A BAYESIAN ESTIMATION FORMULATION TO VOXEL-BASED LESION-SYMPTOM MAPPING

Mame Diarra Fall ⋆*, Nicolas Dobigeon* †*, Pascal Auzou* ‡

⋆ Institut Denis Poisson, Université d'Orléans, Université de Tours, CNRS, 45100 Orléans, France.
† Université de Toulouse, IRIT/INP-ENSEEIHT, 31071 Toulouse, France.
‡ Centre Hospitalier Régional d'Orléans, Service de Neurologie, 14 avenue de l'hôpital, Orléans, France.

ABSTRACT

Studying brain-injured patients is important for investigating structure–function relationships using neuroimaging techniques. Voxel-based lesion-symptom mapping (VLSM) has increasingly been advocated as a relevant approach to detect structure–function associations in neuroimaging studies. The VLSM method involves mapping the relationship between brain injuries and behavioral performance on a voxel-by-voxel basis. This means that the statistical relationship between damage and behavior (across patients) is computed separately for each voxel. However, one could expect voxels characterizing group differences to be localized into spatially consistent regions rather than randomly distributed over the brain. Thus, in this paper, we propose to depart from conventional models to characterize and exploit this spatial consistency. More precisely, we derive a Bayesian model that explicitly accounts for spatial correlations between neighboring voxels using a Markov random field. Our results highlight that the proposed approach outperforms the conventional ones. Besides, it has the great advantage of possibly reducing the number of patients and identifying new language areas, which are two crucial insights in the targeted medical context.

Index Terms— Bayesian inference, Markov random field, stroke, LAST test.

1. INTRODUCTION

Voxel-based lesion-symptom mapping (VLSM) [1] aims at identifying the brain areas that are involved in some specific functions such as the language considered hereafter. Language pathology is a common and serious complication after stroke. According to epidemiological studies, nearly 25-30% of stroke patients develop such a defect. Lesion-based analysis has been mostly investigated under the same statistical framework: null hypothesis significance testing (NHST). The latter classically allows one to choose between two competing statistical hypotheses: H_0 (referred to as the null hypothesis) and H_1 (the alternative one). When specifically used in lesion-based analysis, NHST is conducted voxel-by-voxel and the voxels demonstrating significant difference are considered as being involved in language disorders.

In this paper, we follow a significantly different approach. Based on language screening test (LAST) scores and lesions maps from stroke patients, we directly estimate the language areas, without resorting to any hypothesis testing. Instead, the mapping of these brain regions of interest is formulated as an estimation problem which is solved within a Bayesian framework. Besides, this framework has the great advantage of offering the possibility to exploit spatial consistency of the regions within the brain by equipping the Bayesian model with a Markov random field (MRF). Instanced as an Ising potential, this prior allows to spatially regularize the mapping, departing from a more conventional spatially independent voxel-by-voxel analysis. It is particularly suitable when the number of patients is reduced to reduce uncertainty.

The sequel of this paper is organized as follows. The VLSM approach is reviewed in Section 2. Section 3 introduces the proposed Bayesian estimation model. The Gibbs sampler that generates samples asymptotically distributed according to the posterior distribution of the proposed model is derived in Section 4. The clinical data, their processing, the competing statistical tests and the numerical results are described in Section 5. Section 6 concludes this paper.

2. PROBLEM STATEMENT

Given a cohorte of N subjects, let y_n denote a score representing the task performance reached by the nth subject ($n = 1, \ldots, N$). In this paper, since we are mainly interested in locating the brain regions involved in the language, this score is considered as the outcome of the language screening test (LAST) [2]. Thus we denote by $\mathbf{y} = [y_1, \ldots, y_N]$ the observed LAST scores from the N subjects. In the LAST test, there are two subscores, namely an expression score (with a maximum of 8 points) and a receptive score (with a maximum of 7 points). The two subscores are summed and the corresponding total (with a maximum of 15) is considered as the observed score. When it is less than 15, the patient is considered as symptomatic (i.e., aphasic) while he/she is considered

as asymptomatic when the score reaches the maximum, i.e., for all $n = 1, \ldots, N$,

$$y_n = \begin{cases} \{0, \ldots, 14\}, & \text{if individual } \#n \text{ is symptomatic;} \\ \{15\}, & \text{if individual } \#n \text{ is asymptomatic.} \end{cases}$$

Within a lesion-based study, the lesions are assumed to be (manually or automatically) delineated for each subject to produce binary lesion maps which are co-registered to a common stereotaxic space. These maps are gathered in the $N \times J$-matrix $\mathbf{Z} = [\mathbf{z}_1, \ldots, \mathbf{z}_J]$ with $\mathbf{z}_j = [z_{1,j}, \ldots, z_{N,j}]^T$ (where J is the total number of voxels) such that

$$z_{n,j} = \begin{cases} 1, & \text{if voxel } \#j \text{ in patient } \#n \text{ is lesioned;} \\ 0, & \text{if voxel } \#j \text{ in patient } \#n \text{ is non-lesioned.} \end{cases}$$

Given these two sets of observations $\mathbf{y}$ and $\mathbf{Z}$, one aims at identifying the voxels composing the language areas[1]. In most of the lesion-based studies, a statistical analysis is conducted following NHST. For each voxel, patients are partitioned into two groups according the state of this voxel (i.e., lesioned vs. non-lesioned). Based on this partioning, two sets of scores can be defined according to group membership. More precisely, we denote

$$\mathbf{y}_1^j = \{y_n, n : z_{n,j} = 1\} \quad \text{and} \quad \mathbf{y}_2^j = \{y_n, n : z_{n,j} = 0\}$$

the sets of scores obtained by patients whose jth voxel is lesioned and non-lesioned, respectively, with

$$n_1^j = \sum_{n=1}^{N} \mathbf{1}(z_{n,j} = 1) \quad \text{and} \quad n_2^j = \sum_{n=1}^{N} \mathbf{1}(z_{n,j} = 0)$$

their respective cardinalities. In lesion-symptom mapping, one is interested in whether these two groups are significantly different. In other words, for each voxel $j \in \{1, \ldots, J\}$, given the sets of samples $\mathbf{y}_1^j \overset{iid}{\sim} F_1$ and $\mathbf{y}_2^j \overset{iid}{\sim} F_2$ with F_1 and F_2 unknown, one wants to decide between the two competing hypotheses

$$H_0^j : F_1 = F_2 \quad \text{versus} \quad H_1^j : F_1 \neq F_2.$$

The significant voxels, i.e, such that H_0^j is rejected, are considered as those belonging to language areas. In the medical literature, this two-sample problem is usually addressed via Student t-tests [1]. However, this strategy generally faces to two limitations that may question its reliability. First, the normality assumption does not always hold. Second, due to the fact that the number of subjects with lesions at a given voxel is often small, the central limit theorem does not apply. One alternative is to resort to nonparametric tests that do not rely on asymptotic results. For instance, such an approach has been adopted in [3] where a Bayesian nonparametric framework was derived. However, all aforementioned methods rely on a voxel-wise analysis, i.e., by analyzing each voxel independently of its neighbours. The aim of this work is to propose a more realistic spatial model that includes spatial interactions among neighbouring voxels. To this aim, the lesion-symptom mapping task is formulated as a statistical estimation problem, rather than a testing procedure. The language areas are located from stroke patients directly from the data. The next section describes the proposed Bayesian estimation method.

[1]Hereafter, we refer to "language areas" the brain areas involved in language disorder.

3. BAYESIAN ESTIMATION MODEL

Contrary to the statistical test-based procedure described above, locating the language areas is formulated as an inference task. It consists in estimating a set $\omega = [\omega_1, \ldots, \omega_J]^T$ of unobserved binary variables such that, for all $j = 1, \ldots, J$,

$$\omega_j = \begin{cases} 1, & \text{if voxel } \#j \text{ belongs to language areas;} \\ 0, & \text{otherwise.} \end{cases}$$

In what follows, $J_1 = \{1, \ldots, J\}$ denotes the set of indexes of voxels belonging to language areas with $J_1 = \text{card}\{J_1\}$. Conversely, $J_2 = \{1, \ldots, J\}$ denotes the set of indexes of voxels not belonging to language areas with $J_2 = \text{card}\{J_2\}$ with $J_1 + J_2 = J$. These labels are estimated within a Bayesian framework, whose key ingredients are introduced below.

Likelihood – In the LAST test, any patient with a total outcome less than 15 is considered as symptomatic. Thus, to simplify the analysis, we propose to threshold the overall score by introducing a set $\bar{\mathbf{y}} = [\bar{y}_1, \ldots, \bar{y}_N]$ of diagnosis variables such that, for all $n = 1, \ldots, N$,

$$\bar{y}_n = \begin{cases} 1, & \text{if individual } \#n \text{ is symptomatic} \\ 0 & \text{if individual } \#n \text{ is asymptomatic.} \end{cases}$$

In what follows, N_1 (resp. N_2) will denote the number of symptomatic (resp. asymptomatic) patients, ($N_1 + N_2 = N$). Then, the elements $z_{n,j}$ of the matrix $\mathbf{Z}$ defining the binary lesion maps are assumed to be independent and identically distributed according to Bernoulli distributions

$$z_{n,j} | \omega_j = 0 \quad \text{Ber}(\theta) \tag{1}$$

$$z_{n,j} | \omega_j = 1, \bar{y}_n = 1 \quad \text{Ber}(\theta_1) \tag{2}$$

$$z_{n,j} | \omega_j = 1, \bar{y}_n = 0 \quad \text{Ber}(\theta_0) \tag{3}$$

In (1), θ denotes the probability, across subjects, that a voxel outside the language area is lesioned. In (2) and (3), θ_1 and θ_0 denote the lesion probability in the language area for symptomatic and asymptomatic patients, respectively. The set of unknown parameters is $\boldsymbol{\Psi} = \{\omega, \theta, \theta_0, \theta_1\}$. We now endow them with prior distributions.

Prior modelling – We assume that the parameters ω, θ, θ_0 and θ_1 are a priori independent. The lesion probabilities are assigned Beta distributions

$$\theta \quad \text{Beta}(a, b) \tag{4}$$

$$\theta_0 \quad \text{Beta}(a_0, b_0) \tag{5}$$

$$\theta_1 \quad \text{Beta}(a_1, b_1). \tag{6}$$

Regarding the parameter of interest ω, it is quite legitimate to consider that the components ω_j and ω_k ($j \neq k$) are not independent but there are some spatial correlations between neighboring voxels. To model this correlation, we consider an Ising model with hyperparameter β

$$\mathrm{P}(\omega_j = k | \omega_{\setminus j}) \propto \exp\left[\beta \sum_{j' \in \mathcal{V}(j)} \delta(k, \omega_{j'})\right], \quad k \in \{0, 1\}$$

where $\omega_{\setminus j}$ is the vector ω whose jth element has been removed, $\mathcal{V}(j)$ denotes the set of indexes of neighboring voxels of the jth voxel and δ is the Kronecker symbol. The corresponding prior for ω can be written as

$$\mathrm{P}(\omega|\beta) = Z(\beta)^{-1} \exp\left[\sum_{j=1}^{J} \sum_{j' \in \mathcal{V}(j)} \beta \delta(\omega_j, \omega_{j'})\right]. \tag{7}$$

The amount of spatial correlation is controlled by the Ising field parameter β. This latter determines the level of spatial homogeneity between neighboring voxels. A value close to zero would imply that neighboring voxels are independent. In this work, its value will be assumed to be fixed and ajusted beforehand. Its estimation can be conducted following the strategies in [4–6].

4. GIBBS SAMPLER

By using the Bayes rule, the posterior distribution can be expressed as follows

$$\begin{aligned}\pi(\boldsymbol{\Psi}|\mathbf{Z}, \bar{\mathbf{y}}) &\propto f(\mathbf{Z}|\omega, \bar{\mathbf{y}}, \theta, \theta_0, \theta_1)\pi(\theta, \theta_0, \theta_1, \omega) \\ &\propto f(\mathbf{Z}|\omega, \bar{\mathbf{y}}, \theta, \theta_0, \theta_1)\pi(\theta)\pi(\theta_0)\pi(\theta_1)\pi(\omega)\end{aligned}$$

where the likelihood $f(\mathbf{z}|\omega, \bar{\mathbf{y}}, \theta, \theta_0, \theta_1)$ has been defined in (1), (2) and (3), and the prior distributions $\pi(\theta)$, $\pi(\theta_0)$, $\pi(\theta_1)$ and $\pi(\omega|\beta)$ in (4), (5), (6), and (7). It leads to

$$\begin{aligned}\pi(\boldsymbol{\Psi}|\mathbf{Z}, \bar{\mathbf{y}}) \propto\ & \prod_{n=1}^{N} \prod_{j \in \mathcal{J}_2} \theta^{z_{nj}}(1-\theta)^{1-z_{nj}} \prod_{n=1}^{N_1} \prod_{j \in \mathcal{J}_1} \theta_1^{z_{nj}}(1-\theta_1)^{1-z_{nj}} \\ & \prod_{n=1}^{N_2} \prod_{j \in \mathcal{J}_1} \theta_0^{z_{nj}}(1-\theta_0)^{1-z_{nj}} \theta^{a-1}(1-\theta)^{b-1}\theta_0^{a_0-1}(1-\theta_0)^{b_0-1} \\ & \theta_1^{a_1-1}(1-\theta_1)^{b_1-1} \exp\left[\beta \sum_{j=1}^{J} \sum_{j' \in \mathcal{V}(j)} \delta(\omega_j, \omega_{j'})\right].\end{aligned}$$

Since computing the Bayesian estimates associated to this posterior distribution is not straightforward, we propose to generate samples asymptotically distributed according to this posterior using a Gibbs sampler. It consists in drawing samples according to the full conditional of each parameter successively. These steps are detailed below.

Sampling according to $\pi(\theta|\omega, \mathbf{Z})$ – The posterior probability of having a lesioned voxel outside the language area is

$$\begin{aligned}\pi(\theta|\omega, \mathbf{Z}) &\propto \theta^{a-1}(1-\theta)^{b-1} \prod_{n=1}^{N} \prod_{j \in \mathcal{J}_2} \theta^{z_{nj}}(1-\theta)^{1-z_{nj}} \\ &\propto \theta^{a+(\sum_{n=1}^{N}\sum_{j \in \mathcal{J}_2} z_{nj})-1}(1-\theta)^{b+NJ_2-(\sum_{n=1}^{N}\sum_{j \in \mathcal{J}_2} z_{nj})-1}.\end{aligned}$$

This conditional is recognisable as a Beta distribution with parameters $a + \sum_{n=1}^{N}\sum_{j \in \mathcal{J}_2} z_{nj}$ and $b + NJ_2 - \sum_{n=1}^{N}\sum_{j \in \mathcal{J}_2} z_{nj}$.

Sampling according to $\pi(\theta_1|\omega, \mathbf{Z})$ – The posterior probability of having a lesioned voxel in language areas of symptomatic subjects is

$$\begin{aligned}\pi(\theta_1|\omega, \mathbf{Z}) &\propto \theta_1^{\alpha-1}(1-\theta_1)^{b_1-1} \prod_{n=1}^{N_1} \prod_{j \in \mathcal{J}_1} \theta_1^{z_{nj}}(1-\theta_1)^{1-z_{nj}} \\ &\propto \theta_1^{a_1+\sum_{n=1}^{N_1}\sum_{j \in \mathcal{J}_1} z_{nj}-1}(1-\theta_1)^{b_1+N_1J_1-\sum_{n=1}^{N_1}\sum_{j \in \mathcal{J}_1} z_{nj}-1}.\end{aligned}$$

This is also recognisable as a Beta distribution with parameters $a_1 + \sum_{n=1}^{N_1}\sum_{j \in \mathcal{J}_1} z_{nj}$ and $b_1 + N_1J_1 - \sum_{n=1}^{N_1}\sum_{j \in \mathcal{J}_1} z_{nj}$.

Sampling according to $\pi(\theta_0|\omega, \mathbf{Z})$ – The posterior probability of having a lesioned voxel in language areas of asymptomatic subjects is

$$\begin{aligned}\pi(\theta_0|\ldots) &\propto \theta_0^{a_0-1}(1-\theta_0)^{b_0-1} \prod_{n=1}^{N_2} \prod_{j \in \mathcal{J}_1} \theta_0^{z_{nj}}(1-\theta_0)^{1-z_{nj}} \\ &\propto \theta_0^{a_0+\sum_{n=1}^{N_2}\sum_{j \in \mathcal{J}_1} z_{nj}-1}(1-\theta_0)^{b_0+N_2J_1-\sum_{n=1}^{N_2}\sum_{j \in \mathcal{J}_1} z_{nj}-1}.\end{aligned}$$

This conditional is again a Beta distribution with parameters $a_0 + \sum_{n=1}^{N_2}\sum_{j \in \mathcal{J}_1} z_{nj}$ and $b_0 + N_2J_1 - \sum_{n=1}^{N_2}\sum_{j \in \mathcal{J}_1} z_{nj}$.

Sampling according to $\pi(\omega|\theta, \theta_0, \theta_1, \mathbf{y}, \mathbf{Z})$ – The vector ω can be updated coordinate-by-coordinate using Gibbs moves. For each voxel $j \in \{1, 2, \ldots, J\}$, ω_j is a binary variable whose conditional distribution is defined by the probability

$$\mathrm{P}(\omega_j = 1|\omega_{\setminus j}, \mathbf{z}_j, \bar{\mathbf{y}}) = \frac{p_1}{p_1 + p_0}$$

where, for $k \in \{0, 1\}$,

$$p_k = \mathrm{P}(\omega_j = k|\omega_{\setminus j}) f(\mathbf{z}_j|\omega_j = k, \bar{\mathbf{y}})$$

with

$$\mathrm{P}(\omega_j = k|\omega_{\setminus j}) \propto \exp\left[\beta \sum_{j' \in \mathcal{V}(j)} \delta(k, \omega_{j'})\right]$$

$$f(\mathbf{z}_j|\omega_j = 0, \theta, \bar{\mathbf{y}}) = \prod_{n=1}^{N} \theta^{z_{nj}}(1-\theta)^{1-z_{nj}} = \theta^{\sum_{n=1}^{N} z_{nj}}(1-\theta)^{N-\sum_{n=1}^{N} z_{nj}},$$

$$f(\mathbf{z}_j|\omega_j = 1, \theta_0, \theta_1, \bar{\mathbf{y}}) = \prod_{\{n:\bar{y}_n=1\}} \theta_1^{z_{nj}}(1-\theta_1)^{1-z_{nj}} + \prod_{\{n:\bar{y}_n=0\}} \theta_0^{z_{nj}}(1-\theta_0)^{1-z_{nj}}.$$

Note that the posterior distribution of $\boldsymbol{\omega}$ also defines a MRF.

We next provide the numerical results demonstrating the suitability of the proposed approach to VLSM studies.

5. EXPERIMENTS

Data and image processing – The analyzed data were collected from $N = 58$ participants (47 men, 11 women) who had suffered a single left-hemispheric stroke in the acute phase (< 7 days) and admitted at the Neurology Department of the Orléans Hospital in France. All patients, regardless of the arterial distribution of their stroke, were included. The mean age standard deviation is 66.1 13.4 years. The range is 19-91 years. The patients were evaluated with the LAST test to provide the threshold score vector $\bar{\mathbf{y}}$. MRI was carried out and we used the diffusion-weighted imaging (DWI) and the fluid-attenuated inversion-recovery (FLAIR) sequences. The lesions were drawn directly on each patient's DWI or Flair digital MRI image (choosing the best contrasted one) using OSIRIX software. DWI or FLAIR images were yoked to the T1 images so that the extent of the lesion could be verified on these image sequences. Then, we obtain a mask for each patient. MRI images were registered into MNI space (standard template of the Montreal Neurological Institute) using the standard nonlinear spatial normalization procedure from Statistical Parametric Mapping (SPM12) (Wellcome Trust Centre for Neuroimaging) running under Matlab 2017a. We re-aligned and co-registered 3D images with a 5-th Degree B-Spline interpolation method in SPM12 and then averaged them. Masks were re-sliced and normalized to the native space of the averaged 3D images with trilinear interpolation by voxels of 1mm^3. We finally obtain the matrix $\mathbf{Z}$ gathering the N binary maps locating the lesioned and non-lesioned voxels.

Compared methods – We aim at assessing the performance of the proposed Bayesian estimation method in contrast to some competing methods. The proposed model is compared to some frequentist tests: the t-test used in the original VLSM of [1], two common frequentist nonparametric procedures, namely the Kolmogorov-Smirnov and the Mann-Whitney tests, and the Bayesian nonparametric test proposed in [3]. The total number of analysed voxels in the competitive methods is $V = 7.109.137$.

The original VLSM technique has been implemented using the nonparametric mapping (NPM) software, distributed as part of the MRIcron toolset. In VLSM, t-tests are performed in each voxel. It is usual to confine tests to voxels in which there are at least five patients with a lesion and five patients without a lesion; this is the approach we have taken in the current study for all frequentist tests. The significance level is 5%. As for the BNP test, the procedure described in [3] has been applied.

Regarding the proposed estimation model we consider in this preliminary study a 2D MRF model. The 2D MRF are considered for single slices of size 181 227. A 4-pixel neighbourhood structure is considered for the MRF prior model. The hyperparameter β is set to 2.2. This value has been chosen by testing several values of β on synthetic data (results not shown here due to lack of space) and choosing the one that minimizes the mean squared error. Future work will study the estimation of β jointly with the other unknown parameters of the model. The other hyperparameters have been chosen as follows in agreement with the expected values by the neurologists: $a = b = a_0 = b_0 = a_1 = b_1 = 10^{-3}$; $\theta = 0.15$, $\theta_1 = 0.75$, $\theta_0 = 0.02$. The Gibbs sampler has been run using 1000 iterations including 500 iterations of burn-in period. We obtain the following posterior mean estimates using MCMC post burn-in samples: $\hat{\theta} = 0.0072$, $\hat{\theta_0} = 0.0012$ and $\hat{\theta_1} = 0.1543$. A trivial choice for the point estimate $\boldsymbol{\omega}$ of the map $\boldsymbol{\omega}$ would be the posterior mean. However as mentioned in [7], for many types of images, particularly binary images, the use of posterior means does not provide sensible solutions. Other possibilities are available, such as the maximum a posteriori (MAP) estimators or the marginal posterior modes (MPM) considered in this work. The jth component of the MPM vector $\boldsymbol{\omega}^{\mathrm{MPM}}$ is the modal value of the posterior marginal. It can be computed as

$$\hat{\omega}_j^{\mathrm{MPM}} = \begin{cases} 1 \text{ if } \mathrm{P}(\omega_j = 1|\theta, \theta_0, \theta_1, \mathbf{y}, \mathbf{Z}) > 1/2 \\ 0 \text{ if } \mathrm{P}(\omega_j = 1|\theta, \theta_0, \theta_1, \mathbf{y}, \mathbf{Z}) \quad 1/2 \end{cases}$$

Given M samples $\{\omega_j^{(1)}, \ldots, \omega_j^{(M)}\}$ produced by the Gibbs sampler, this MPM can be approximated as

$$\hat{\omega}_j^{\mathrm{MPM}} = \begin{cases} 1 \text{ if } \mathrm{card}\{\omega_j^{(m)}|\hat{\omega}_j^{(m)} = 1\} \quad M/2 \\ 0 \text{ otherwise.} \end{cases}$$

Numerical Results – As in [3], we first investigated the results obtained using the total number of patients, $n = 58$. Results not shown here reveal that all competing methods located more or less the classical language regions known as Broca's and Wernicke's areas. The former is involved mostly in the production of speech, while receptive speech has traditionally been associated with the latter. Second, since our goal was to reduce the number of patients involved in VLSM studies, we tried smaller sample sizes for which the uncertainty would be higher. The most remarkable results are the