properties of true activation maps in fMRI experiments. However, the relationship between the potential functions and the statistical properties of samples drawn from the model is a complex one. In our earlier work (Ou et al., 2005), we used the frequency counts to set the pairwise potentials:

$$\Psi(a, b) = \frac{1}{2(\text{Total number of edges in the image})} \sum_{(ij)} \delta(\tilde{x}_i - a) \delta(\tilde{x}_j - b), \quad (10)$$

where an edge connects a pair of adjacent voxels in a 6-voxel neighborhood.

We observed in later experiments that it is easy to generate undesired configurations whose posterior probability, $P_{\tilde{\mathbf{X}}|\tilde{\mathbf{Y}}}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}})$, is greater than that of a desired configuration under this model. A relatively large value of $\Psi(a, b)$ indicates that co-occurrence of the corresponding states a and b is more likely than those with smaller values of the pairwise potential function. A typical Ψ estimated from our real fMRI data is around $\begin{bmatrix} 0.96 & 0.04 \\ 0.04 & 0.96 \end{bmatrix}$. However, the potential functions estimated from noisy frequency counts can over- or under-emphasize the relative likelihood of certain configurations. To enable training of the model, we add a modeling parameter λ that controls the overall ‘‘sharpness’’ of the joint component of the prior, replacing $\Psi(a, b)$ with $\tilde{\Psi}(a, b) = \Psi^\lambda(a, b)$. Varying this positive parameter λ controls the degree to which more favorable configurations are emphasized, but does not alter their ranking order. This parameter is important especially when working with noisy data because it can alleviate the effects of inaccurate estimation of Ψ .

In synthetic experiments, we estimate λ from half of the data by performing an exhaustive search with respect to the true positive rate while fixing the false positive rate at 10^{-3} . Since ground truth activations are not available in real fMRI experiments, we cannot form a training set to estimate λ . In this work, we apply the estimated λ from the synthetic data to activation detection in real fMRI experiments. In order to obtain a sensible estimate of λ , we match the simulation parameters, such as the ratio of gray and white matter volumes and the proportion of active voxels, of the synthetic data with the real one.

The detection performance reported in this paper represents significant improvement over our earlier results (Ou et al., 2005),

which did not include this modeling parameter λ (in other words, it used $\lambda = 1$). The improvement is particularly large for low SNR data. The augmented model also achieves a significant improvement in the detection accuracy for real fMRI data.

5. Anatomically guided detection

The general nature of the MRF framework enables a straightforward extension of the probabilistic model in the previous section to include a tissue type for each voxel. We define $V = [V_1, \dots, V_N]$ to be the tissue types of all voxels, and $W = [W_1, \dots, W_N]$ to be the tissue type observations, e.g., results of automatic segmentation of an MRI scan. Here we model both the hidden tissue type variables V and the observed anatomical labels W on the same grid of the fMRI scan.

Now each voxel has two hidden attributes: the activation state X_i and the tissue type V_i . We combine these attributes into a single hidden node U_i , as illustrated in Fig. 2. For example, for binary activation states (active or not active) and three tissue types (gray matter, white matter, or other), U_i has six possible values. Since the segmentation labels W and the fMRI observation Y are obtained from two different images, it is reasonable to assume that they are conditionally independent given the hidden state.

We choose to treat the segmentation results as a noisy observation rather than the true tissue type. Imperfect registration between the fMRI images and the anatomical scan, partial volume effects, and the errors in the segmentation itself contribute to the “noise” in the anatomical information. Among these factors, partial volume has the strongest influence in our experiments because some brain structures have smaller dimensions than an fMRI voxel. For example, since the highly folded cortical sheet is approximately 3 mm thick, an fMRI voxel (2–4 mm) may correspond to multiple tissue types.

A straightforward extension of the model described in Section 2 leads to the MAP estimator for the new hidden variable

$$\vec{u}^* = \arg \max_{\vec{u}} \frac{1}{Z} \prod_{(i,j)} \Psi_{ij}(u_i, u_j) \prod_i \Phi_i(u_i) P_{Y_i|U_i}(\mathbf{y}_i|u_i) P_{W_i|U_i}(w_i|u_i) \quad (11)$$

and the Mean Field update rule

$$b_i^{t+1}(u_i) \leftarrow \gamma P_{W_i|U_i}(w_i|u_i) P_{Y_i|U_i}(\mathbf{y}_i|u_i) \Phi_i(u_i) e^{2 \sum_{j \in \mathcal{N}(i)} \sum_{u_j=0}^{M-1} b_j^t(u_j) \log \Psi_{ij}(u_i, u_j)} \quad (12)$$

This update rule is similar to Eq. (5), with the exception of the extra likelihood term $P_{W_i|U_i}(w_i|u_i)$ for the tissue type observation. The compatibility matrix Ψ_{ij} is $M \times M$, where M is the number of states in U_i .

Similar to the procedures for estimating the likelihood and prior model parameters in Sections 3 and 4, we estimate the MRF model




Fig. 2. Graphical model for $P_{\vec{U}, \vec{Y}, \vec{W}}$.

from thresholded GLR images. We estimate the likelihood term for the noisy segmentation W given the true hidden states from the corresponding counts of tissue labels in the high-resolution segmentations, the coarse resolution subsampled images, and the initially thresholded GLM results. This accounts for the partial volume effects between the anatomical and the functional data. We also note that the inference framework can utilize probabilistic segmentations by incorporating them into the likelihood model.

Our procedure for the construction of the MRF prior is as follows:

- (1) Register the EPI images and the anatomical scan.
- (2) Downsample the anatomical images to fMRI resolution to establish correspondence.
- (3) Run standard GLM and obtain an SPM.
- (4) Construct the probabilistic model for noisy segmentation $P_{W_i|U_i}(w_i|u_i)$ from the segmented anatomical scan and a thresholded SPM obtained in Step (3).
- (5) Estimate $\Phi_i(u_i)$ and $\Psi_{ij}(u_i, u_j)$ from the segmented anatomical scan and the thresholded SPM.

6. Experiments

In this section, we first quantitatively evaluate the proposed spatial regularizers using synthetic data, followed by a comparison between the detection results computed using the Min–Max and the Mean Field solvers. We then present the experimental results using real fMRI data from the functional Biomedical Informatics Research Network (fBIRN) (Zou et al., 2005).

To compare the performance of the anatomically guided MRF detection with the traditional methods, we also implemented two types of anatomically driven extensions to the GLM detector and the Gaussian filter. In anatomically guided GLM detection, we suppress activations outside of the gray matter using the anatomical information as guidance. In this case, “soft” masking can also account for mis-registration and errors in segmentation. To incorporate anatomical information into the Gaussian filter, we adjust the weights of the filter based on the tissue types of the voxel’s neighbors. When evaluating the filter at voxel i , we assign higher weights to the neighbors sharing the same segmentation results as voxel i . In this implementation, neighbors sharing the same segmentation results contribute twice as much as those sharing different segmentation results. In the later sections, we will refer to these two methods as GLM with anatomical information and anatomically guided Gaussian smoothing, respectively.

We initialize the MRF by assigning equal beliefs to all states, i.e., $b_i(a) = 1/M$ for all activation states a . We stop the algorithm when the belief for each voxel changes by less than 0.01 between consecutive iterations. In all experiments reported in this paper, the Mean Field solver converged in 10–20 iterations.

6.1. Artificial data

To quantitatively evaluate the performance of the method, we generated phantom data sets by applying the expectation-maximization (EM) segmentation algorithm (Pohl et al., 2004) to a 1 mm isotropic anatomical MRI scan and randomly placing activation areas of variable size (average diameter of 15 mm) in the gray matter. We then downsampled the scan to a typical fMRI resolution, 4 mm isotropic ($64 \times 64 \times 64$). In this synthetic data set, the gray matter voxels represent 10% of the total number of voxels in the volume, and the active voxels represent about 10% of the gray matter voxels. We then created simulated fMRI scans based on a fixed parametric hemodynamic response function (the two-gamma function) (Jezzard et al., 2002), using a block-design protocol. In




Fig. 3. Example signals (blue, light) and their noise-free counterparts (black, dark) at two different SNR levels. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

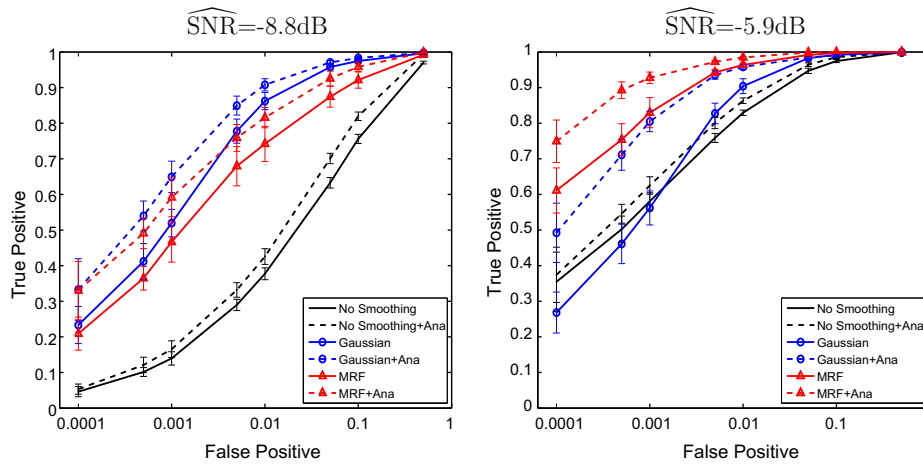


Fig. 4. ROC curves for GLM, GLM with Gaussian smoothing, the MRF-based detector and their anatomically guided variants. False positive rate, shown on logarithmic scale, is computed using the brain voxels only.

this simplified model, we used the same noise level for all voxels in the volume to enable statistical comparison among different methods by averaging the detection rates over the volume.

In our real fMRI studies, the estimated SNR, $\widehat{\text{SNR}} = 10 \log_{10} (\|B\hat{\beta}\|^2 / \hat{\sigma}^2)$, varied from -4 dB to -6.5 dB across sessions and across subjects. We used the estimated SNR to determine an observed level of the simulated noise as the true SNR is unavailable for real fMRI scans. To create realistic simulated signals, we used two levels of noise in these experiments. The resulting noisy signals have the estimated SNR of -5.9 dB and -8.8 dB that bracket the lowest SNR (-6.5 dB) in our real data. We chose such a conservative range of noise levels since the true SNR is typically lower than the estimated one in real fMRI data as part of the noise energy is assigned to the estimated signal during detection when the signal and the noise overlap in some frequency bands. Fig. 3 shows example signals from our simulated data sets.

To decouple testing from phantom generation, we used a GLM detector based on a 10-bin finite impulse response (FIR) hemodynamic response function in testing.

To compare the methods statistically, we generated 16 independent artificial data sets using the procedure mentioned above. We used eight data sets to learn the optimal model parameter λ for the MRF prior, as described in Section 4, and the remaining eight data sets for the evaluation of the methods. In the training phase, we explored the range of values for λ between 1 and 4. Most of our data sets show the best performance when λ is around 3. All 16 data sets were employed to evaluate the GLM detector with and without Gaussian smoothing, as well as their anatomically guided variants. Gaussian kernels of 7 mm full-width-half-maximum (FWHM) achieved the best performance in most simulated data sets when

we varied the kernel width from 4 mm to 12 mm. We therefore report the results for the 7 mm FWHM Gaussian filter.

Fig. 4 illustrates the ROC analysis for GLM, GLM with Gaussian smoothing, and GLM with MRF regularizer (solid curves) and their anatomically guided variants (dashed curves). The ROC curves were created by varying the threshold applied to the GLM statistic or the posterior probability. The error bars indicate the standard deviation of the true positive rate.

Since we expect only a small fraction of voxels in the image volume to be active, a low false positive rate is necessary to prevent false detections from substantially outnumbering true detections (Fig. 4). Here we focus on the false positive rates below 10^{-3} , which correspond to about 10% of the total number of the active voxels, or approximately 250 voxels.

Let us first concentrate on the detection results without anatomical information (solid curves in Fig. 4). As expected, the accuracy of all methods improves with increasing SNR. At high noise levels (low SNR), Gaussian smoothing outperforms MRFs. As the simplest smoothing technique, Gaussian smoothing is more robust to noise. As the SNR increases, the MRF provides better regularization of the activation state. For example, at SNR = -5.9 dB and at the false positive rate of 10^{-3} , the MRF outperforms the Gaussian smoothing by about 30% in true detection accuracy; at 60% true detection, the MRF approximately achieves one tenth of the false detections compared to the Gaussian smoothing. We believe our current estimation procedure for the MRF potential functions Ψ_{ij} and Φ_i partially degrades the MRF performance when SNR is relatively low. As mentioned earlier, these two parameters are estimated from the thresholded GLM activation map. Higher noise levels lead to increasingly inaccurate estimates for the model

parameters and to sub-optimal detection results. However, with ongoing improvements of the imaging technology, we believe MRFs will become even more helpful in reducing spurious false detection islands.

Anatomical information significantly boosts the performance of all detectors at all noise levels (dashed curves). At high noise levels ($\widehat{\text{SNR}} = -8.8$ dB) and false positive rates between 10^{-4} and 10^{-3} (26 and 262 false positive voxels, respectively), both Gaussian smoothing and MRF gain at least 10% in true detection rate while using the anatomical information. At the lower noise level ($\widehat{\text{SNR}} = -5.9$ dB), the basic GLM detector augmented with anatomical information exceeds the performance of the Gaussian smoothing. Since false detections occur relatively uniformly throughout the volume but activations are located in the gray matter, masking out irrelevant regions improves the performance substantially. At the 10^{-4} false positive rate, the anatomically guided MRF outperforms the anatomically guided Gaussian smoothing by about 20% in true detection rate, achieving over 90% detection accuracy.

In addition to the quantitative ROC analysis, we find it useful to visually inspect the resulting activation maps. Fig. 5 shows one axial slice of the activation phantom, the anatomical segmentation, and the detection results for the three methods and their anatomically guided versions at two different SNR levels. All the detections were performed at the false positive rate of 5×10^{-4} . The basic GLM produces a fragmented activation map that shows very little of the original activation at low SNR (Fig. 5c) and contains fragmented activation islands at high SNR (Fig. 5i). Given either of these maps, users would have trouble inferring the true activation areas and disambiguating them from spurious false detections. Gaussian smoothing leads to a reasonable estimate of the ground truth; however, it tends to make the detections “spherical”, which may change the shape of the detected activation areas (Fig. 5e and k). Consequently, many false positive voxels in the Gaussian smoothing occur at the boundaries of the activation regions. Anatomical information reduces this over-smoothing effect for some areas (Fig. 5f and l). At low SNR (-8.8 dB), the MRF model fills in many of the active pixels that were missed by GLM (Fig. 5g). At higher SNR (-5.9 dB), MRF produces relatively accurate results (Fig. 5m). Not all of the scattered activation islands are removed through regularization, but the activation map looks more similar to the ground truth. The activation map is further improved when the anatomical information is integrated into the model (Fig. 5h and n). Similar to our findings from the ROC analysis, at low SNR anatomically guided Gaussian smoothing achieves the best results, while at high SNR the anatomically guided MRF outperforms the rest of the detectors.

6.2. Comparison to the optimal solution

In the binary case, we can evaluate the quality of the Mean Field solution by comparing it to the exact solution. The Min–Max solver mentioned in Section 2.2 runs in polynomial time and finds the activation map that achieves the global maximum of the posterior probability distribution for binary MRFs. We refer the readers to previously published work (Cosman et al., 2004) for implementation details on the Min–Max algorithm in application to MRF-based detection in fMRI.

Fig. 6 shows the ROC curves for the Mean Field and the Min–Max solvers applied to the binary activation data sets introduced in Section 6.1. The Mean Field ROC curve is identical to the corresponding one in Fig. 4. The error bars for the Min–Max ROC curve were computed from the 16 artificial data sets. The Min–Max ROC curve is incomplete for large false positive rates, because extreme threshold values lead to demanding computation (more than 10 h for a single threshold value), and their corresponding false positive rate exceeds the acceptable region in practice. Because each run

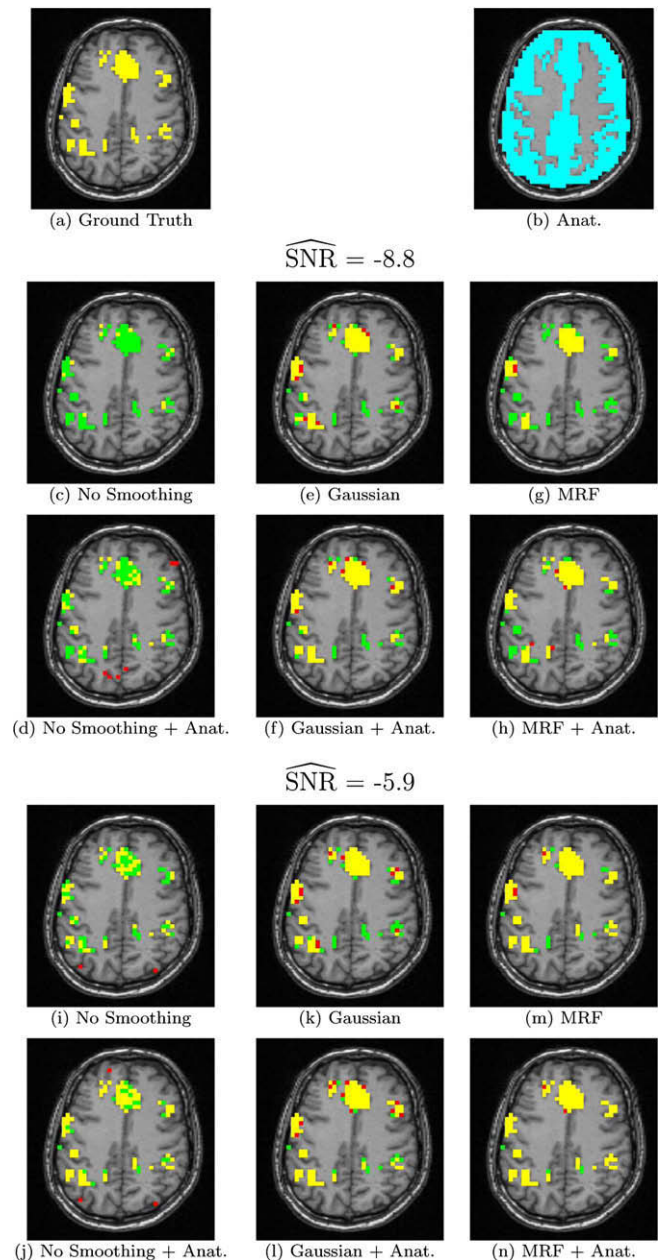


Fig. 5. Axial slice of the ground truth, the segmentation and the estimated activation maps at the false positive rate of 5×10^{-4} . True detections (yellow), false positive (red), and false negative (green) areas are shown. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

takes a long time, we did not optimize the shaping parameter λ for the Min–Max algorithm. Instead, we experimented with different values of λ , and displayed the best performance. As presented in Fig. 6, Min–Max solution is roughly 5% more accurate than the Mean Field solution, suggesting that the Markov model captures the underlying activation pattern reasonably well. On the other hand, Min–Max’s computational requirements are too demanding for the detection application. The Mean Field algorithm offers a good approximation to the globally optimal solution.

The maximum flow in the Min–Max solver can be obtained via Edmonds–Karp algorithm (Cormen et al., 1997). The run time complexity of the algorithm is $O(NE^2)$, where E is the number of edges in the graph. In our case, the number of edges per node in the graph is bounded by a constant, leading to $O(N^3)$ run time complexity.




Fig. 6. ROC curves for the MRF-based detection results using the Mean Field and the Min–Max algorithms. False positive rate is shown on logarithmic scale.

The run time complexity of the Mean Field update is $O(NM)$, where M is the number of state values. In our experiments, the algorithm converged in 10 to 20 iterations. Our Mean Field with Matlab implementation usually converged in 5–10 min on a standard PC (2 GHz CPU and 2 GB RAM), but it was slower than Gaussian smoothing which took less than a second. The Min–Max solver with Matlab implementation took 30–40 min for reasonable thresholds. The Mean Field algorithm requires $O(NM)$ memory, which is comparable to the memory necessary to implement the Gaussian smoother and the standard GLM detector.

We note that the graph cuts algorithm (Boykov et al., 2001) is an extension of the Min–Max solver to multi-state MRF; it employs a series of binary Min–Max steps to achieve a good quality approximation. But since even a single run of a Min–Max solver on a 3D fMRI volume is quite slow, we did not explore the graph cuts algorithms for approximate MRF inference any further.

6.3. Real fMRI experiments

We demonstrate the method in application to a sensory-motor fMRI study. The study included five healthy right-handed subjects. Each subject performed four identical runs of a sensory-motor task. Each run consisted of nine rest epochs and eight task epochs (total 17 epochs), with each epoch lasting 15 s. A flashing checkerboard

and a series of tones were presented at 3 Hz during a task epoch, and the subject was asked to tap along in a regular 3 Hz rhythm. The original purpose of this fMRI study was to characterize cross-scanner variability (Magnotta and Friedman, 2006; Zou et al., 2005), but we only use data from one scanner in the study. The fMRI images were acquired using a Siemens 3T machine (TR = 3 s, $64 \times 64 \times 35$, $3.4 \text{ mm} \times 3.4 \text{ mm} \times 5 \text{ mm}$, dual echo EPI, echos at 20 and 50 ms, EPI bandwidth of 2441 Hz, 90° flip angle, single channel head coil). The estimated SNR in expert-selected ROIs relevant to the task varied between -4 dB and -6.5 dB . Each subject also underwent anatomical scanning (transverse three-dimensional magnetization prepared rapid acquisition gradient-echo protocol, 9.8/minimum, 15° flip angle, $22 \times 16.5 \text{ cm}$ field of view, 256×192 matrix, 124–128 sections, $1 \text{ mm} \times 1 \text{ mm} \times 1.2 \text{ mm}$, $T_1 = 300 \text{ ms}$, bandwidth of 15,625 kHz, two signals acquired).

Since the ground truth activation is unknown in real fMRI experiments, we cannot construct ROC curves. In this work, we evaluated the proposed fMRI analysis methods with respect to their ability to detect activations from reduced-length time courses. Fig. 7 illustrates the flow chart of our evaluation process. We first applied the GLM detector, parameterized by the two-gamma function, without any smoothing to each of the four full-length runs for the same subject. We then employed majority voting among the resulting four activation maps to construct the pseudo ground truth activation map for the experiment. To evaluate the six methods we compared in Section 6.1, we first quantitatively compared each of their activation maps produced with varying length of time courses to the pseudo ground truth activation map (Fig. 8). We then visually inspect the quality of the activation maps shown in Figs. 9–11.

The pseudo ground truth activation map (Figs. 9–11a) exhibits a substantial amount of fragmentation. It often shows bias towards to the GLM detector, especially for the results obtained with relatively long time courses. At the same time, other detectors might produce perceptually more meaningful results. Hence, extra caution is needed while interpreting results based on this pseudo ground truth reference.

In contrast to our experiments with synthetic data, where the activation thresholds were chosen according to the required bound on the false positive rate, in real data we chose to select the top 1% of the voxels with the highest statistic or posterior probability when comparing the activation maps across detectors. Sophisticated threshold selection methods, such as controlling the false discovery rate (Benjamini and Hochberg, 1995) or controlling family-wise error rate, can be applied to statistics generated from the GLM detectors. When we introduce smoothing or spatial prior, the theory of statistical testing becomes substantially more complex. For the purpose of comparing different detectors, our threshold

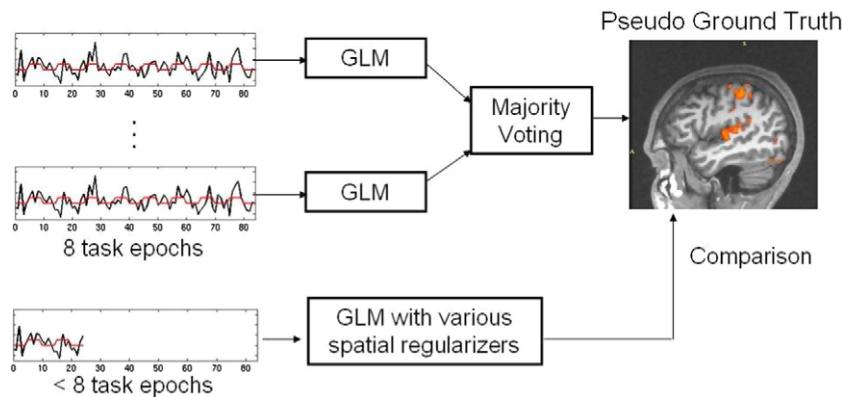


Fig. 7. Our validation procedure for real fMRI data.




Fig. 8. Detection sensitivity with respect to the length of time courses used in the detection. Results shown in Figs. 9–11 correspond to detection using seven epochs. The map obtained through majority voting on four runs was used as “pseudo ground truth”.

selection method effectively reveals the detected voxels and the spatial pattern when equal number of voxels pass the threshold.

Fig. 8 shows the fraction of detected activation matching with those in the pseudo ground truth versus the number of epochs included in the detection. We omit the corresponding figure for the false positive rate as the false positive rate can be straightforwardly derived from Fig. 8 since we select the same number of voxels active in each experiment. We can see that spatial regularization, of either MRF or Gaussian smoothing, is crucial for detection with short time courses. Based on three epochs, GLM achieves 0.4 true positive rate, MRF and Gaussian achieve 0.48 and 0.44, respectively. The results of the MRF-based detectors show a much better agreement with the pseudo ground truth than the baseline methods. Furthermore, anatomical information significantly improves the performance of all detectors, boosting it by additional 5–8%.

We now inspect the activation maps from the six detectors with seven-epoch time courses (out of 17) in Figs. 9–11c–h. Subfigure

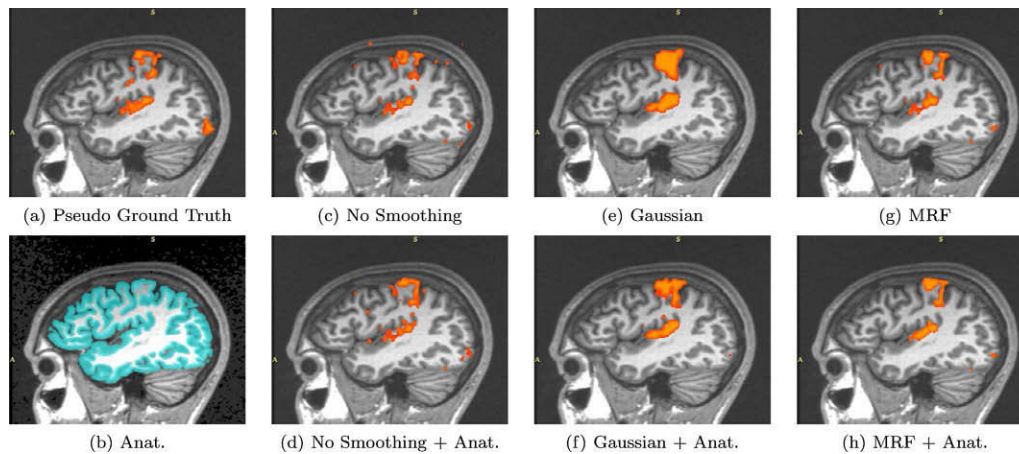


Fig. 9. Real fMRI study. One sagittal slice in the estimated activation maps. (a) Pseudo ground truth created from four full-length sessions (17 epochs each). (b) High-resolution segmentation result overlaid onto the corresponding MR image. (c)–(h) The activation maps obtained from the first seven epochs. (c) Plain GLM detection results. (d) A gray matter mask is employed to suppress GLM-detected activations which do not lie in the gray matter. (e) Isotropic Gaussian smoothing with 7 mm FWHM is applied prior to the GLM detector. (f) A weighted Gaussian smoothing is applied prior to the GLM detector. (g) Generated by the GLM detector with MRF spatial prior. (h) Further incorporates the anatomical information into the MRF prior.

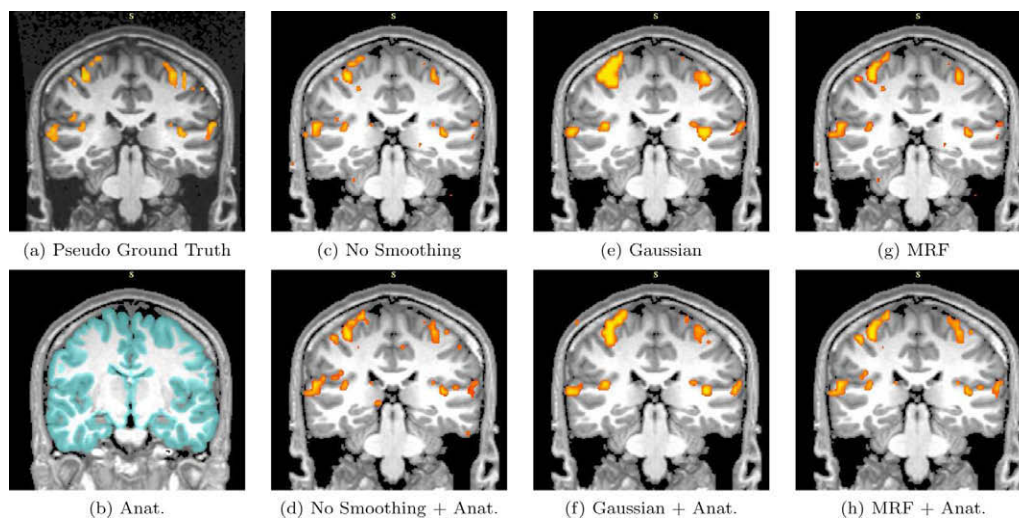


Fig. 10. Real fMRI study. A coronal slice in the estimated activation maps. The description of each subfigure follows the conventions of Fig. 9.

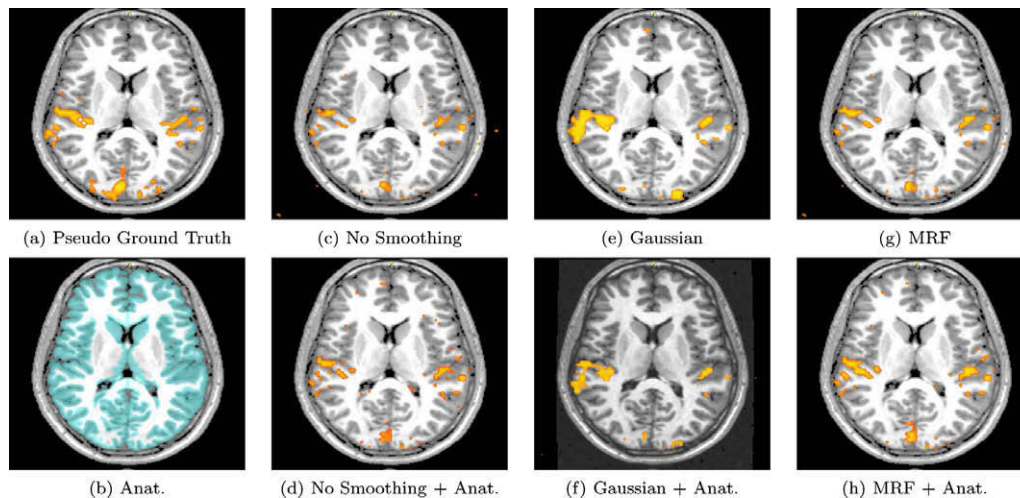


Fig. 11. Real fMRI study. An axial slice in the estimated activation maps. The description of each subfigure follows the conventions of Fig. 9.

(b) in all three figures displays the high-resolution anatomical segmentation. Without spatial regularization, the GLM detector yields more fragmented activation map (c) due to loss in statistical power from reducing the length of the signal. Spatial masking removes some false detections outside of the gray matter in (d). Since we keep the number of active voxels constant, extra gray matter voxels are labelled as active in (d). The other four images illustrate the results of applying GLM with Gaussian smoothing and the MRF-based detector, as well as their anatomically guided versions. Although Gaussian smoothing (kernel width of 7 mm) removes most of the single-voxel activation islands, its activation map in (e) significantly over-smooths the activation areas when compared to (a). Anatomical weighting somewhat reduces the over-smoothing effect. MRF regularization ((g) and (h)) captures activations with elongated spatial structures. This highlights the potential benefit of using the Markov priors in fMRI detection. Furthermore, Fig. 9h illustrates that a coherent activation pattern along the motor cortex is detected by the anatomically guided MRF detector. Similar to the Gaussian smoothing, the MRF model benefits from using anatomical information to remove spurious activations. Our experiments demonstrate that employing anatomically guided MRF spatial regularization leads to accurate detection from time courses of substantially reduced length.

7. Discussion and conclusions

Our experiments confirm the importance of spatial regularization in reducing fragmentation of activation maps. We investigated two types of improvement in spatial modeling for fMRI detection: MRF priors and anatomical bias. The MRF provides a spatial prior that refines the structure of the resulting activation maps over Gaussian smoothing, as demonstrated by our experiments on phantom and real data. We explored the fast Mean Field approximate solver in application to MRF-based fMRI detection and showed that it provides good detection accuracy while taking substantially less time to evaluate than the exact solution. We note that since the Markov model itself is an approximation of the real geometry of the activation regions, we should not dwell on the small differences in the activation maps introduced by the approximate solvers but rather focus on their performance relative to the ground truth or the pseudo ground truth.

We also demonstrated that anatomical information provides helpful bias in fMRI detection. We derived an algorithm for anatomically guided MRF estimation and showed how Gaussian smoothing can be straightforwardly augmented with anatomical

information by rescaling the coefficients of the smoothing kernel. Partial volume effects in reducing the anatomical segmentations to the fMRI resolution should be investigated further in the future. In this work, we downsample the anatomical scan to match the resolution of the functional scan. A better solution would be to use the high-resolution anatomical scans to resolve the activation in the functional voxels that are on the boundary of the gray matter, leading to a “super-resolution” detector.

In this paper, the detection method treats the anatomical segmentation as a noisy observation of the true tissue types. In addition, it utilizes anatomical information to estimate the pairwise and the singleton potential functions Ψ_{ij} and Φ_i . If we treat the segmentation labels as the true tissue types for inference purposes, i.e., employ the graphical model from Fig. 1 with spatially variant potential functions based on the segmentation label at each voxel, the method fails to produce accurate results. We believe that the mis-registration and subsampling noise in producing anatomical data on the same grid with fMRI measurements requires additional modelings and cannot be ignored in anatomically guided fMRI analysis. We can also use the anatomical information for inference only, but employ spatially stationary potential functions that are ignorant of anatomy. The structure of the corresponding graphical model is similar to that in Fig. 2, except that each hidden node contains the activation state X_i only. In this model, the tissue type does not propagate between voxels. The resulting MRF prior does not encourage strong spatial coherency within tissue types and weak spatial coherency across tissue types. In our experiments, we found this method to produce similar results to those of the full model. To summarize, modeling the mismatch in alignment and resolution between the two modalities is crucial for accurate anatomically guided fMRI detection.

It has been suggested that anatomical information can be obtained from breath-hold fMRI experiments (Li et al., 1999), which would remove the effects of mis-registration and resolution mismatch between the functional data and the anatomical structure. The BOLD signal in the gray matter is highly correlated with the breath-hold experimental protocol, allowing gray matter segmentation with an appropriately selected threshold. In addition, a rough gray matter segmentation can be extracted from the AR(1) temporal model of noise (Worsley et al., 2002). In our experiments, we found that the anatomical prior produced by these two approaches in general was of much lower quality than that obtained from anatomical segmentation followed by downsampling. We therefore focused on anatomical segmentation as a source of tissue type information.

We evaluated the methods by performing the ROC analysis on phantom data, and by studying their ability to recover activation from significantly shorter time courses on real data. In the high noise setting for the binary activation states, the Gaussian smoothing outperformed other methods, but the difference between Gaussian and MRF was marginal. As the SNR in the images increased, the Markov prior offered a substantial improvement in the detection accuracy.

Many fMRI studies aim to detect and characterize positive and negative activations separately (Allison et al., 2000; Aron et al., 2004; Seghier et al., 2004). The negatively activated regions have also been shown to correlate with disease-induced neurodegeneration (Buckner et al., 2005; Greicius et al., 2004). Most literature defines positive/negative activations as a comparison of responses to two different stimuli or as a comparison of responses to a stimuli versus baseline. Our model can capture both definitions by extending the activation state variables X_i to three states, i.e., $\{-1, 0, 1\}$. If more than one comparison of stimuli is of neuroscientific interest, our model can be adapted by handling each comparison independently. Recent work by Makni et al. (2008) and Vincent et al. (2007) jointly estimate the positive/negative activations and the shape of HRF through a bilinear model. Their work requires an estimation on a set of pre-defined regions, such as the cortical parcellation from anatomical scans. We have completed preliminary evaluation of the six detectors discussed in Section 6.1 in the context of positive/negative activations (Ou, 2005). The overall performance of these detectors is similar to the results for the binary activation cases, except that the MRF-based detectors show better performance in preserving the boundary between positively and negatively activated regions while Gaussian smoothing tends to average away positive and negative activation signals.

We used block design experiments to validate the method. We expect the MRF regularizer to perform somewhat worse in the event-related paradigms due to lower SNR. In our experiments, we observed that the proposed regularization methods improve detection accuracy at both SNR levels, but the effect is more substantial for higher SNR levels, suggesting that the utility of the method in analyzing real fMRI data will grow as the quality of the data improves.

We also note that localization response curve (LROC) analysis (Swensson, 1996) could explicitly incorporate location information into the ROC analysis for simulated data. In general, it is easy to detect some active voxels in or near the real activation region. However, correctly detecting the boundaries of the activation region is difficult, and it is not captured by the LROC. Our current approach to quantifying the detection performance by averaging the performance at each voxel using a standard ROC approach lacks the localization information, but it serves as a better evaluation tool as improving the boundary detection is one of the main goals of this work.

Evaluation is challenging in real fMRI data since no ground truth is available even for limited experiments. Alternative evaluation methods for fMRI detectors have been proposed, based on repeatability of activation results over multiple sessions (Genovese et al., 1997). By modeling the distribution of detection results for a voxel over all sessions as a mixture of two binomial distributions, this method estimates the ROC curve. The underlying hidden activation state is estimated as part of the mixture model. Unfortunately, this model relies on a large number of independent sessions and on independence among voxels to provide reasonable estimates. Both assumptions may be too optimistic in fMRI experiments, and were certainly violated by our data and methods. Extensions of this procedure address multiple tasks (Liou et al., 2006) and non-parametric models (Strother et al., 2002) and allow characterization of prediction accuracy and the SNR associated with repro-

ducible SPMs. However, the detectors in our comparison produce substantially different types of statistics (F -statistic for the GLM and posterior probability for the MRF-based detectors), presenting additional challenges for such methods. Here we chose to compare the method's ability to detect activations from reduced-length time courses as an evaluation criterion; further work in comparing activation detectors is needed. Maitra et al. extended work in (Genovese et al., 1997) to handle spatially-dependent fMRI detection (Maitra et al., 2002), and it is of our interest to compare their approach with ours in future work.

The anatomically guided Markov spatial prior enabled us to shorten the fMRI scan length from four runs of 17 epochs to nine-or-fewer epochs in one run while still retaining comparable detection power. We expect a similar effect to occur with respect to the spatial resolution when we extend the method to utilize the anatomical information at the original scan resolution. As the quality of the scanning equipment improves, the sophisticated spatial models, such as MRFs, will become even more important in recovering the details of the activation regions.

Acknowledgments

We thank Dr. Eric Cosman and Dr. Kilian Pohl for help on this article and Dr. Douglas Greve for pre-processing the fMRI data. This work was partially supported by the NIH National Center for Biomedical Computing Program, National Alliance for Medical Imaging Computing (NAMIC) Grant U54 EB005149, the NSF IIS 9610249 grant, the NIH NCRR mBIRN U24-RR021382 grant, the NIH NCRR NAC P41-RR13218 grant, the NIH NINDS R01-NS051826 grant and the NSF CAREER 0642971. fMRI acquisition was supported by the NCRR FIRST-BIRN (U24 RR021992) grant. Wanmei Ou is partially supported by the NSF graduate fellowship, and Dr. William M. Wells III is supported by the NCRR FIRST-BIRN grant and the Neuroimaging Analysis Center.

Appendix A. Alternative GLM formulation

Here we show that the formulation of Eq. (1) is equivalent to the more typically employed application of a contrast matrix in conjunction with GLM. GLM models the BOLD response as a linear time-invariant system (Friston et al., 1995):

$$\mathbf{y} = H\boldsymbol{\beta} + \boldsymbol{\epsilon} \quad (\text{A.1})$$

where \mathbf{y} is the voxel's fMRI time course and H is the design matrix, which may contain responses to different tasks or the same task but different types of stimulus. Since GLM models each voxel separately, we omit the voxel index i in the derivations here. Noise $\boldsymbol{\epsilon}$ is assumed to be i.i.d. Gaussian noise. The vector of regression coefficients $\boldsymbol{\beta}$ includes both the coefficients in the signal space and the nuisance parameters. The contrast row vector \mathbf{c} is defined to select the relevant components of the ML estimate of the regression vector $\hat{\boldsymbol{\beta}}$ for statistical testing. For example, we can construct the t -statistic as $t = (\mathbf{c}\hat{\boldsymbol{\beta}}) / \sqrt{\hat{\Sigma}_{\mathbf{c}\hat{\boldsymbol{\beta}}}}$ with $T - \text{rank}(C)$ degrees of freedom, and $\hat{\Sigma}_{\mathbf{c}\hat{\boldsymbol{\beta}}}$ is the estimated variance of $\mathbf{c}\hat{\boldsymbol{\beta}}$.

To show that Eq. (1) is equivalent to this GLM formulation, we introduce transformation matrix T_H . The first row of matrix T_H is identical to the contrast row \mathbf{c} ; the remaining $N_\beta - 1$ rows are constructed such that each row in T_H is orthonormal to the remaining rows. We now consider $\tilde{\boldsymbol{\beta}} = T_H\boldsymbol{\beta}$. The first element of $\tilde{\boldsymbol{\beta}}$ is the projection of the regression vector onto the signal space of interest, defined by \mathbf{c} . The remaining elements of $\tilde{\boldsymbol{\beta}}$ can be considered nuisance parameters. Therefore, we can rewrite Eq. (A.1) in the form of Eq. (1) by setting $\tilde{H} = HT_H^{-1}$:

$$\mathbf{y} = H\boldsymbol{\beta} + \boldsymbol{\epsilon} = \tilde{H}\tilde{\boldsymbol{\beta}} + \boldsymbol{\epsilon} = \tilde{H}^{(1)}\tilde{\boldsymbol{\beta}}^{(1)} + \tilde{H}^{(2,\dots,N_{\#})}\tilde{\boldsymbol{\beta}}^{(2,\dots,N_{\#})} + \boldsymbol{\epsilon}. \quad (\text{A.2})$$

The first term in the resulting sum corresponds to the signal space components. The second term in the sum represents contributions orthogonal to the signal space of interest.

Appendix B. Mean Field derivation

In this appendix, we demonstrate one possible derivation of the Mean Field iteration as a variational approximation of the original MAP problem, closely following the development in (Jaakkola, 2000). The basic idea is to approximate the posterior probability $P_{\vec{x}|\vec{y}}(\vec{x}|\vec{y})$ by a simpler distribution $Q(\vec{x}; \vec{b}) = \prod_{i=1}^N b_i(x_i)$ through minimizing the KL-divergence between them and to use the mode of $Q(\cdot; \vec{b})$ as an approximation for the mode of $P_{\vec{x}|\vec{y}}(\cdot|\vec{y})$. $\vec{b} = [b_1, b_2, \dots, b_N]$ is the vector of belief functions: $b_i(a)$ denotes the probability that voxel i is in state a . Clearly, $\sum_{a=0}^{M-1} b_i(a) = 1$, where M is the number of possible values that the state variable X_i can take. For example, $M = 2$ for binary MRFs. We note that the resulting distribution $Q(\cdot; \vec{b})$ implicitly depends on the observations \vec{y} .

The KL-divergence

$$D(Q||P_{\vec{x}|\vec{y}}) = \sum_{\vec{x}} Q(\vec{x}) \log \left(\frac{Q(\vec{x})}{P_{\vec{x}|\vec{y}}(\vec{x}|\vec{y})} \right) \quad (\text{B.1})$$

serves as a distance between $Q(\cdot; \vec{b})$ and $P_{\vec{x}|\vec{y}}(\cdot|\vec{y})$; it is non-negative and is equal to zero only if $Q = P_{\vec{x}|\vec{y}}$. As an aside, the statistical inference theory implies that we should minimize the KL-divergence $D(P_{\vec{x}|\vec{y}}||Q)$, which often leads to intractable problems. To overcome this difficulty, variational methods are usually formulated as a minimization of the KL-divergence $D(Q||P_{\vec{x}|\vec{y}})$.

It is easy to see that the minimum of $D(Q(\cdot; \vec{b})||P_{\vec{x}|\vec{y}})$ is achieved for the same belief vector \vec{b}^* that minimizes the so-called *free energy*,

$$F_{MF}(\vec{b}) = D(Q||P_{\vec{x}|\vec{y}}) - \log(P_{\vec{y}}(\vec{y})) - \log(Z) \quad (\text{B.2})$$

since the last two terms are independent of \vec{b} . Substituting the definitions for $Q(\vec{x}; \vec{b})$ and for $P_{\vec{x}|\vec{y}}(\vec{x}|\vec{y})$, we obtain,

$$\begin{aligned} F_{MF}(\vec{b}) = & - \sum_i \sum_{j \in \mathcal{V}(i)} \sum_{x_i=0}^{M-1} \sum_{x_j=0}^{M-1} b_i(x_i) b_j(x_j) \log(\Psi_{ij}(x_i, x_j)) \\ & + \sum_i \sum_{x_i=0}^{M-1} b_i(x_i) [\log(b_i(x_i)) - \log(P_{\mathbf{y}_i|X_i}(\mathbf{y}_i|x_i)\Phi_i(x_i))] \end{aligned} \quad (\text{B.3})$$

leading to the following constrained optimization problem:

$$\vec{b}^* = \arg \min_{\vec{b}} F_{MF}(\vec{b}) \quad (\text{B.4})$$

$$\text{s.t. } \sum_{a=0}^{M-1} b_i(a) = 1, \quad \forall i.$$

Using Lagrange multipliers, we reduce the problem above to minimize

$$\begin{aligned} J(\vec{b}) = & - \sum_i \sum_{j \in \mathcal{V}(i)} \sum_{x_i=0}^{M-1} \sum_{x_j=0}^{M-1} b_i(x_i) b_j(x_j) \log(\Psi_{ij}(x_i, x_j)) \\ & + \sum_i \sum_{x_i=0}^{M-1} b_i(x_i) [\log(b_i(x_i)) - \log(P_{\mathbf{y}_i|X_i}(\mathbf{y}_i|x_i)\Phi_i(x_i))] \\ & + \sum_i \xi_i \sum_{x_i=0}^{M-1} (b_i(x_i) - 1). \end{aligned}$$

Differentiating with respect to $b_k(x_k)$ yields

$$\begin{aligned} \frac{\partial J(\vec{b})}{\partial b_k(x_k)} = & -2 \sum_{j \in \mathcal{V}(k)} \sum_{x_j=0}^{M-1} b_k(x_k) b_j(x_j) \log(\Psi_{kj}(x_k, x_j)) + 1 \\ & + \log(b_k(x_k)) - \log(P_{\mathbf{y}_k|X_k}(\mathbf{y}_k|x_k)\Phi_k(x_k)) + \xi_k. \end{aligned} \quad (\text{B.5})$$

By setting the derivatives to zero and manipulating the expression above, we arrive at the fixed-point iteration for the belief functions

$$b_k(x_k) \leftarrow e^{1+\xi_k} P_{\mathbf{y}_k|X_k}(\mathbf{y}_k|x_k)\Phi_k(x_k) e^{2 \sum_{j \in \mathcal{V}(k)} \sum_{x_j=0}^{M-1} b_j(x_j) \log(\Psi_{kj}(x_k, x_j))} \quad (\text{B.6})$$

that immediately leads to the Mean Field update rule of Eq. (5) if we set $\gamma = e^{1+\xi_k}$.

Appendix C. Equivalence to the joint MAP solution

In this appendix, we present the connection between the MAP solution and the classical GLM inference in the binary activation scenario. In particular, we show that for a prior that models all voxels independently, there is a setting of the prior that forces the MAP solution to coincide with the classical solution. Furthermore, the two-step estimation procedure we outlined in Section 3, which first finds the ML estimates of the GLM parameters $\{\boldsymbol{\alpha}, \boldsymbol{\beta}, \sigma^2\}$ and then uses them in the MAP estimation of the state variables X , is asymptotically equivalent to the optimal simultaneous MAP estimation of the activation state and the regression parameters in the case of the voxel-wise independent prior.

To simplify the derivations, we let $\Theta_{0i} = \{\boldsymbol{\alpha}_i, \sigma_i^2\}$ be the parameters of the likelihood under H_0 for voxel i , $\Theta_{1i} = \{\boldsymbol{\alpha}_i, \beta_i, \sigma_i^2\}$ be the parameters of the likelihood under H_1 for voxel i , $\hat{\Theta}_{0i}^{\text{MAP}} (\hat{\Theta}_{0i}^{\text{ML}})$ and $\hat{\Theta}_{1i}^{\text{MAP}} (\hat{\Theta}_{1i}^{\text{ML}})$ be the MAP (ML) estimates of Θ_{0i} and Θ_{1i} , respectively. The activation state estimate obtained through the joint MAP estimation of the activation state and the GLM parameters

$$\begin{aligned} \{\vec{x}^* \vec{\theta}_0^* \vec{\theta}_1^*\} = & \arg \max_{\vec{x}, \vec{\theta}_0, \vec{\theta}_1} P_{\vec{x}, \vec{\theta}_0, \vec{\theta}_1}(\vec{x}, \vec{\theta}_0, \vec{\theta}_1 | \vec{y}) \\ = & \arg \max_{\vec{x}, \vec{\theta}_0, \vec{\theta}_1} P_{\vec{x}, \vec{\theta}_0, \vec{\theta}_1, \vec{y}}(\vec{x}, \vec{\theta}_0, \vec{\theta}_1, \vec{y}) \end{aligned} \quad (\text{C.1})$$

can also be expressed as

$$\vec{x}^* = \arg \max_{\vec{x}} \max_{\vec{\theta}_0, \vec{\theta}_1} P_{\vec{x}, \vec{\theta}_0, \vec{\theta}_1, \vec{y}}(\vec{x}, \vec{\theta}_0, \vec{\theta}_1, \vec{y}) \quad (\text{C.2})$$

$$= \arg \max_{\vec{x}} P_{\vec{x}}(\vec{x}) \prod_{i=1}^N \max_{\theta_0, \theta_1} P_{\Theta_{0i}, \Theta_{1i}, \mathbf{y}_i|X_i}(\theta_0, \theta_1, \mathbf{y}_i | x_i) \quad (\text{C.3})$$

For a voxel-wise independent prior $P_{\vec{x}}(\vec{x}) = \prod_{i=1}^N P_{X_i}(x_i)$, the problem above can be solved for each voxel separately. In the remainder of this section, we omit the voxel index i .

$$x^* = \arg \max_{x \in \{0,1\}} P_X(x) \max_{\theta_0, \theta_1} P_{\Theta_{0i}, \Theta_{1i}, \mathbf{y}|X}(\theta_0, \theta_1, \mathbf{y} | x) \quad (\text{C.4})$$

$$= \arg \max_{x \in \{0,1\}} P_X(x) \max_{\theta} P_{\mathbf{y}, \Theta_{x|X}}(\mathbf{y}, \theta | x) \quad (\text{C.5})$$

$$= \arg \max_{x \in \{0,1\}} P_X(x) P_{\mathbf{y}, \Theta_{x|X}}(\mathbf{y}, \hat{\Theta}_x^{\text{MAP}} | x) \quad (\text{C.6})$$

$$= \arg \max_{x \in \{0,1\}} P_X(x) P_{\mathbf{y}|\Theta_{x|X}}(\mathbf{y} | \hat{\Theta}_x^{\text{MAP}}, x) \quad (\text{C.7})$$

$$\approx \arg \max_{x \in \{0,1\}} P_X(x) P_{\mathbf{y}|\Theta_{x|X}}(\mathbf{y} | \hat{\Theta}_x^{\text{ML}}, x) \quad (\text{C.8})$$

We obtained Eq. (C.6) by assuming a uniform prior distribution for Θ_x . Eq. (C.8) is true asymptotically as the number of time points in a time course increases. Alternatively, we can express Eq. (C.8) in terms of hypothesis testing:

$$\frac{P_{Y|X}(\mathbf{y}|\hat{\Theta}_1^{\text{ML}}, x=1)}{P_{Y|X}(\mathbf{y}|\hat{\Theta}_0^{\text{ML}}, x=0)} \underset{H_0}{\overset{H_1}{\geq}} \frac{P_X(x=0)}{P_X(x=1)}. \quad (\text{C.9})$$

As shown in (Rencher, 2002), the left-hand side of Eq. (C.9) is a monotonic function of the F -statistic. Thus, we obtain the corresponding hypothesis testing threshold, η , according to the selected p value threshold in the classical GLM procedure:

$$\frac{P_{Y|X}(\mathbf{y}|\hat{\Theta}_1^{\text{ML}}, x=1)}{P_{Y|X}(\mathbf{y}|\hat{\Theta}_0^{\text{ML}}, x=0)} \underset{H_0}{\geq} \eta. \quad (\text{C.10})$$

This implies that selecting $P_X(x=0) = \frac{\eta}{1+\eta}$ and $P_X(x=1) = \frac{1}{1+\eta}$ causes the MAP with independent prior and the classical inferences to produce identical results.

References

- Allison, J.D., Meador, K.J., Loring, D.W., Figueroa, R.E., Wright, J.C., 2000. Functional MRI cerebral activation and deactivation during finger movement. *Neurology* 54, 135–142.
- Amedi, A., Malach, R., Hendler, T., Peled, S., Zohary, E., 2001. Visuo-haptic object-related activation in the ventral visual pathway. *Nature Neuroscience* 4, 324–330.
- Andrade, A., Kherif, F., Mangin, J.-F., Worsley, K.J., Paradis, A.-L., Simon, O., Dehaene, S., Bihan, D.L., Poline, J.-B., 2001. Detection of fMRI activation using cortical surface mapping. *Human Brain Mapping* 12, 79–93.
- Aron, A.R., Shohamy, D., Clark, J., Myers, C., Gluck, M.A., Poldrack, R.A., 2004. Human midbrain sensitivity to cognitive feedback and uncertainty during classification learning. *Journal of Neurophysiology* 92, 1144–1152.
- Ashtari, M., Perrine, K., Elbaz, R., Syed, U., Thaden, E., McIlree, C., Dolgoff-Kaspar, R., Clarke, T., Diamond, A., Ettinger, A., 2005. Mapping the functional anatomy of sentence comprehension and application to presurgical evaluation of patients with brain tumor. *American Journal of Neuroradiology* 26, 1461–1468.
- Benjamini, J., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B* 57, 289–300.
- Besag, J., 1974. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society Series B* 36, 192–236.
- Boykov, Y., Veksler, O., Zabih, R., 2001. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 1222–1239.
- Buckner, R.L., Snyder, A.Z., Shannon, B.J., LaRossa, G., Sachs, R., Fotenos, A.F., Sheline, Y.I., Klunk, W.E., Mathis, C.A., Morris, J.C., Mintun, M.A., 2005. Molecular, structural, and functional characterization of Alzheimer's disease: evidence for a relationship between default activity, amyloid, and memory. *Journal of Neuroscience* 25, 7709–7717.
- Burock, M.A., Dale, A.M., 2000. Estimation and detection of event-related fMRI signals with temporally correlated noise: a statistically efficient and unbiased approach. *Human Brain Mapping* 11, 249–260.
- Cormen, T.H., Leiserson, C.E., Rivest, R.L., 1997. *Introduction to Algorithms*. The MIT Press and McGraw-Hill.
- Cosman, E.R., Fisher, J.W., Wells, W.M., 2004. Exact MAP activity detection in fMRI using a GLM with an Ising spatial prior. In: *Proceedings of the Medical Image Computing and Computer-Assisted Intervention, LNCS*, vol. 3217, pp. 703–710.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis. I: segmentation and surface reconstruction. *NeuroImage* 9, 179–194.
- Descombes, X., Kruggel, F., von Cramon, D.Y., 1998. Spatio-temporal fMRI analysis using Markov random fields. *IEEE Transactions on Medical Imaging* 17, 1028–1039.
- Flandin, G., Penny, W.D., 2007. Bayesian fMRI data analysis with sparse spatial basis function priors. *NeuroImage* 34, 1108–1125.
- Freeman, W., Pasztor, E.C., Carmichael, O.T., 2000. Learning low-level vision. *International Journal of Computer Vision* 40, 25–47.
- Friston, K.J., Worsley, K.J., Frackowiak, R.S.J., Massiotto, J.C., Evans, A.C., 1993. Assessing the significance of local activations using their spatial extent. *Human Brain Mapping* 1, 210–220.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.-B., Frith, C.D., Frackowiak, R.S.J., 1995. Statistical parametric maps in functional imaging: a general linear approach. *Human Brain Mapping* 2, 189–210.
- Genovese, C.R., Noll, D.C., Eddy, W.F., 1997. Estimating test–retest reliability in functional MR imaging. I: statistical methodology. *Magnetic Resonance in Medicine* 38, 497–507.
- Goebel, R., Esposito, F., Formisano, E., 2006. Analysis of functional image analysis contest (FIAC) data with brainvoyager QX: from single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. *Human Brain Mapping* 27, 392–401.
- Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V., 2004. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. *National Academy of Sciences* 101, 4637–4642.
- Greig, D.M., Porteous, B.T., Gramon, D.Y., 1989. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society Series B* 51, 271–279.
- Hartvig, N.V., 2002. A stochastic geometry model for functional magnetic resonance images. *Scandinavian Journal of Statistics* 29, 333–353.
- Jaakkola, T., 2000. *Tutorial on variational approximation methods*. *Advanced Mean Field Methods: Theory and Practice*. MIT Press.
- Jezzard, P., Balaban, R.S., 1995. Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine* 34, 65–73.
- Jezzard, F., Matthews, P.M., Smith, S.M., 2002. *Functional MRI – An Introduction to Methods*. Oxford University Press, Oxford.
- Kapur, T., Grimson, W.E.L., Wells, W.M., Kikinis, R., 1998. Enhanced spatial priors for segmentation of magnetic resonance imagery. In: *Proceedings of the Medical Image Computing and Computer-Assisted Intervention, LNCS*, vol. 1496, pp. 457–468.
- Kiebel, S.J., Goebel, R., Friston, K.J., 2000. Anatomically informed basis functions. *NeuroImage* 11, 656–667.
- Langan, D.A., Molnar, K.J., Modestino, J.W., Zhang, J., 1992. Use of the mean-field approximation in an EM-based approach to unsupervised stochastic model-based image segmentation. In: *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 3, pp. 57–60.
- Li, T.Q., Kastrup, A., Takahashi, A.M., Moseley, M.E., 1999. Function MRI of human brain during breath holding by BOLD and FAIR techniques. *NeuroImage* 9, 243–249.
- Lin, F.S., Huang, T.Y., Chen, N.K., Wang, F.N., Stufflebeam, S.M., Belliveau, J.W., Wald, L.L., Kwong, K.K., 2005. Functional MRI using regularized parallel imaging acquisition. *Magnetic Resonance in Medicine* 54, 343–353.
- Liou, M., Su, H.R., Lee, J.D., Aston, J.A.D., Tsai, A.C., Cheng, P.E., 2006. A method for generating reproducible evidence in fMRI studies. *NeuroImage* 29, 383–395.
- Magnotta, V.A., Friedman, L., 2006. FIRST BIRN measurement of signal-to-noise and contrast-to-noise in the fBIRN multi-center imaging study. *Journal of Digital Imaging* 19, 140–147.
- Maitra, R., Roys, S., Gullapalli, R., 2002. Test–retest reliability estimation of functional MRI data. *Magnetic Resonance in Medicine* 48, 62–70.
- Makni, S., Idier, J., Vincent, T., Thirion, B., Dehaene-Lambertz, G., Ciuciu, P., 2008. A fully Bayesian approach to the parcel-based detection-estimation of brain activity in fMRI. *NeuroImage* 41, 941–969.
- Ou, W., Golland, P., 2005. From spatial regularization to anatomical priors in fMRI analysis. In: *Proceedings of the IPMI, LNCS*, vol. 3565, pp. 88–100.
- Ou, W., 2005. *fMRI detection with spatial regularization*. MIT Master Thesis.
- Paris, G., 1998. *Statistical Field Theory*. Addison-Wesley.
- Penny, W.D., Trujillo-Barreto, N.J., Friston, K.J., 2005. Bayesian fMRI time series analysis with spatial priors. *NeuroImage* 24, 350–362.
- Penny, W.D., Flandin, G., Trujillo-Barreto, N.J., 2007. Bayesian comparison of spatially regularised general linear models. *Human Brain Mapping* 28, 275–293.
- Pohl, K.M., Wells, W.M., Guimond, A., Kasai, K., Shenton, M.E., Kikinis, R., Grimson, W.E.L., Warfield, S.K., 2002. Incorporating non-rigid registration into expectation maximization algorithm to segment MR images. In: *Proceedings of the Medical Image Computing and Computer-Assisted Intervention, LNCS*, vol. 2488, pp. 564–571.
- Pohl, K.M., Bouix, S., Kikinis, R., Grimson, W.E.L., 2004. Anatomical guided segmentation with non-stationary tissue class distributions in an expectation-maximization framework. In: *Proceedings of the IEEE Symposium on Biomedical Imaging*, vol. 1, pp. 81–84.
- Rajapakse, J.C., Piyaratna, J., 2001. Bayesian approach to segmentation of statistical parametric maps. *IEEE Transactions on Biomedical Engineering* 48, 1186–1194.
- Rencher, A.C., 2002. *Methods of Multivariate Analysis*. Wiley.
- Salli, E., Aronen, H.J., Savolainen, S., Korvenoja, A., Visa, A., 2001. Contextual clustering for analysis of functional MRI data. *IEEE Transactions on Medical Imaging* 20, 403–414.
- Seghier, M.L., Lazeyras, F., Zimine, S., Maier, S.E., Hanquinet, S., Delavelle, J., Volpe, J.J., Huppi, P.S., 2004. Combination of event-related fMRI and diffusion tensor imaging in an infant with perinatal stroke. *NeuroImage* 21, 463–472.
- Strother, S.C., Anderson, J., Hansen, L.K., Kjems, U., Kustra, R., Sidtis, J., Frutiger, S., Muley, S., LaConte, S., Rottenberg, D., 2002. The quantitative evaluation of functional neuroimaging experiments: the NPAIRS data analysis framework. *NeuroImage* 15, 747–771.
- Svensen, M., Kruggel, F., von Cramon, D.Y., 2000. Probabilistic modeling of single-trial fMRI data. *IEEE Transactions on Medical Imaging* 19, 25–35.
- Swensson, R.G., 1996. Unified measurement of observer performance in detecting and localizing target objects on images. *Medical Physics* 23, 1025–1029.
- Van De Ville, D., Seghier, M.L., Lazeyras, F., Blu, T., Unser, M., 2007. WSPM: wavelet-based statistical parametric mapping. *NeuroImage* 37, 1205–1217.

- Vincent, T., Ciuciu, P., Idier, J., 2007. Spatial mixture modelling for the joint detection-estimation of brain activity in fMRI. In: Proceedings of the ICASSP, vol. 1, pp. 325–328.
- Woolrich, M.W., Ripley, B.D., Brady, M., Smith, S.M., 2001. Temporal autocorrelation in univariate linear modelling of fMRI data. *NeuroImage* 14, 1370–1386.
- Woolrich, M.W., Behrens, T.E., Beckmann, C.F., Smith, S.M., 2005. Mixture models with adaptive spatial regularization for segmentation with an application to fMRI data. *IEEE Transactions on Medical Imaging* 24, 1–11.
- Worsley, K.J., Liao, C.H., Aston, J., Petre, V., Duncan, G.H., Morales, F., Evans, A.C., 2002. A general statistical analysis for fMRI data. *NeuroImage* 15, 1–15.
- Wu, F., 1982. The Potts model. *Reviews of Modern Physics* 54, 235–268.
- Zou, K.H., Greve, D.N., Wang, M., Pieper, S.D., Warfield, S.K., White, N.S., Manandhar, S., Brown, G.G., Vangel, M.G., Kikinis, R., Wells, W.M., 2005. Reproducibility of functional MR imaging: preliminary results of prospective multi-institutional study performed by biomedical informatics research network. *Radiology* 237, 781–789.