NeuroImage 63 (2012) 1519–1531

Contents lists available at SciVerse ScienceDirect

NeuroImage

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

Bayesian hierarchical multi-subject multiscale analysis of functional MRI data

Nilotpal Sanyal, Marco A.R. Ferreira *

Department of Statistics, University of Missouri, 146 Middlebush Hall, Columbia, MO 65211-6100, United Sates

ARTICLE INFO

Article history:
Accepted 15 August 2012
Available online 21 August 2012

Keywords:
Bayesian inference
Image smoothing
Mixture prior
Multiple subjects
Spatiotemporal analysis
Wavelet modeling

ABSTRACT

We develop a methodology for Bayesian hierarchical multi-subject multiscale analysis of functional Magnetic Resonance Imaging (fMRI) data. We begin by modeling the brain images temporally with a standard general linear model. After that, we transform the resulting estimated standardized regression coefficient maps through a discrete wavelet transformation to obtain a sparse representation in the wavelet space. Subsequently, we assign to the wavelet coefficients a prior that is a mixture of a point mass at zero and a Gaussian white noise. In this mixture prior for the wavelet coefficients, the mixture probabilities are related to the pattern of brain activity across different resolutions. To incorporate this information, we assume that the mixture probabilities for wavelet coefficients at the same location and level are common across subjects. Furthermore, we assign for the mixture probabilities a prior that depends on a few hyperparameters. We develop an empirical Bayes methodology to estimate the hyperparameters and, as these hyperparameters are shared by all subjects, we obtain precise estimated values. Then we carry out inference in the wavelet space and obtain smoothed images of the regression coefficients by applying the inverse wavelet transform to the posterior means of the wavelet coefficients. An application to computer simulated synthetic data has shown that, when compared to single-subject analysis, our multi-subject methodology performs better in terms of mean squared error. Finally, we illustrate the utility and flexibility of our multi-subject methodology with an application to an event-related fMRI dataset generated by Postle (2005) through a multi-subject fMRI study of working memory related brain activation.

1. Introduction

Functional Magnetic Resonance Imaging (fMRI) experiments have become a ubiquitous tool for the study of human brain activation patterns. These activation patterns usually exhibit spatial dependence that can be well modeled by including the spatial features as part of a probabilistic model within a Bayesian framework (Badillo et al., 2011; Bowman et al., 2008; Flandin and Penny, 2007; Gössl et al., 2001; Groves et al., 2009; Harrison and Green, 2010; Harrison et al., 2008; Ng et al., 2010a,b; Penny et al., 2003, 2005; Quirós et al., 2010a,b; Smith et al., 2003; Vincent et al., 2010). In that regard, a particularly effective Bayesian framework assigns spatial priors to the regression coefficients of a general linear model, such as for example Gaussian Markov random field priors (Penny et al., 2005), wavelet basis priors (Flandin and Penny, 2007), anatomical parcellation-based hierarchical Gaussian priors (Bowman et al., 2008), and diffusion-based spatial priors (Harrison et al., 2008). In particular, the wavelet basis priors of Flandin and Penny (2007) are able to handle spatial variations in smoothness and, as a consequence, better recover the original shapes of the activation regions. However, the methodology of Flandin and Penny (2007) is restricted to the analysis of *single-subject* fMRI data. Here, we build upon their approach and develop a novel hierarchical multiscale Bayesian analysis for data from *multiple subjects*. Our methodology borrows information across subjects and thus leads to more precisely estimated activation maps.

Following the hierarchical approach of Flandin and Penny (2007), we begin by modeling the brain images temporally with a standard general linear model (GLM). The GLM accounts for the temporal information contained in the data. The estimated regression coefficient images represent the spatial variation contained in the data. For the sake of comparability across voxels, we standardize the estimated regression coefficients using their standard deviation. Next, similar to Flandin and Penny (2007), we transform the estimated standardized regression coefficient images through a discrete wavelet transformation to obtain a sparse representation in the wavelet space. The use of the wavelet transform captures the spatial variation across different resolution levels. Subsequently, we assign to the wavelet coefficients a prior that is a mixture of a point mass at zero and a Gaussian component (Abramovich et al., 1998; Clyde et al., 1998). An important element of this prior is the probability of the Gaussian component, that we refer to as the mixture probability. Specifically, this is the prior probability that a wavelet coefficient is different than zero. Previous works in the literature (Abramovich et al., 1998; Clyde et al., 1998) usually assume that the mixture probabilities vary across scales but are constant across locations within each scale. Here we take a novel approach that assumes that the mixture probabilities vary across locations within each scale.

We model the mixture probabilities using a hierarchical specification that borrows information about the activation pattern within

* Corresponding author. Fax: +1 573 884 5524.
E-mail addresses: ns5x2@mail.missouri.edu (N. Sanyal), ferreiram@missouri.edu (M.A.R. Ferreira).

1053-8119/$ – see front matter © 2012 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.neuroimage.2012.08.041

different resolution levels from multiple subjects. Specifically, as the data are co-registered and normalized, the wavelet coefficients from different subjects at the same position within a given resolution level correspond to the same anatomical position. To incorporate this information into the model, we assume that the mixture probabilities for wavelet coefficients at the same location and level are common across subjects. Furthermore, for the mixture probabilities, we consider a Beta prior with shape parameters depending on resolution level. We parameterize these shape parameters in terms of a few hyperparameters that encode how sparse is the wavelet decomposition across resolution levels. Finally, the application of our methodology to a real dataset (Section 3.2) shows that our methodology is able to identify the common activation pattern at coarser resolutions across subjects and concurrently respects the natural variability among subjects at finer resolution levels.

For the analysis of our hierarchical model, we develop an inference procedure that is performed in three steps. In the first step, we develop an empirical Bayes methodology to obtain estimates of the model hyperparameters as well as their standard errors. As shown in the applications in Section 3, these standard errors are orders of magnitude smaller than the hyperparameter estimates. Because of their high estimation precision, in the subsequent steps we take an empirical Bayes approach, that is, we hold the hyperparameters fixed at their estimates. Because the hyperparameters encode the sparsity of the wavelet decomposition, in our methodology the degree of sparsity is estimated from the data. In the second step, conditional on the estimated hyperparameters, our approach leads to explicit posterior distributions for the wavelet coefficients. Using these posterior distributions, we carry out inference in the wavelet domain and obtain the posterior mean of the wavelet coefficients. Finally, in the third step we obtain smoothed images of the regression coefficients by applying the inverse wavelet transform to the posterior means of the wavelet coefficients. These regression coefficient maps can then be used to obtain group level posterior images. Finally, for uncertainty assessment we develop a simulation-based method for the computation of the posterior variance of the regression coefficients.

Our inferential approach is quite different and much faster than the variational Bayes approach previously used in the literature (Flandin and Penny, 2007; Penny et al., 2003, 2005). The variational Bayes method estimates all the parameters jointly by approximating high dimensional integrals. Moreover, the variational Bayes method can be seen as an extension of the EM algorithm and, as such, may be slow to converge, may converge to a local optimum, and may underestimate the inferential uncertainty. In contrast, as we show in Section 2.2, by taking an empirical Bayes approach we are able to solve exactly most of the high-dimensional integrals analytically. Further, in our approach the few integrals that cannot be solved analytically are all one-dimensional and can be solved numerically with arbitrary precision. Therefore, our methodology is fast and quite accurate with respect to the characterization of the model uncertainty.

A notable merit of our methodology is the development of an inference procedure that does not use Markov chain Monte Carlo (MCMC) simulation. Many of the existing Bayesian methods of fMRI data analysis perform inference using MCMC simulation (Badillo et al., 2011; Bowman et al., 2008; Costafreda et al., 2009; Genovese, 2000; Gössl et al., 2001; Makni et al., 2008; Rajapakse and Zhou, 2007; Smith and Fahrmeir, 2007; Vincent et al., 2010; Woolrich et al., 2004, 2005). The use of MCMC simulation makes multi-subject analysis computationally expensive. Conversely, as we show in Section 2.2, in our work most of the integrals can be solved analytically. Moreover, in our work the integrals that have to be approximated are all one-dimensional and, consequently, can be solved numerically with arbitrary precision using fast numerical integration methods such as Newton–Cotes (Press et al., 1992). Considering the usually large size of fMRI datasets, this feature of our methodology is greatly useful and considerably reduces the cost of computation.

The use of the wavelet transform for the analysis of fMRI data has been proposed in the literature in a number of works. Most of this previous literature applies the wavelet transform directly to the BOLD response in the temporal domain (Alexander et al., 2000; Fadili and Bullmore, 2002; Luo and Puthusserypady, 2008; Meyer, 2003; Ruttimann et al., 1998; Sendur et al., 2005; Van De Ville et al., 2004; Wink and Roerdink, 2004; Zaroubi and Goelman, 2000). Another possible approach, proposed in the context of positron emission tomography by Turkheimer et al. (2003, 2006) and that may be adapted for fMRI studies, is to apply a 2D or 3D wavelet transform to the spatially distributed BOLD response at each time point. Then, the resulting wavelet coefficients are related to the underlying physiological process through a general linear model. Finally, a quite promising approach applies the wavelet transform to the regression coefficients in the spatial domain (Flandin and Penny, 2007; Long et al., 2004; Soleymani et al., 2009). In particular, Flandin and Penny (2007) use wavelets to define a prior for the regression coefficients obtained by fitting a GLM to the BOLD response. However, while the methodology of Flandin and Penny (2007) is for single-subject analysis, our methodology is for multi-subject analysis. See Bullmore et al. (2004), Van De Ville et al. (2006a) and Wongsawat (2009) for comprehensive reviews of the use of wavelet analysis for fMRI data.

The remainder of this article is organized as follows. Section 2 describes the model at different stages of the hierarchy, including the choice of the prior and hyperparameters, the inference procedure including estimation of the hyperparameters and the exact posterior distribution for the empirical wavelet coefficients, and the posterior reconstruction of the regression coefficient images. Section 3 presents applications of our methodology to two datasets: a simulated dataset, and a real dataset generated by Postle (2005) from an event-related fMRI study of working memory related brain activation. Finally, Section 4 contains an overall discussion of the merits of our methodology and concluding remarks.

2. Theory

In this section, we describe the different components of our Bayesian hierarchical multi-subject multiscale model along with the assumptions made, the inference method in the wavelet domain, and finally the reconstruction of the images back to the regression coefficients space. The probabilistic model along with the different parameters and their dependencies is represented in a diagram in Fig. 1.

2.1. Model

2.1.1. Temporal model for the BOLD response

We start by modeling the BOLD response for each subject with a standard general linear model (Friston et al., 1994). Specifically, the whole volume of a subject's brain is divided into 3D volumetric pixels or voxels. During a task that the subject performs, the whole brain is scanned at multiple time points. Thus, corresponding to each voxel of the brain, the experiment yields a time-series of BOLD response. For the ith subject, $i=1,\ldots,I$, we model the BOLD response as

$$\mathbf{y}_i = \mathbf{X}_i\beta_i + \epsilon_i, \tag{1}$$

where $\mathbf{y}_i$ is the $T\times N$ matrix of BOLD response, $\mathbf{X}_i$ is a $T\times K$ matrix of regressors, β_i is a $K\times N$ matrix of regression coefficients, ϵ_i is a $T\times N$ matrix of errors, T is the number of time points, N is the number of voxels and K is the number of regressors.

The nth column of the data matrix $\mathbf{y}_i$ contains the BOLD response time series for the nth voxel. In the case where there are multiple runs corresponding to the ith subject, the time series for all the runs for a voxel are stacked into a single column of the data matrix $\mathbf{y}_i$. For each subject, the design matrix $\mathbf{X}_i$ is the same for all the voxels, and each column of $\mathbf{X}_i$ corresponds to the values of one regressor. The regressors considered in the model are convolutions of indicator variables for each experimental condition with an empirically derived subject-specific hemodynamic response function. With this specification of the design matrix $\mathbf{X}_i$, Eq. (1)