



HAL
open science

Bayesian Spatiotemporal Segmentation of Combined PET-CT Data Using a Bivariate Poisson Mixture Model

Zacharie Irace, Hadj Batatia

► **To cite this version:**

Zacharie Irace, Hadj Batatia. Bayesian Spatiotemporal Segmentation of Combined PET-CT Data Using a Bivariate Poisson Mixture Model. 22rd European Signal and Image Processing Conference (EUSIPCO 2014), European Association for Signal Processing (EURASIP), Sep 2014, Lisbonne, Portugal. pp.1–6. hal-03256422

HAL Id: hal-03256422

<https://hal.science/hal-03256422>

Submitted on 11 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

BAYESIAN SPATIOTEMPORAL SEGMENTATION OF COMBINED PET-CT DATA USING A BIVARIATE POISSON MIXTURE MODEL

Zacharie Irace, Hadj Batatia

University of Toulouse, IRIT/INP-ENSEEIH, 2, rue Charles Camichel, 31071 Toulouse, France

ABSTRACT

This paper presents an unsupervised algorithm for the joint segmentation of 4-D PET-CT images. The proposed method is based on a bivariate-Poisson mixture model to represent the bimodal data. A Bayesian framework is developed to label the voxels as well as jointly estimate the parameters of the mixture model. A generalized four-dimensional Potts-Markov Random Field (MRF) has been incorporated into the method to represent the spatio-temporal coherence of the mixture components. The method is successfully applied to 4-D registered PET-CT data of a patient with lung cancer. Results show that the proposed model fits accurately the data and allows the segmentation of different tissues and the identification of tumors in temporal series.

Index Terms— multimodality, data fusion, 4-D segmentation, PET-CT, bivariate Poisson distribution

1. INTRODUCTION

Cancer diagnosis and its treatment require precise localization of lesions before any quantification. In clinics, tumors are manually or semi-automatically delineated in Positron Emission Tomography (PET) and Computed Tomography (CT) images, separately. Robust automatic segmentation methods are therefore necessary to help medical decision making. However, the segmentation of PET images is a difficult task due to their poor contrast and lack of anatomical information. PET and CT images are usually fused a posteriori to allow the visual localization of tumors. The advent of combined imaging systems that allow the acquisition of registered PET-CT images makes feasible the fusion of functional and anatomical modalities during segmentation [1, 2]. However, research efforts have focused more on labeling techniques in separate modality; only a few recent studies have been devoted to the automatic segmentation based on both modalities.

Wojak et al. [3] proposed a segmentation technique on 3-D bimodal (PET-CT). The method uses PET data to guide the segmentation of CT images. Ballangan et al. [4] proposed a technique where the tumor region obtained from PET is used as a mask to resolve a 2-component Gaussian mixture model in the corresponding CT image. These two methods do not make full use of PET and CT modalities. They rather

use the PET image to inform the segmentation of the CT image. Han et al. [5] formulated the problem as a graph-based segmentation. Their method consists in minimizing a binary Markov Random Field (MRF) energy for both PET and CT images and penalizing the segmentation difference between the two images. However, the method requires user interaction. More recently, Bagci et al. [6] presented an automated graph-based method for simultaneously segmenting functional and anatomical structures. Gribben et al. [7] performed a multimodal segmentation by using a maximum a posteriori Markov random field (MAP-MRF) approach. A vectorial representation is used for CT and PET data. The voxel intensities are assumed to follow a bivariate Gaussian distribution. A 3-D Potts model is proposed as prior spatial information. An EM algorithm is used to estimate the resulting posterior.

The problem of accurate segmentation of PET-CT images is particularly exacerbated when target organs exhibit natural motions, such as breathing. In such situation, prior spatio-temporal registration between the two modalities is required [8, 9]. Spatio-temporal segmentation of coherent 4D PET-CT offer the advantage of localizing and tracking the tumors, especially important for radiotherapy applications. Recently, Bai et al. [10] proposed a co-segmentation of 4D CT and PET images, based on MRF including a regularization between PET and CT. Their method consists in segmenting the CT image based on PET data as prior using a single maximum flow optimization algorithm.

This paper proposes a fully automatic technique to label voxels in 4D bimodal PET-CT images. A hierarchical Bayesian model is proposed where the bimodal data is represented as a mixture of bivariate Poisson distributions. A 4-D Markov Random Field prior is incorporated to enforce the spatio-temporal correlation of voxels. An MCMC algorithm is developed to jointly estimate the mixture model parameters as well as classification labels of the voxels.

2. BAYESIAN FRAMEWORK

Let \mathbf{x} be a 4-D PET-CT image (a time series of 3-D images). The value of the n_{th} voxel $\mathbf{x}_n \in \mathbb{N}^2$ is written as a 2-component vector corresponding to the CT and PET data $\mathbf{x}_n = (x_n^{\text{CT}}, x_n^{\text{PET}})^T$. The image \mathbf{x} is assumed to be composed of K distinct biological tissues $\{C_1, \dots, C_K\}$. Let $\mathbf{z} =$

$\{z_1, \dots, z_N\}$ be a hidden label variable defined such as $z_n = k$ if $x_n \in C_k$. The image segmentation problem can be formulated as a *maximum a posteriori* (MAP) problem:

$$\hat{z} = \underset{z}{\operatorname{argmax}} p(\theta, z | x)$$

where θ is an unknown parameter vector. This problem can be solved in a Bayesian framework by estimating the posterior and jointly estimating θ and z . For this purpose, the data likelihood and priors are established in the following sections.

2.1. Data likelihood model

In [7], a bivariate-Gaussian mixture has been used to model PET-CT bimodal data. However, the Poisson mixture model (PMM) has been considered as more appropriate to model the PET projections and PET reconstructed images [11], [12]. CT data is also considered to follow a mixture of Poisson distributions [13]. Given a class k , let X_k^{CT} and X_k^{PET} the random variables associated to the CT and PET data respectively. Considering that each modality follows a Poisson distribution, if the two channels are independent, the joint likelihood can be expressed as the product of the two distributions. However, it has been established that PET and CT data are intrinsically dependent [7]. Accordingly, we propose a bivariate Poisson model for the bimodal data to take into consideration their dependence. The data is considered as the result of the following process:

$$X_k^{\text{CT}} = X_k^1 + X_k^0 \text{ and } X_k^{\text{PET}} = X_k^2 + X_k^0$$

where X_k^i are Poisson random variables:

$$\forall i, X_k^i \sim \mathcal{P}(\theta_i^k)$$

and X_k^0 represents the common dependence between PET and CT intensities. This yields the Poisson marginal distributions,

$$X_k^{\text{CT}} \sim \mathcal{P}(\theta_1^k + \theta_0^k) \text{ and } X_k^{\text{PET}} \sim \mathcal{P}(\theta_2^k + \theta_0^k)$$

Consequently, the bimodal PET-CT data in the class C_k follows a bivariate Poisson distribution:

$$(X_k^{\text{CT}}, X_k^{\text{PET}}) \sim \mathcal{BP}(\theta_1^k, \theta_2^k, \theta_0^k)$$

where

$$p(x_n | \theta_1^k, \theta_2^k, \theta_0^k, z_n = k) = e^{-(\theta_1^k + \theta_2^k + \theta_0^k)} \frac{(\theta_1^k)^{x_n^{\text{CT}}} (\theta_2^k)^{x_n^{\text{PET}}}}{x_n^{\text{CT}}! x_n^{\text{PET}}!} \cdot \sum_{j=0}^{\min(x_n^{\text{CT}}, x_n^{\text{PET}})} \binom{x_n^{\text{CT}}}{j} \binom{x_n^{\text{PET}}}{j} j! \left(\frac{\theta_0^k}{\theta_1^k \theta_2^k} \right)^j$$

By assuming that all the observations from each channel are independent, the likelihood can be defined as:

$$p(x | \theta, z) = \prod_{k=1}^K \prod_{n \in \mathcal{I}_k} p(x_n | \theta^k, z_n = k) \quad (1)$$

where \mathcal{I}_k is the set of voxels belonging to class C_k and $\theta^k = (\theta_1^k, \theta_2^k, \theta_0^k)$.

2.2. Prior distributions of θ

The matrix $\theta = \{\theta_i^k\}_{i=0,1,2; k=1 \dots K}$ is the set of unknown parameters of the proposed model. In this study, the number of classes K is assumed to be known. An inverse Gamma distribution with hyperparameters a_0 and b_0 is chosen as a prior distribution for each θ_i^k

$$\theta_i^k \sim \mathcal{IG}(a_0, b_0), \quad \forall k = 1 \dots K, \forall i = 0 \dots 2,$$

The hyperparameters are set to $a_0 = 1$ and $b_0 = 1$, leading to a flat non-informative prior. Assuming independent mixture parameters, the joint prior distribution for θ is written:

$$p(\theta) = \prod_{k=1}^K \prod_{i=0}^2 p(\theta_i^k) \quad (2)$$

2.3. Prior distribution of z

It is natural to consider that the label of a voxel is correlated with those of its neighbors. In this study, a first order 4-D Potts Markov random field is proposed for the prior density of z . Each voxel is then connected to its 6 spatial neighbors and 2 temporal neighbors (see figure 1). The prior density on




Fig. 1. 4-D neighborhood

the labels is expressed as the Potts distribution:

$$p(z) = \frac{1}{C(\beta)} \exp \left[\sum_{n=1}^N \sum_{n' \in \mathcal{V}(n)} \beta \delta(z_n - z_{n'}) \right] \quad (3)$$

where β is the granularity coefficient, $C(\beta)$ is the partition function, $\delta(\cdot)$ is the Kronecker function and $\mathcal{V}(\cdot)$ represents the neighborhood structure (Fig. 1).

2.4. Posterior distribution $p(\theta, z | x)$

Considering the parameter vector θ and z independent, using Bayes' theorem, we can express the posterior distribution of the vector (z, θ) as:

$$p(\theta, z | x) \propto p(x | \theta, z) p(\theta) p(z) \quad (4)$$

where the likelihood $p(x | \theta, z)$ and the prior distributions $p(\theta)$ et $p(z)$ have been defined in (1), (2) and (3) respectively.

3. HYBRID GIBBS SAMPLER

A Metropolis-within-Gibbs sampler is proposed to solve the segmentation problem. Samples are iteratively drawn according to the conditional densities of the posterior (4) and then used to estimate the maximum a posteriori (MAP). The algorithm (1) summarizes the proposed process. For more details on MCMC methods, the reader is referred to [14].

Algorithm 1 Proposed Gibbs sampler

Initialisation:

- Draw $\theta^{(0)}$ according to (2).
- Generate $z_1^{(0)}, z_2^{(0)}, \dots, z_N^{(0)}$ randomly.

for $t = 1, 2, \dots$ to T **do**

– Update θ –

for $k = 1, 2, \dots$ to K **do**

1. Propose $\theta^{k*} \sim \mathcal{TN}(\theta^{k(t-1)}, \epsilon_k \cdot \mathbb{I}_3)$ (see (6)).
2. Compute the acceptance ratio a (see (3.2)).
3. Draw $u \sim \mathcal{U}(0, 1)$.

if $(u < a)$ **then**

4. Set $\theta^{k(t)} = \theta^{k*}$.

else

5. Set $\theta^{k(t)} = \theta^{k(t-1)}$.

end if

end for

– Update z –

for $n = 1, 2, \dots$ to N **do**

7. Draw z_n in $\{1, \dots, K\}$ according to (5).

end for

end for

3.1. Conditional distribution $p(z|\theta, \mathbf{x})$

The conditional distribution of the discrete label z_n can be expressed as:

$$p(z_n = k | x_n, \theta_k, \mathbf{z}_{-n}) \propto p(x_n | \theta^k, z_n = k) p(z_n | \mathbf{z}_{-n}) \quad (5)$$

where $k = 1, \dots, K$ and \mathbf{z}_{-n} represents the vector \mathbf{z} whose n^{th} element has been deleted. Because (5) defines a Markov random field, samples are generated according to this conditional distribution by randomly drawing a discrete value in the finite set $\{1, \dots, K\}$ with probabilities (5).

3.2. Conditional distribution $p(\theta|z, \mathbf{x})$

To generate samples asymptotically distributed according to $p(\theta|z, \mathbf{x})$, a Metropolis-Hastings (MH) algorithm is used, leading to a Metropolis-within-Gibbs algorithm [14, p. 317]. More precisely, θ is updated by the mean of a Random-Walk Metropolis-Hastings (RWMH) algorithm [14, p. 245] with the following proposal distribution:

$$\theta^{k*} \sim \mathcal{TN}(\theta^{k(t-1)}, \epsilon_k \cdot \mathbb{I}_3). \quad (6)$$




Fig. 2. Data fit. (a) 2-D histogram of the data ; (b) Estimated Bivariate Poisson Mixture Model

where \mathcal{TN} refers to the trivariate normal distribution, $\theta_k^{(t-1)}$ is the last value of the chain and the parameter ϵ_k is chosen so the acceptance ratio tends to $\frac{1}{2}$, as recommended in [15]. Moreover, by considering the fact that the proposal is symmetrical, the acceptance ratio can be expressed as the product of the likelihood ratio and the prior ratio:

$$a = \min \left\{ 1, \prod_{n: z_n = k} \frac{p(x_n | \theta^{k*}, z_n = k) p(\theta^{k*})}{p(x_n | \theta^{k(t-1)}, z_n = k) p(\theta^{k(t-1)})} \right\}$$

where the prior $p(\theta_k)$ is defined in (2).

4. EXPERIMENTAL RESULTS

Experiments have been made on real chest 4-D PET-CT data, reconstructed according to the methodology described in [8, 9]. The original PET and CT images have been obtained by a GE Discovery ST scanner on a patient suffering from a lung cancer. Six images have been reconstructed in each modality, corresponding to different temporal bins in the respiratory cycle. The CT images have been re-scaled by linear interpolation to match the dimensions of the PET images. Finally, both images have been registered and cropped to a $(56 \times 91 \times 11)$ voxels sized Region Of Interest (ROI). Results from fitting various models to the data are presented as well as the associated labeling results.

4.1. Validation of the data model

This section reports results on how well the proposed model fits the real data. First, figure 2 represents the 2-D histogram of the measured data and the corresponding estimated model. One notice that the model closely fits the data.

The goodness-of-fit of the proposed model has been assess quantitatively and compared to other models from the literature. The proposed algorithm has been applied on the data by considering mixture models of different statistical distributions:

- $\mathcal{N} \times \mathcal{N}$: mixture of normal distributions for PET and CT channels independently.

- \mathcal{BN} : mixture of bivariate normal distributions without independence assumption.
- $\mathcal{P} \times \mathcal{P}$: mixture of Poisson distributions for PET and CT channels independently.
- \mathcal{BP} : the proposed mixture of bivariate Poisson distributions.

The MAP values of each parameter has been measured as the mean of 6000 generated samples of the Monte-Carlo chains, for $K = 5$ classes. The measured log-likelihood, the Akaike information criterion (AIC) as well as the Bayes information criterion (BIC) are reported in table 1. According to the three

	log-likelihood	AIC	BIC
$\mathcal{N} \times \mathcal{N}$	-2.4485	5.5697	9.1774
\mathcal{BN}	-2.4433	5.5594	9.1672
$\mathcal{P} \times \mathcal{P}$	-2.4222	5.5170	9.1246
\mathcal{BP}	-2.3875	5.4476	9.0553

Table 1. log-likelihood ($\times 10^6$), AIC ($\times 10^6$) and BIC ($\times 10^6$) for each mixture model.

criteria, the proposed bivariate Poisson mixture fits better the data, although all four models provide interesting fits. It is particularly interesting to notice that the two Poisson mixture models have better scores than the two normal mixture models, although the normal models have more parameters making them more flexible. Moreover, the relative advantage of the bivariate Poisson over the univariate Poisson model confirms the dependence assumption.

4.2. Segmentation results

Figure 3 presents the original PET (Fig.3.a) and CT (Fig.3.b) images as well as the results of the separate segmentation by the univariate Poisson mixture applied to PET (Fig.3.c) and CT (Fig.3.d), respectively. The PET results show anatomically incoherent classes due to the coarse resolution and lack of anatomical information. In the CT image, misclassifications are visible (different tissues associated to same classes). For instance, background pixels and the tumor are inaccurately associated to the same class.

Figures (3.e) to (3.h) report the labeling results obtained by the four models described in section 4.1. It is visually obvious that the segmentation results of both bivariate and univariate Poisson distributions are better than those of the normal distributions. The shapes of the classes are in better accordance with anatomical structures. And interestingly, regions showing different functional activities (see the original PET image) have distinct labels. This indicates the benefit of PET and CT bimodal data fusion during segmentation.

Moreover, the proposed bivariate Poisson model (\mathcal{BP}) gives better results (Fig.3.h) than the univariate Poisson ($\mathcal{P} \times \mathcal{P}$) (Fig.3.g). One can notice that voxels corresponding to the




Fig. 3. Segmentation results. (a) Original PET channel ; (b) Original CT channel ; (c) Segmentation of the PET channel only, based on a mixture of Poisson model ; (d) Segmentation of the CT channel only, based on a mixture of Poisson model ; (e) to (h) Segmentation of both channels based on the 4 tested mixture models

bones (see the white zones in the original CT image) belong to the same class in the result of the \mathcal{BP} segmentation (3.h). This indicates that the level of dependence between the PET and CT data is significant and correctly taken into consideration by the \mathcal{BP} model. Another remarkable result is the single class that surrounds the tumor in both \mathcal{BP} and ($\mathcal{P} \times \mathcal{P}$) segmentations. This class seems to represent the partial volume effect that produces statistics distinct from the tumor and the healthy tissues.

4.3. Discussion

The goodness of fit reported in table (1) showed that the univariate and especially bivariate Poisson models fit slightly better the data than the univariate and bivariate normal mixtures. However, the two Poisson models give better segmentation results and the bivariate Poisson mixture shows the importance of considering the dependence of the two modalities (PET and CT) when doing bimodal joint segmentation.

One might consider the poor segmentation of the normal mixtures paradoxical with their good enough fit to the data. This can be explained by the flexibility of the normal distribution given the number of its parameters compared to Poisson. This flexibility leads to faster estimation of fitting parameters without being able to find optimal labels. This behavior suggests that the labels space has not been fully explored by the Markov chains for the normal based models, favoring the estimation of the model parameters, at the expense of the labels.

5. CONCLUSION AND FUTURE WORK

This paper proposed a bivariate Poisson mixture model to perform multimodal data fusion of PET and CT images. A hierarchical Bayesian model has been elaborated that includes a 4-D Potts-Markov field to consider spatial and temporal coherence in a time series of PET-CT images. An MCMC algorithm based on Metropolis-within-Gibbs sampler has been developed to jointly estimate the parameters of the Bayesian model while segmenting distinct tissues. In addition to the proposed model, this work brings two contributions consisting in performing statistical multimodal data fusion, and spatio-temporal dynamic images segmentation. The method has been applied to real four-dimensional PET-CT images of a patient having a lung tumor. The proposed model shows a good goodness of fit compared to three other models. The segmentation results obtained are visually in excellent concordance with anatomical structures while delineating the tumor.

The proposed framework can be generalized to other modalities and other statistical models can be integrated. Future work will consider distributions that handle heterogeneity within tumors in PET data [16]. In addition, more reliable sampling methods [17] will be investigated to consider the unsymmetrical role of the parameters θ_1 , θ_2 and θ_0 , which slows the exploration of the posterior distribution.

REFERENCES

- [1] Gerald Antoch, Jrg Stattaus, Andre T. Nemat, Simone Marnitz, Thomas Beyer, Hilmar Kuehl, Andreas Bockisch, Jrg F. Debatin, and Lutz S. Freudenberg, "Non-small cell lung cancer: Dual-modality pet/ct in preoperative staging," *Radiology*, vol. 229, no. 2, pp. 526–533, 2003, PMID: 14512512.
- [2] Sung Shine Shim, Kyung Soo Lee, Byung-Tae Kim, Myung Jin Chung, Eun Jung Lee, Joungho Han, Joon Young Choi, O Jung Kwon, Young Mog Shim, and Seonwoo Kim, "Non-small cell lung cancer: Prospective comparison of integrated fdg pet/ct and ct alone for preoperative staging," *Radiology*, vol. 236, no. 3, pp. 1011–1019, 2005, PMID: 16014441.
- [3] J. Wojak, E.D. Angelini, and I. Bloch, "Joint variational segmentation of ct-pet data for tumoral lesions," in *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*, April 2010, pp. 217–220.
- [4] C. Ballangan, Xiuying Wang, M. Fulham, S. Eberl, and Dagan Feng, "Lung tumor delineation in pet-ct images using a downhill region growing and a gaussian mixture model," in *Image Processing (ICIP), 2011 18th IEEE International Conference on*, Sept 2011, pp. 2173–2176.
- [5] Dongfeng Han, John Bayouth, Qi Song, Aakant Taurani, Milan Sonka, John Buatti, and Xiaodong Wu, "Globally optimal tumor segmentation in pet-ct images: A graph-based co-segmentation method," in *Information Processing in Medical Imaging*, Gbor Szekely and HorstK. Hahn, Eds., vol. 6801 of *Lecture Notes in Computer Science*, pp. 245–256. Springer Berlin Heidelberg, 2011.
- [6] Ulas Bagci, JayaramK. Udupa, Jianhua Yao, and DanielJ. Mollura, "Co-segmentation of functional and anatomical images," in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2012*, Nicholas Ayache, Herv Delingette, Polina Golland, and Kensaku Mori, Eds., vol. 7512 of *Lecture Notes in Computer Science*, pp. 459–467. Springer Berlin Heidelberg, 2012.
- [7] H. Gribben, P. Miller, G.G. Hanna, K.J. Carson, and A.R. Hounsell, "Map-mrf segmentation of lung tumours in pet/ct images," in *Biomedical Imaging: From Nano to Macro, 2009. ISBI '09. IEEE International Symposium on*, June 2009, pp. 290–293.
- [8] Z. Ouksili and H. Batatia, "4d ct image reconstruction based on interpolated optical flow fields," in *IEEE Int. Conf. on Image Processing (ICIP 2010)*, Hong Kong, Sep 2010, pp. 633 – 636.
- [9] D. Didierlaurent, S. Ribes, H. Batatia, C. Jaudet, L.O. Dierickx, S. Zerdoud, S. Brillouet, O. Caselles, and F. Courbon, "The retrospective binning method improves the consistency of phase binning in respiratory-gated pet/ct," *Physics in Medicine and Biology*, vol. 57, no. 23, pp. 7829–7841, novembre 2012.
- [10] Junjie Bai, Qi Song, Sudershan K. Bhatia, and Xiaodong Wu, "Globally optimal lung tumor co-segmentation of 4d ct and pet images," 2013.
- [11] A. Rangarajan, I.T. Hsiao, and G. Gindi, "A Bayesian joint mixture framework for the integration of anatomical information in functional image reconstruction," *J. Math. Im. Vis.*, vol. 12, pp. 199–217, 2000.
- [12] N. Pustelnik, C. Chaux, and J.-C. Pesquet, "A wavelet-based quadratic extension method for image deconvolution in the presence of poisson noise," in *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*, april 2009, pp. 701 –704.
- [13] I. A. Elbakri and J. A. Fessler, "Statistical image reconstruction for polyeenergetic x-ray computed tomography," *IEEE Trans. Med. Imaging*, vol. 21, no. 2, pp. 89–99, Feb 2002.
- [14] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, Springer-Verlag, New York, 1999.
- [15] G. O. Roberts, "Markov chain concepts related to sampling algorithms," in *Markov Chain Monte Carlo in Practice*, W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, Eds., pp. 259–273. Chapman & Hall, London, 1996.
- [16] Z. Irace, M. Pereyra, N. Dobigeon, and H. Batatia, "Bayesian segmentation of chest tumors in pet scans using a poisson-gamma mixture model," in *Statistical Signal Processing Workshop (SSP), 2011 IEEE*, June 2011, pp. 809–812.
- [17] Mark Girolami and Ben Calderhead, "Riemann manifold langevin and hamiltonian monte carlo methods," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 73, no. 2, pp. 123–214, 2011.