Anatomically Informed Bayesian Model Selection for fMRI Group Data Analysis

Merlin Keller^{1,2,3}, Marc Lavielle^{2,5}, Matthieu Perrot^{1,4}, and Alexis Roche¹

¹ LNAO, Neurospin, CEA, F-91191 Gif-sur-Yvette, France
`merlin.keller@cea.fr`

² Department of Probability and Statistics, University of Paris Sud, France

³ PARIETAL team, INRIA Saclay, France

⁴ INSERM U.797, Orsay, France

⁵ University René Descartes, Paris, France

Abstract. A new approach for fMRI group data analysis is introduced to overcome the limitations of standard voxel-based testing methods, such as Statistical Parametric Mapping (SPM). Using a Bayesian model selection framework, the functional network associated with a certain cognitive task is selected according to the posterior probabilities of mean region activations, given a pre-defined anatomical parcellation of the brain. This approach enables us to control a Bayesian risk that balances false positives and false negatives, unlike the SPM-like approach, which only controls false positives. On data from a mental calculation experiment, it detected the functional network known to be involved in number processing, whereas the SPM-like approach either swelled or missed the different activation regions.

1 Introduction

One of the goals of fMRI group studies is to identify the brain structures that are consistently involved in a given cognitive task across individuals. Mass univariate, or voxel-based, detection [1] is to date the most widely used approach to address such questions. It starts with normalizing individual images onto a common brain template using nonrigid image registration. Next, a t -statistic is computed in each voxel to locally assess mean group effects. The candidate regions are then defined as the connected components (clusters) of the resulting statistical map above an arbitrary threshold called the cluster-forming threshold.

To account for the multiple testing problem [2], only the clusters whose sizes exceed a critical value are reported. This critical size acts as a second threshold, and is generally tuned to control the probability of one or more clusters being detected by chance. Such statistical calibration is achieved in practice using either analytical approximations or resampling techniques [3]. For scientific reporting purposes, the detected clusters may finally be related to known anatomical regions based on expert knowledge, or using a digital brain atlas such as the Automated Atlas Label (AAL) [4], or the Cortical Sulci Atlas (CSA) [5].

Fig. 1. Clusters detected at different cluster-forming thresholds by the SPM-like approach (axial slices $z = 37\text{mm}$ in Talairach, with the subjects' mean anatomical image in the background). From left to right, the threshold is tuned to control the false positive rate (FPR) respectively at 10^{-2} , 10^{-3} and 10^{-4} uncorrected. Each cluster with corrected p -values less than 5% is represented with a specific color, showing how distinct functional regions are merged. Far right: the CSA atlas.

This approach, which we will refer to in the following as Statistical Parametric Mapping-like (SPM-like), is simple and widely applicable. However it suffers from the following drawbacks:

Arbitrary Cluster-Forming Threshold. The fact that a suprathreshold cluster is found significant by the cluster size test only implies that it contains *some* active voxels [3]. Low values of the cluster-forming threshold may result in merging functionally distinct regions, thus yielding poor localization power, while high values may result in missing active regions. This is illustrated in Figure 1 (the dataset and procedure are detailed in Section 3.2).

Exclusive Control of False Positives. False negative regions are not controlled, meaning that the absence of activations outside the detected clusters cannot be assessed. Consequently, there is no guarantee that the whole functional network can be recovered. This partly explains the poor reproducibility of group analyses across datasets [6].

Assumption of Perfect Match Between Individual Brains. Due to unavoidable intersubject registration errors, the observed activations are not well-localized, and possibly displaced across distinct functional regions, which may result in blurring the group activation map and creating unhandled false positives [7].

To date, these issues have been tackled separately. In [8], the SPM-like approach is extended to control false negatives using Bayesian inference, while [9] addresses the arbitrary threshold problem. Uncertainty on the localization of individual effects is accounted for in [10] by comparing high-level features extracted from the individual images and in [7] by extending the mass univariate model.

This paper introduces a new approach for fMRI group data analysis that jointly addresses all the above issues. Our method rests upon a Bayesian model selection framework in which models are functional networks, defined as binary

partitions of a given brain parcellation. Each network is then rated in terms of a posterior probability. Because regions are defined beforehand, the method is threshold-free, while taking advantage of the prior knowledge about functionally homogeneous regions provided that the parcellation is sensible. A similar use of pre-defined regions with Bayesian modeling can be found in [11,12], while other authors have used spatial priors rather than explicit parcellations [13].

By controlling a Bayesian risk, our approach balances false positives and false negatives, with relative weights that may be tuned depending on the application. Furthermore, the generative model, described in Section 2.1, accounts for spatial uncertainty on the individual effects due to spatial normalization errors. Results of this procedure are presented on both simulated and real fMRI data in Section 3, and we conclude with a brief discussion in Section 4.

2 Method

2.1 Generative Model

After scanning several subjects during an fMRI experiment, individual datasets are processed separately so that, for each subject $i = 1, \dots, n$, we are given an image of estimated effects $Y_i \in \mathbb{R}^p$, and an image of estimation variances $s_i^2 \in \mathbb{R}^p$, where p is the number of voxels in the search volume $V = \{v_k, k = 1, \dots, p\} \subset \mathbb{R}^3$. Following [7], we extend the mass univariate model in [14] by relaxing the assumption that the individual images are perfectly aligned so that, at voxel k ,

$$Y_i(v_k) = \mu(v_k + u_{ik}) + \varepsilon_{ik}. \quad (1)$$

Here, we note $Y_i(v_k) = Y_{ik}$ to emphasize that it is a spatial map; $\mu \in \mathbb{R}^p$ is the map of mean population effects; the vector u_{ik} is a hidden variable that models the subject-to-atlas registration error for subject i at voxel k ; finally, the ε_{ik} are independently distributed Gaussian variables $\mathcal{N}(0, \sigma_I^2 + s_{ik}^2)$, where σ_I^2 is the between-subject variance and s_{ik}^2 the within-subject variance due to estimation errors in the effects. This generalizes the model in [14], which corresponds to the special case where the estimation variances s_{ik}^2 and the registration errors u_{ik} are neglected.

The displacements u_{ik} are modeled by a Karhunen-Loève expansion using elementary displacements w_{ib} in a limited number of fixed control points $\{v_{kb}, b = 1, \dots, B\}$. As is a common approach in deformable models, those control points are interpolated to the whole brain image using a radial basis function K ,

$$u_{ik} = \sum_{b=1}^B K(v_k, v_{kb}) w_{ib}, \quad (2)$$

where $K(v_k, v_{kb}) = \exp\{-\|v_k - v_{kb}\|^2/\gamma^2\}$, and γ controls the displacement field's smoothness. The w_{ib} 's are independent trivariate Gaussian variables with zero mean and common spherical covariance matrix $\sigma_S^2 \mathbf{I}_3$, where σ_S models the standard registration error.

2.2 Bayesian Model Selection Framework

Our approach to functional network selection is based on an *a priori* partition of the search volume into N regions of interest $V = V_1 \cup \dots \cup V_N$, assumed to be homogeneous functional areas. Thus, throughout region V_j the population mean effects $\{\mu_k, v_k \in V_j\}$ are modeled as independent random variables with common distribution $\mathcal{N}(\eta_j, \sigma_j^2)$, where η_j is the regional mean effect and σ_j^2 the regional variance.

Based on the above model (1) and (2), our final task is to select the regions involved in the considered task. To this end, we define region V_j as being *involved* in the task if its regional mean effect is nonzero, *i.e.* $\eta_j \neq 0$, and *inactive* if $\eta_j = 0$. The functional network we wish to recover is thus represented in the following by the unknown binary label vector $\Gamma^* \in \{0, 1\}^N$, such that $\Gamma_j^* = 1$ if $\eta_j \neq 0$ and $\Gamma_j^* = 0$ if $\eta_j = 0$. Adopting a Bayesian model selection viewpoint, we estimate Γ^* by selecting the most probable network given the data Y :

$$\hat{\Gamma}^* = \arg \max_{\Gamma} p(\Gamma|Y) \quad (3)$$

$$= \arg \max_{\Gamma} p\{\forall j, \Gamma_j = 1 : \eta_j \neq 0; \quad \forall j, \Gamma_j = 0 : \eta_j = 0|Y\}. \quad (4)$$

However, computing $p(\Gamma|Y)$ for all 2^N possible choices of Γ is intractable. To cut down computational complexity, we use the following independence approximation:

$$p(\Gamma|Y) \approx \prod_{\Gamma_j=1} p(\eta_j \neq 0|Y) \prod_{\Gamma_j=0} p(\eta_j = 0|Y). \quad (5)$$

It can be shown that (5) would be exact in absence of spatial uncertainty ($\sigma_S^2 = 0$, $u_{ik} \equiv 0$) and given the inter-subject variance σ_I^2 . Using (5), the most probable network $\hat{\Gamma}^*$ is the one for which all regions j with $p(\eta_j \neq 0|Y) > 0.5$. are active.

2.3 Posterior Probability Computation

For all $j = 1, \dots, N$, the posterior probabilities $P_j = p(\eta_j \neq 0|Y)$ and $p(\eta_j = 0|Y) = 1 - P_j$ are computed through their Bayes' factor:

$$\frac{P_j}{1 - P_j} = \frac{p(Y|\eta_j \neq 0)}{p(Y|\eta_j = 0)}, \quad (6)$$

where $p(Y|\eta_j \neq 0) = \int_{(\eta_j, \sigma_j^2) \in \mathbb{R} \times \mathbb{R}^+} p(Y|\eta_j, \sigma_j^2) \pi(\eta_j, \sigma_j^2) d(\eta_j, \sigma_j^2)$ is the marginal likelihood in the model where region j is involved in the considered task and $p(Y|\eta_j = 0) = \int_{\sigma_j^2 \in \mathbb{R}^+} p(Y|\eta_j = 0, \sigma_j^2) \pi(\sigma_j^2) d(\sigma_j^2)$ the marginal likelihood in the model where region j is inactive, assuming a uniform prior $\pi(\eta_j \neq 0) = \pi(\eta_j = 0) = 0.5$ on the state of region j . The priors $\pi(\eta_j, \sigma_j^2)$ and $\pi(\sigma_j^2)$ are normal-inverse-Gamma and inverse-Gamma distributions, as in [7].

Fig. 2. Results on simulated data. From left to right: original activation pattern (with CSA atlas in the background), t -score map, posterior mean estimate $\hat{\mu} = \mathbb{E}(\mu|Y, \hat{\theta})$ of the mean population effect map μ and histogram of posterior probabilities P_j found by our approach.

Computing the marginal likelihoods $p(Y|\eta_j \neq 0)$ is non trivial, since it involves evaluating integrals on high dimensional spaces. Following [15], we use the basic marginal equality :

$$p(Y|\eta_j \neq 0) = \frac{p(Y|\theta_j, \eta_j \neq 0)\pi(\theta_j)}{\pi(\theta_j|Y, \eta_j \neq 0)}, \quad (7)$$

valid for any value of the parameter $\theta_j = (\sigma_I^2, \sigma_S^2, \eta_j, \sigma_j^2)$, and apply it to the maximum *a posteriori* $\hat{\theta}_j$, which we estimate using a MCMC-SAEM algorithm [16]. The “null” likelihood $p(Y|\eta_j = 0)$ is computed in a similar fashion.

3 Results

3.1 Simulated Data

We designed an artificial activation map μ simulating the mean population effect defined in (1), based on the CSA atlas, which we used to analyze real fMRI data (see Section 3.2). To reflect realistic situations, we placed two activations in neighboring regions, one at the intersection of several regions, and a smaller one inside the largest atlas region (see Figure 2, far left).

$n = 40$ images were generated by warping this map according to the deformation model defined by (2), with one control point in each voxel, choosing $\gamma = 40\text{mm}$ and $\sigma_S = 8\text{mm}$. Homoscedastic noise ε_{ik} was then added to each image according to the model in Section 2.1, with $\sigma_I^2 = 1$, and the s_{ik}^2 ’s generated as independent chisquare variables.

We then applied our Bayesian model selection algorithm to this dataset, to recover the atlas regions $j = 1, \dots, N$ with a nonzero regional mean activation η_j , still using the deformation model defined in (2), except this time control points were restricted to a regular grid with regular spacing equal to γ along each axis. We also used the SPM-like approach, described in Section 1. The cluster-forming threshold was tuned to control the voxel-level false positive rate

(FPR) at 10^{-3} uncorrected. Clusters with corrected p -values less than 5% in the cluster-size test were reported, and labeled according to their maximum statistic.

Results from this simulation study are illustrated in Figure 2. The posterior mean estimate $\hat{\mu}$ of the mean effect map μ is more contrasted than the t -score map, indicating a better fit of our model, which accounts for localization uncertainty. The histogram of posterior probabilities illustrates the discriminative power of our Bayesian model selection algorithm in separating involved from inactive regions, with more than 90% of the posterior probabilities P_j either above 0.9 or under 0.1. Furthermore, regions j with regional means η_j above 0.04, were all correctly reported as active and the regional means η_j were estimated with an average relative error $\epsilon = \frac{1}{N} \sum_j |\hat{\eta}_j - \eta_j|/\eta_j$ of 7%. In contrast, while at the chosen threshold the SPM-like method detected all activations, with corrected p -values smaller than 10^{-3} in the cluster-size test, neighboring active regions were merged in both the left and right hemispheres.

3.2 Real fMRI Data

We now present results of our method on a real fMRI dataset. We considered the activation maps of 38 subjects for a ‘Calculation–Sentences’ contrast, which subtracts activations due to audio or video instructions from the overall activations during the computation tasks. This contrast may thus reveal regions specifically involved in number processing. As in the simulation study, the data was analyzed using both the SPM-like approach and our model selection algorithm, based on the CSA atlas [5], derived from the anatomical images of 63 subjects and comprising 125 regions which correspond to subdivisions of cortical sulci.

As in the simulation study, the histogram of posterior probabilities suggests a good discriminative power, with over 80% of the posterior probabilities either above 0.9 or below 0.1. As shown in Figure 3 and Table 1, our method detected the bilateral intra-parietal and fronto-cingular networks known to be active during number processing [17], with posterior probabilities higher than 0.94 for all frontal regions and higher than 0.83 for all parietal regions. Interestingly, the

Fig. 3. Real fMRI data, number processing task. From left to right: classical t -score map, posterior mean estimate $\hat{\mu} = \mathbb{E}(\mu|Y, \hat{\theta})$ of the mean population effect map μ , map of posterior probabilities P_j , histogram of posterior probabilities.

Table 1. Number processing task, regions detected using the Bayesian model selection approach. Reported regions have a posterior probability of being involved in the task greater than 0.5, and constitute the most probable functional network given the data. $\hat{\eta}_j$ is the posterior estimate of the regional mean effect. Asterisks (*) mark regions that are found significant at 5% by the SPM-like approach (corrected cluster-level inference, cluster-forming threshold set to $\text{FPR} = 10^{-3}$), and correspond to the middle activation map in Figure 1.

Frontal lobe		
Sulcus/Fissure	P_j	$\hat{\eta}_j$
Left middle frontal	1.00	2.94
Right middle frontal*	1.00	3.17
Left superior frontal	0.94	2.11
Right superior frontal	1.00	1.93
Left middle precentral	1.00	4.11
Right middle precentral	1.00	1.98
Left inferior precentral	1.00	5.30
Right inferior precentral	0.97	2.69
Left anterior cingular	1.00	3.39
Right anterior cingular	1.00	4.15
Left inferior frontal*	1.00	4.38

Parietal lobe		
Sulcus/Fissure	P_j	$\hat{\eta}_j$
Left intra-parietal*	1.00	4.87
Right intra-parietal	1.00	3.03
Left precuneus	0.98	5.58
Right precuneus	0.85	3.84
Right inferior postcentral	0.88	2.30
Right parieto-occipital	0.83	1.30

Other		
Sulcus/Fissure	P_j	$\hat{\eta}_j$
Right callosal	0.79	1.99

bilateral precuneus sulci where also detected. Although not considered as part of the core numerical system, the precuneus has been linked to memory access and a wide range of high-level tasks [18]. In contrast, only three activated clusters were detected by the SPM-like approach at the chosen cluster-forming threshold. Each cluster contained over a thousand voxels, and extended over several atlas regions, hence merging several functionally distinct areas. Also, no activations were detected in the right frontal area. Using different thresholds could not solve these problems, as illustrated in Figure 1.

4 Discussion

We have introduced a new approach for fMRI group data analysis which addresses several limitations of standard voxel-based methods, such as SPM. It is threshold-free, accounts for the imperfect match between individual images, and controls both false positive and false negative risks. In a calculation experiment, our method correctly detected the functional network associated with basic number processing, while the SPM-like approach either swelled or missed the different activation regions. Beside illustrating the benefits of combining functional and anatomical information in neuroimaging, these results plead for a new paradigm for fMRI group data analysis.

Acknowledgments. Support was provided by the EC-funded IMAGEN project.

References

1. Friston, K.J.: 2. In: *Human Brain Function*, pp. 25–42. Academic Press, London (1997)
2. Nichols, T., Hayasaka, S.: Controlling the Familywise Error Rate in Functional Neuroimaging: A Comparative Review. *Statistical Methods in Medical Research* 12(5), 419–446 (2003)
3. Hayasaka, S., Nichols, T.: Validating Cluster Size Inference: Random Field and Permutation Methods. *Neuroimage* 20(4), 2343–2356 (2003)
4. Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M.: Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage* 15(1), 273–289 (2002)
5. Perrot, M., Rivière, D., Mangin, J.F.: Identifying cortical sulci from localizations, shape and local organization. In: *5th Proc. ISBI*, Paris, France, pp. 420–423 (2008)
6. Thirion, B., Pinel, P., Mériaux, S., Roche, A., Dehaene, S., Poline, J.B.: Analysis of a large fMRI cohort: Statistical and methodological issues for group analyses. *Neuroimage* 35(1), 105–120 (2007)
7. Keller, M., Roche, A., Tucholka, A., Thirion, B.: Dealing with spatial normalization errors in fMRI group inference using hierarchical modeling. *Statistica Sinica* 18(4), 1357–1374 (2008)
8. Friston, K., Glaser, D.E., Henson, R.N.A., Kiebel, S., Phillips, C., Ashburner, J.: Classical and Bayesian inference in neuroimaging: Applications. *Neuroimage* 16(2), 484–512 (2002)
9. Smith, S.M., Nichols, T.E.: Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage* 44(1), 83–98 (2009)
10. Thirion, B., Tucholka, A., Keller, M., Pinel, P., Roche, A., Mangin, J.F., Poline, J.B.: High level group analysis of FMRI data based on Dirichlet process mixture models. In: Karssemeijer, N., Lelieveldt, B. (eds.) *IPMI 2007*. LNCS, vol. 4584, pp. 482–494. Springer, Heidelberg (2007)
11. Bowman, D.F., Caffo, B., Bassett, S.S., Kilts, C.: A bayesian hierarchical framework for spatial modeling of FMRI data. *Neuroimage* 39(1), 146–156 (2008)
12. Makni, S., Idier, J., Vincent, T., Thirion, B., Dehaene-Lambertz, G., Ciuciu, P.: A fully Bayesian approach to the parcel-based detection-estimation of brain activity in fMRI. *Neuroimage* 41(3), 941–969 (2008)
13. Penny, W.D., Trujillo-Barreto, N., Friston, K.J.: Bayesian fMRI time series analysis with spatial priors. *Neuroimage* 23(2), 350–362 (2005)
14. Friston, K., Ashburner, J., Frith, C., Poline, J.B., Heather, J., Frackowiak, R.: Spatial registration and normalization of images. *Hum. Brain Mapp.* 3(3), 165–189 (1995)
15. Chib, S.: Marginal likelihood from the Gibbs output. *J. Amer. Statist. Assoc.* 90, 1313–1321 (1995)
16. Kuhn, E., Lavielle, M.: Coupling a stochastic approximation of EM with a MCMC procedure. *ESAIM P&S* 8, 115–131 (2004)
17. Pinel, P., Thirion, B., Mériaux, S., Jobert, A., Serres, J., Le Bihan, D., Poline, J.B., Dehaene, S.: Fast reproducible identification and large-scale databasing of individual functional cognitive networks. *BMC Neurosci.* 8(1), 91 (2007)
18. Cavanna, A.E., Trimble, M.R.: The precuneus: a review of its functional anatomy and behavioural correlates. *Brain* 129(3), 564–583 (2006)