A Hierarchical Bayesian Model for the Identification of PET Markers Associated to the Prediction of Surgical Outcome after Anterior Temporal Lobe Resection

Sharon Chiang[1,2], Michele Guindani[3], Hsiang J. Yeh[4], Sandra Dewar[4], Zulfi Haneef[5], John M. Stern[4] and Marina Vannucci[1]*

[1] Department of Statistics, Rice University, Houston, TX, United States, [2] School of Medicine, Baylor College of Medicine, Houston, TX, United States, [3] Department of Statistics, University of California, Irvine, Irvine, CA, United States, [4] Department of Neurology, University of California, Los Angeles, Los Angeles, CA, United States, [5] Department of Neurology, Baylor College of Medicine, Houston, TX, United States

We develop an integrative Bayesian predictive modeling framework that identifies individual pathological brain states based on the selection of fluoro-deoxyglucose positron emission tomography (PET) imaging biomarkers and evaluates the association of those states with a clinical outcome. We consider data from a study on temporal lobe epilepsy (TLE) patients who subsequently underwent anterior temporal lobe resection. Our modeling framework looks at the observed profiles of regional glucose metabolism in PET as the phenotypic manifestation of a latent individual pathologic state, which is assumed to vary across the population. The modeling strategy we adopt allows the identification of patient subgroups characterized by latent pathologies differentially associated to the clinical outcome of interest. It also identifies imaging biomarkers characterizing the pathological states of the subjects. In the data application, we identify a subgroup of TLE patients at high risk for post-surgical seizure recurrence after anterior temporal lobe resection, together with a set of discriminatory brain regions that can be used to distinguish the latent subgroups. We show that the proposed method achieves high cross-validated accuracy in predicting post-surgical seizure recurrence.

Keywords: Bayesian hierarchical model, positron emission tomography (PET), spatially-informed prior, mixture model, variable selection, Pólya-Gamma distribution

1. INTRODUCTION

In the era of precision medicine, in order to deliver targeted therapies for neurological disorders, the development of methods to identify reliable and quantifiable biomarkers that are associated to individual clinical outcomes has become of paramount importance (Insel and Cuthbert, 2015). Temporal lobe epilepsy (TLE) is the most common form of adult epilepsy and the most common epilepsy refractory to anti-epileptic drugs. Surgery provides an effective treatment for many patients, yielding a seven-fold greater probability of seizure freedom 1 year after surgery than patients treated with medications alone (Wiebe et al., 2001). Despite its effectiveness, 30–50% of patients with TLE continue to experience seizures after surgery (Spencer et al., 2005; de Tisi et al., 2011).

As interictal ^{18}F-fluorodeoxyglucose positron emission tomography (FDG-PET) has traditionally been used for seizure focus localization (Wieser, 2004), there is substantial interest in identifying methods that utilize PET for prediction of post-surgical seizure relief (Willmann et al., 2007). Mesial TLE with hippocampal sclerosis is defined by the presence of neuronal cell loss and gliosis in the CA1 region and endfolium of the hippocampus, a particular part of the temporal lobe (Wieser, 2004). Therefore, prediction of post-surgical outcome using FDG-PET has traditionally focused on specific regions selected a priori within the temporal lobe (Dupont et al., 2000; Lin et al., 2007). Such studies have demonstrated predictive value of FDG-PET for identifying mesial TLE. Increasing evidence, however, points at TLE as a network disorder that includes abnormality distributed beyond the temporal lobe, rather than a focal disorder (Bonilha et al., 2005; McDonald et al., 2008; Mueller et al., 2010; Chiang and Haneef, 2014). This suggests that whole-brain statistical approaches may allow for improved identification of quantifiable features from neuroimaging data that can be reliably associated with individual clinical outcomes and improve clinical decision-making.

Traditional predictive modeling approaches for neuroimaging data have included the use of pattern recognition techniques, such as Linear Discriminant Analysis (Haynes and Rees, 2005), Support Vector Machines (Mitchell et al., 2004; LaConte et al., 2005) and Bayesian classifiers (Burge et al., 2009; Arribas et al., 2010). In particular, pattern recognition techniques have been used with varying success to predict post-surgical outcome in TLE, ranging from 50 to 75% accuracy using random forests (Njiwa et al., 2015) to 70% accuracy using elastic net and support vector machines (Munsell et al., 2015). Recently, Bayesian spatial hierachical models have also been used to improve prediction accuracy from PET data by borrowing strength from spatial correlations between neighboring voxels/regions (Derado et al., 2013). Several approaches for dynamic PET data have also been proposed. O'Sullivan (2006) and Jiang and Ogden (2008), for example, utilize mixture modeling and conditional autoregressive models to incorporate spatial information into PET analysis, while other work has used functional principal components (Jiang et al., 2009) or wavelets (Millet et al., 2000; Alpert et al., 2006) to analyze dynamic PET signal. Although each of these approaches represents an important advance in neuroimaging methods development, these methods do not quantify the relative importance of selected regions, which may impact the effectiveness of related clinical decisions. Recently, Bayesian scalar-on-image regression methods have been proposed that associate a univariate outcome to massive multi-dimensional image predictors, particularly for functional magnetic resonance imaging (fMRI) data (van Gerven et al., 2010; Goldsmith et al., 2014; Li et al., 2015). All the methods above, however, do not consider the heterogeneity of the population of individuals and implicitly assume that, given a set of discriminatory regions, their association to the outcome is the same across the population. In reality, however, the strength of the association can vary across subgroups of subjects.

In this paper, we develop a statistical model to identify whole-brain biomarkers from PET imaging which are associated to the prediction of post-surgical seizure recurrence following anterior temporal lobe resection. Post-surgical seizure recurrence is thought to be due to incomplete resection of the epileptogenic zone, which is defined as the area of cortex necessary and sufficient for initiating seizures, and whose removal is necessary for seizure abolition (Lüders et al., 1993). While the epileptogenic zone was historically thought to arise from discrete focal sources, more recent evidence suggests that seizure activity arises from the activity of epileptogenic cortical networks that are distributed beyond the temporal lobe (Franaszczuk et al., 1994; Franaszczuk and Bergey, 1998; Baccalá et al., 2004; Worrell et al., 2004, 2008; Jirsch et al., 2006; Kramer et al., 2008; Chiang et al., 2017a). Patients with different epileptogenic zone configurations are expected to exhibit different likelihoods of post-surgical seizure recurrence. Different epileptogenic zone configurations are also expected to produce different interictal metabolic patterns of FDG uptake, due to the effect of epileptogenic activity on neuronal loss and postictal metabolic depression (Luders, 2008). The epileptogenic zone, however, cannot be identified pre-operatively, due to the fact that parts of an epileptogenic lesion may not be implicated in the preoperatively recorded seizure, but will continue to generate seizures post-operatively if not resected (Rosenow and Lüders, 2001). In our model formulation, we look at the observed PET brain measurements as the phenotypic manifestation of latent individual pathological states that are assumed to vary across the population. We then factor the joint distribution of the data into the product of two conditionally independent submodels, an *outcome model* that relates the post-surgical outcome to the latent states, and a *measurement model* that relates those latent states to the observed brain measurements. For the latter, we employ mixture models for clustering and variable selection priors that capture spatial correlation among neighboring brain regions. This allows us to cluster subjects into subgroups with different latent pathological states, while simultaneously identifying discriminatory brain regions that characterize the subgroups. A logistic regression model relates the latent states to the binary clinical outcome.

We apply the proposed approach to PET data collected at the University of California, Los Angeles (UCLA) as part of a clinical study on post-surgical outcomes in temporal lobe epilepsy. We also incorporate into the analysis connectivity information from resting-state functional magnetic resonance imaging (fMRI) data, to inform the selection of discriminatory brain regions. Integrative models that take into account neuroscientific information from multi-modal data sources, such as fMRI, electroencephalography (EEG), or diffusion tensor imaging (DTI), are a pressing issue in the field, in particular given the limited number of patient samples collected in many neuroimaging experiments (Bowman et al., 2012; Hinne et al., 2014; Jorge et al., 2014). Bayesian inference provides a powerful way to incorporate multi-modal imaging into computational anatomy by inclusion through network priors. In our case study, we identify a subgroup of patients at high risk for post-surgical seizure recurrence, together with several discriminatory brain regions which can be used in clinical decisions to maximize interventional treatments. Furthermore, we show that the proposed approach achieves high cross-validated accuracy in

predicting post-surgical seizure recurrence. Further assessment of the performance of our method is performed in the Supplementary Material by conducting a comparison study on synthetic data against multi-step approaches and/or approaches that do not condition on latent states.

2. MATERIALS AND METHODS

2.1. Case Study on Temporal Lobe Epilepsy

Positron emission tomography (PET) is a type of *in vivo* nuclear medicine imaging which uses radioactive tracers to quantify tissue function. The subject is injected with a positron-emitting isotope, such as [18]F-FDG, and a PET image is reconstructed of the isotope concentration based on the incidence of gamma rays from the positron-electron annihilation. In this work, we analyze data on 19 adult patients with drug resistant MTLE and radiological evidence of unilateral hippocampal sclerosis (MTLE-HS), who underwent pre-operative interictal [18]F-FDG PET and anterior temporal lobe resection (ATL) at the UCLA Seizure Disorder Center between 2007 and 2012. Patients were identified from the UCLA video-EEG Epilepsy Monitoring Unit. As the primary outcome of this study was post-operative seizure freedom after epilepsy surgery, a healthy control group was not obtained as anterior temporal lobe resections are not performed in healthy patients without indication for surgery. Diagnostic evaluation included video-EEG monitoring, high resolution MRI, interictal [18]F-FDG PET, and neuropsychological testing. PET/CT scans were acquired on a Siemens Biograph scanner as described in Kerr et al. (2013). Patients fasted for at least 6 h before each scan except for water and medications. Patients received 0.14 mCi/kg of [18]F-FDG intravenously and rested in a quiet, dimly lit room with their eyes open during the ensuing 40 min uptake period with concomitant EEG monitoring to confirm interictal status. The iterative reconstruction program Ordered Subset Expectation Maximization (OSEM) available through NeuroQ (Syntermed, GA, USA) was used for reconstruction of PET images. Iterative reconstruction was halted after two iterations using eight subsets. CT images were reconstructed using filtered back projection at 3.4 mm axial intervals to match the slice separation of the PET data, and used for attenuation correction. Post-operative seizure freedom was assessed 1 year after surgery and classified as either seizure-free (SF; Engel Class 1) or not seizure-free (NSF; Engel Class 2–4). The binary outcome of complete freedom from disabling seizures (Engel Class 1) is the standard primary outcome of interest evaluated in epilepsy surgery treatment trials (Engel et al., 2012). The use of this primary outcome in epilepsy surgery trials results from the goal of epilepsy surgery, which is complete seizure freedom. In addition, we have available resting state fMRI (rs-fMRI) data collected on a separate set of 32 TLE patients recruited from the UCLA Seizure Disorder Center. Details on fMRI data are described in section 3.1.

2.2. PET Pre-processing

In PET studies, the quantity that is clinically assessed is a scalar rate of regional glucose uptake, based on a method described by Sokoloff et al. (1977). This quantity is then normalized relative to an internal reference standard, such as the whole-brain or cerebellar activity, and compared to the expected level for a reference normal subject (Silverman et al., 2008). The cerebellum is commonly used as the reference PET region for diseases of interest in which the cerebellum is thought not to be affected, such as diseases involving diffuse forebrain involvement. However, cerebellar atrophy is a very well described phenomenon in epilepsy, and is moreover associated with longer duration of epilepsy as well as younger age of epilepsy onset (Sandok et al., 2000). Given that the cerebellum could be more involved in epilepsy than traditionally thought (Fountas et al., 2010), we chose to normalize by the average whole-brain uptake rather than by the cerebellum. The assessed quantity therefore provides a measure of the level of metabolic activity in each region, relative to that expected in healthy controls. Uptake levels may be quantified on the single-voxel level or based on the mean uptake within fixed regions of interest. However, because single-voxel measurements are adversely affected by noise, the use of regions of interest (ROIs) in FDG-PET has been suggested as a more robust alternative for clinical practice (Wahl et al., 2009), which additionally facilitates standardized comparisons of affected regions across subjects. NeuroQ (Syntermed, GA, USA) is a software approved by the FDA in 2004 for quantitative assessment of brain PET imaging in clinical practice and was used to pre-process PET images. Following transformation into template Montreal Neurological Institute (MNI) space by a method previously described by Tai et al. (1997), images were segmented into 47 predefined regions of interest using a predefined NeuroQ atlas (Silverman and Melega, 2004; Ercoli et al., 2012) which has been previously considered for quantitative assessment of PET data in clinical practice (Smith et al., 2007; McCallum et al., 2010; Torosyan and Silverman, 2012; Kerr et al., 2013; Akdemir et al., 2014). ROI abbreviations are listed in the Supplementary Material. Pre-processing consisted of scalp removal, rigid registration to a reference PET image to correct for head tilt, and reformatting of transaxial slices to fit normal template transaxial slices using 10 iterations. Maximization of the mutual information between the image volumes was used to identify the registration parameter. A mean count was calculated in each ROI, normalized by the whole-brain counts and standardized relative to the mean and standard deviation of each ROI among healthy controls. Greater magnitude of PET image intensities indicate more pathological levels of metabolic activity, with positive values indicating greater levels of hypermetabolism (i.e., greater metabolism than in healthy controls) and negative values indicating greater levels of hypometabolism (i.e., lower metabolism than in healthy controls). Consequently, different patterns of nonzero signal characterize different pathological patterns of metabolic activity. Imaging patterns of hyper- and hypometabolism were of interest in this study rather than the raw PET signal intensities, due to the association of hypermetabolic activity with epileptic activity. Lateralized ROIs were recoded from left and right to ipsilateral or contralateral with respect to the side of subsequent resection. A histogram of the normalized and standardized PET image intensities (Figure not shown) indicated a bell-shaped, unimodal, and fairly symmetrical distribution, with a skewness of −0.39.

2.3. Statistical Model

Let $\boldsymbol{X}_i$ denote the $R \times 1$ vector of normalized PET image predictors on R brain regions of interest (ROIs) for subject i and let Y_i denote the corresponding post-surgical outcome, for $i = 1, \ldots, n$. We propose to study the association between the PET image predictors and the outcome via a *measurement error* model formulation. As described above, non-zero values of $\boldsymbol{X}$ indicate the level of PET metabolic activity, with different non-zero intensity patterns indicating different pathological imaging profiles. Accordingly, we assume that the brain's observed profile of metabolic activity is the manifestation of a latent (i.e., unobserved) pathological state. In epilepsy, the latent pathological state represents the configuration of metabolic activity in regions implicated in the underlying epileptogenic zone, which is in turn associated to post-surgical seizure recurrence. Here, we assume a finite number of pathological states due to the expected modular organization of the brain, which is generally decomposed into a finite number of submodules (Meunier et al., 2010). Let η_i denote the latent pathological state of subject i. Then, we propose to factor the joint distribution of $\boldsymbol{Z}_i = \{Y_i, \boldsymbol{X}_i\}_{i=1}^n$ into the product of two conditionally independent sub-models: an *outcome model* that relates the clinical outcome to the latent pathological state, and a *measurement model* that relates the latent pathological state to the observed imaging data. Therefore, we consider a non-differential measurement error model, i.e., conditionally upon the latent pathological state η_i, the observed surrogate $\boldsymbol{X}_i$ contains no additional information on the outcome Y_i (Richardson and Gilks, 1993), $f(Y_i|\eta_i, \boldsymbol{X}_i) = f(Y_i|\eta_i)$. This model allows us to capture the current understanding in epilepsy that failure of temporal lobe resection results most likely from incomplete resection of the epileptogenic zone (Ryvlin and Kahane, 2005). In other words, if the true epileptogenic zone were known, data contained in the PET image $\boldsymbol{X}_i$ would not provide any additional information on the probability of post-operative seizure recurrence Y_i. Thus,

$$f(\boldsymbol{Z}|\boldsymbol{\eta}) = \prod_{i=1}^n f(Y_i|\eta_i) f(\boldsymbol{X}_i|\eta_i), \qquad (1)$$

where $\boldsymbol{\eta} = (\eta_1, \ldots, \eta_n)$. We specify the measurement model in Equation (1) as a mixture model with variable selection. Subgroups of patients with different epileptogenic zone configurations may be expected to exhibit different risks of post-surgical seizure recurrence. We therefore specify the outcome model in Equation (1) as a logistic regression model that relates the latent states to the binary clinical outcome. There is extensive literature on the use of measurement error models to model data in which risk factors related to the observed disease or treatment status are unknown, but where surrogate measures, which provide information on the unobserved risk factor, are recorded. A review of measurement error models may be found in Carroll et al. (2006). With respect to existing literature, our model formulation allows us to cluster subjects into subgroups with different latent pathological states, i.e., different epileptogenic zone configurations, while simultaneously identifying discriminatory brain regions. In the selection, we also

capture spatial correlation among neighboring brain regions via a spatial prior, as described in section 2.3.3.

2.3.1. Clustering via Finite Mixture Models

We envision that a subject may be characterized by one of K possible pathological states. Let η_i denote a latent random variable that identifies the state of the i-th subject, $i = 1, \ldots, n$. We assume that the latent individual state η_i takes values in $\{1, \ldots, K\}$, where one of the states can be assumed as reference. Then, for each subject i we define an allocation vector $\boldsymbol{\rho}_i = (I(\eta_i = 1), \ldots, I(\eta_i = K - 1))$, where $I(\eta_i = k)$ indicates that subject i has latent state k, i.e., $I(\eta_i = k) = 1$ if $\eta_i = k$, and 0 otherwise. Then, for the measurement model in Equation (1), we choose a finite mixture model that clusters the n subjects into K possible subgroups as

$$f(\boldsymbol{X}_i|\eta_i, \boldsymbol{\pi}, \boldsymbol{\theta}) = \sum_{k=1}^K \pi_k f(\boldsymbol{X}_i|\boldsymbol{\theta}_k),$$

with $\eta_i = k$ if subject i belongs to cluster k and $P[\eta_i = k] = \pi_k$. The η_i's are assumed to be independent and identically distributed, so that $\boldsymbol{\eta} \sim$ Multinomial $(1; \pi_1, \ldots, \pi_K)$. We assume a Dirichlet prior on the mixture weights, $p(\boldsymbol{\pi}) =$ Dirichlet $(\alpha_1, \ldots, \alpha_K)$. We consider the case where $f(\mathbf{x}_i|\boldsymbol{\theta}_k)$ is Gaussian with parameters $\boldsymbol{\theta}_k = (\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$, so that

$$f(\boldsymbol{X}_i|\boldsymbol{\theta}_k) = \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \qquad (2)$$

with $k = 1, .., K$. The component-specific mean $\boldsymbol{\mu}_k$ models the latent state specific random effect and characterizes the mean metabolic profile for subjects with latent state k, whereas $\boldsymbol{\Sigma}_k$ is a variance-covariance matrix that captures general relationships among regions for subjects with latent state k. In summary, the likelihood function for the measurement model is

$$L(\boldsymbol{X}|\boldsymbol{\eta}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = \prod_{k=1}^K (2\pi)^{-n_k R/2} |\Sigma_k|^{-n_k/2}$$

$$\times \exp \left\{ -\frac{1}{2} \sum_{\{i \,:\, \eta_i=k\}} (\boldsymbol{X}_i - \boldsymbol{\mu}_k)^T \Sigma_k^{-1} (\boldsymbol{X}_i - \boldsymbol{\mu}_k) \right\},$$

with n_k denoting the number of subjects in cluster k. Here we assume diagonal variance-covariance matrices $\boldsymbol{\Sigma}_k = \text{diag}(\sigma_{k,1}, \ldots, \sigma_{k,R})$. Even though we make this simplifying assumption at this stage of the hierarchy, our proposed model is still able to capture structural dependencies via the specification of the prior model for the mean components in Equation (4) that we describe in section 2.3.3.

2.3.2. Association with the Treatment Outcome

The outcome model in Equation (1) allows the prediction of the subject-specific outcomes based on the patients' individual latent pathological state η_i. We can relate the latent states with the outcome of interest by employing a generalized linear model. In general, we may have available a vector of baseline covariates $\mathbf{U}_i$

for subject i. Since the post-surgical outcome is binary, we can then use a logistic regression model

$$p(Y_i = y_i | \eta_i, \boldsymbol{\beta}) = \frac{\exp(\boldsymbol{\xi}_i^T \boldsymbol{\beta})^{y_i}}{1 + \exp(\boldsymbol{\xi}_i^T \boldsymbol{\beta})}, \qquad (3)$$

with $\boldsymbol{\beta} = (\beta_0, \dots, \beta_{K-1}, \boldsymbol{\beta}_U)$ and $\boldsymbol{\xi}_i = (1, \boldsymbol{\rho}_i, \mathbf{U}_i)$, where $\boldsymbol{\beta}_U$ is the vector of corresponding regression coefficients for $\mathbf{U} = \{\mathbf{U}_i\}_{i=1}^n$. Here, β_k, $k = 1, \dots, K - 1$ captures the "risk" associated to latent state k relative to the baseline latent state. Each β_k can be interpreted as the log-odds of the outcome for subjects in state k relative to subjects in the reference state, and β_0 as an intercept term yielding the log-odds of the outcome for subjects in the reference state.

The analytically intractable form of the likelihood function using a logit link is known to pose challenges for Bayesian inference in logistic regression models. To address this and to improve posterior sampling, we employ the data augmentation approach recently devised by Polson et al. (2013). Let ω be a Pólya-Gamma random variable, $\omega \sim \text{PG}(b, c)$, with parameters $b > 0$ and $c \in \mathbb{R}$,

$$\omega \overset{D}{=} \frac{1}{2\pi^2} \sum_{k=1}^{\infty} \frac{g_k}{(k - 1/2)^2 + c^2/4\pi^2},$$

where g_k are independently distributed as $\text{Gamma}(b, 1)$. Augmentation with a Pólya-Gamma random variable allows for the likelihood contribution of the ith observation to be written as

$$
\begin{aligned}
L_i(\boldsymbol{\beta}) &= \frac{\exp(\boldsymbol{\xi}_i^T \boldsymbol{\beta})^{y_i}}{1 + \exp(\boldsymbol{\xi}_i^T \boldsymbol{\beta})} \\
&= \frac{1}{2} \exp(\kappa_i \boldsymbol{\xi}_i^T \boldsymbol{\beta}) \int_0^\infty \exp\left(-\frac{\omega_i (\boldsymbol{\xi}_i^T \boldsymbol{\beta})^2}{2} \right) p(\omega_i) \partial \omega_i,
\end{aligned}
$$

where $\kappa_i = y_i - 1/2$, for $\omega_i \sim \text{PG}(1, 0)$. Combining all n terms then gives the following convenient representation for the conditional likelihood in $\boldsymbol{\beta}$, given $\boldsymbol{\omega}$ and $\boldsymbol{\eta}$:

$$L(\boldsymbol{\beta} | \boldsymbol{\eta}, \boldsymbol{\omega}) \propto \exp\left\{ -\frac{1}{2} (z - \Xi \boldsymbol{\beta})^T \Omega (z - \Xi \boldsymbol{\beta}) \right\},$$

where $z = (\kappa_1/\omega_1, \dots, \kappa_m/\omega_m)$, $\kappa_i = y_i - 1/2$, $\boldsymbol{\Omega} = \text{diag}(\omega_1, \dots, \omega_n)$, Ξ is the $n \times K$ matrix $\Xi = (\boldsymbol{\xi}_1^T, \dots, \boldsymbol{\xi}_n^T)$, $\boldsymbol{\xi}_i = (1, \rho_{i,1}, \rho_{i,2}, \dots, \rho_{i,K-1})$, and $\rho_{i,k} = I(\eta_i = k) \, \forall k = 1, \dots, K - 1$. See Polson et al. (2013) for details. We complete the model by imposing a conjugate prior on $\boldsymbol{\beta}$, $p(\boldsymbol{\beta}) = \text{N}(\mathbf{m}_\beta, V_\beta)$, where $\mathbf{m}_\beta$ and V_β denote the prior mean and covariance, respectively.

2.3.3. Spatially-Informed Selection Prior
Not all brain regions are expected to provide information about the subgroup structure of the subjects, in which case the inclusion of non-discriminatory regions in model (Equation 2) may obscure the discovery of true groups. One way to address this issue is through variable selection for clustering. Let $\boldsymbol{\gamma} \in \{0, 1\}^R$ denote a binary vector, where $\gamma_j = 1$ if region j is discriminatory, and $\gamma_j = 0$ otherwise, $\forall j = 1, \dots, R$. We follow Hoff (2006) and

identify discriminatory brain regions by imposing spike-and-slab priors on the random effects $\boldsymbol{\mu}_k = (\mu_{k,1}, \dots, \mu_{k,R})$. Given the spatial contiguity in neuronal glucose consumption, we allow for spatial smoothness among neighboring regions by specifying the slab portion of the prior as an intrinsic conditional autoregressive (ICAR) prior distribution (Banerjee et al., 2014). Our prior on $\mu_{k,j}$ can be written as

$$
\begin{aligned}
p(\mu_{k,j} | \gamma_j, \boldsymbol{\mu}_{k,\backslash j}) = &\gamma_j \text{N}\left(\frac{\sum_{j'=1}^R S_{j,j'} \mu_{k,j'}}{\sum_{j'=1}^R S_{j,j'}}, \frac{c_k}{\sum_{j'=1}^R S_{j,j'}} \right) \\
&+ (1 - \gamma_j) \delta_0(\mu_{k,j}),
\end{aligned}
\qquad (4)
$$

where δ_0 denotes a spike at zero, S is an $R \times R$ symmetric neighborhood matrix, with $S_{j,j'} = 1$ if regions j and j' are neighbors, and $S_{j,j'} = 0$ otherwise, and where $\boldsymbol{\mu}_{k,\backslash j}$ denotes all elements of $\boldsymbol{\mu}_k$ except the jth element. We also impose priors on the diagonal elements of Σ_k in Equation (2) and allow for separate variances for the discriminatory and non-discriminatory regions. In particular, for the parameters corresponding to $\gamma_j = 1$, we have $\sigma_{k,j} = \sigma_k \sim \text{IG}(a_k, b_k)$ for all k, while for $\gamma_j = 0$ we impose $\sigma_{k,j} = \sigma_0 \sim \text{IG}(a_0, b_0)$. Finally, in specifying the prior on the selection indicators, $\boldsymbol{\gamma}$, we allow for external information on the network structure of the brain, for example on connectivity between regions, to be incorporated in the model by imposing an Ising prior of the type

$$p(\boldsymbol{\gamma}) \propto \exp\left\{ e \mathbf{1}_R^T \boldsymbol{\gamma} + f \boldsymbol{\gamma}^T S \boldsymbol{\gamma} \right\}, \qquad (5)$$

with S denoting the neighborhood matrix. If a connection exists between two regions j and j', then selection of one region j (i.e., $\gamma_j = 1$) leads to an increased probability that region j' will also be selected (i.e., $\gamma_{j'} = 1$). The hyperparameter $e \in (-\infty, \infty)$ controls the sparsity of the model and represents the prior expected number of discriminatory regions. The hyperparameter $f > 0$ is a smoothing parameter which represents the prior probability of a region being discriminatory given that its neighbors are too. In particular, if a region has no neighbors, then its prior distribution reduces to an independent Bernoulli distribution with probability $\exp(e)/(1 + \exp(e))$, which is a common prior assumed in Bayesian variable selection literature in the case of independent variables.

The prior construction (Equations 4, 5) allows for sparsity while promoting spatial contiguity in the selection. The ICAR prior, in particular, ensures that each cluster's mean metabolic PET profile varies smoothly in space, as each $\mu_{k,j}$ is modeled to vary around the mean of its neighbors, with variance inversely scaled by the number of neighbors. Spatial prior constructions have been used extensively in neuroimaging applications, particularly with fMRI data (Smith and Fahrmeir, 2007; Zhang et al., 2014; Li et al., 2015).

2.3.4. MCMC Algorithm
In order to sample from the joint posterior distribution of all parameters $(\{\sigma_k\}_{k=1}^K, \sigma_0, \boldsymbol{\eta}, \boldsymbol{\pi}, \boldsymbol{\gamma}, \{\boldsymbol{\mu}_k\}_{k=1}^K, \boldsymbol{\beta}, \boldsymbol{\omega})$, we employ Markov Chain Monte Carlo (MCMC) methods that combine

variable selection stochastic search algorithms that use *add-delete-swap* moves (Savitsky et al., 2011) with efficient Pólya-Gamma sampling for logit models (Polson et al., 2013). We provide full details of the implementation in the Supplementary Material.

2.3.5. Prediction

An important characteristic of our model formulation is that it allows for prediction of the outcome status y_f of a future observation x_f, based on the training data $\{X, Y\}$. In the context of pre-surgical evaluation for epilepsy surgery, this allows for probabilistic, patient-specific predictive estimates of the patient's probability of surgery benefit, in order to assist with clinical decision-making. The predictive distribution is given by

$$p(y_f|x_f, X, Y) = \int_{\beta} \sum_{\eta_f \in \{1,\dots,K\}} p(y_f|\eta_f, \beta) p(\beta|X, Y) p(\eta_f|x_f) \partial\beta,$$

(6)

and cannot be computed in closed form. Following standard Bayesian techniques, these steps can be employed to simulate from Equation (6):

1. Sample T values of μ_k, Σ_k, π_k, β from the joint posterior, using the MCMC algorithm as described in the Supplementary Material.
2. For each posterior draw, $t = 1, \dots, T$:

   - Sample $m \geq 1$ values of $\eta_f \in \{1, \dots, K\}$ from $p(\eta_f|x_f)$, where $\forall k = 1, \dots, K$

     $$p(\eta_f = k|x_f) \propto p(x_f|\eta_f = k) p(\eta_f = k) = p(x_f|\mu_k^{(t)}, \Sigma_k^{(t)}) \pi_k^{(t)}.$$

   - For each sampled value of η_f, sample a value of $y_f \in \{0, 1\}$ from $p(y_f|\eta_f, \beta^{(t)})$.

The posterior predictive probability $p(y_f = 1|x_f, X, Y)$ can then be estimated as the proportion of posterior predictive samples for which $y_f = 1$. In the analyses of this paper, given the limited number of samples available, which does not allow a meaningful splitting of the data into training and validation, we implemented cross-validation prediction via the importance-sampling approach, as proposed by Gelfand (1996), and write the cross-validation predictive density for the ith observation as

$$p(Y_i = 1|X, Y_{-i}) = \int_{\eta, \beta} p(Y_i = 1|X, Y_{-i}, \eta, \beta) p(\eta, \beta|X, Y_{-i}) \partial\beta \partial\eta$$

where we use $p(\eta, \beta|X, Y)$ as an importance sampling density for $p(\eta, \beta|X, Y_{-i})$, and Y_{-i} denotes the non-hold out outcomes. Specific details on implementation are provided in the Supplementary Material.

3. RESULTS

We now apply the proposed model to the data we have available from the University of California, Los Angeles Seizure Disorder Center, where we illustrate the utility of our proposed model for predicting a post-surgical outcome among MTLE-HS patients from pre-surgical FDG-PET imaging.
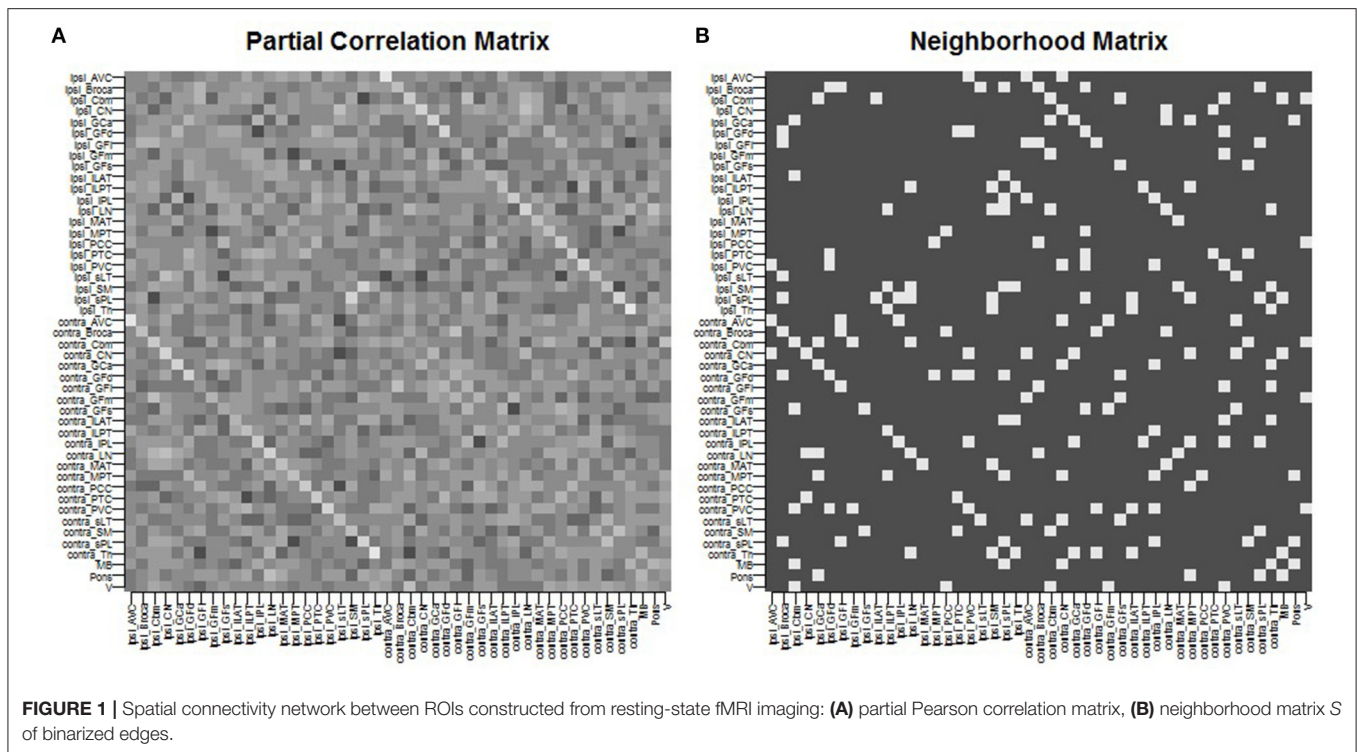
3.1. Prior Connectivity Network

For this analysis, we allowed the spatial network prior Equation (5) to capture information on functional connectivity between the ROIs, which we estimated based on resting-state fMRI data (rs-fMRI), collected on a separate set of 32 unilateral temporal lobe epilepsy patients from the UCLA Seizure Disorder Center. Rs-fMRI was performed on the subjects after a comprehensive epilepsy surgery evaluation and prior to epilepsy surgery. None of the patients had a seizure in the 24 h preceding the imaging or had seizures during the study, as confirmed by the simultaneous EEG obtained during fMRI. There were no post-surgical outcome data available for these patients. External or historical information is often used to formulate priors in Bayesian analysis. There is extensive literature which demonstrates the general replicability of Pearson correlation estimation of functional connectivity from rs-fMRI in temporal lobe epilepsy (Centeno and Carmichael, 2014). Furthermore, despite increasing evidence that functional connectivity is dynamic (Honey et al., 2009; Ma et al., 2014; Chiang et al., 2016), recent research indicates a large proportion of the information present in functional connectivity is contained in static estimates (Chiang et al., 2017b).

We give full details of the rs-fMRI data and the process to estimate a connectivity network in the Supplementary Material. In brief, preprocessing of rs-fMRI imaging was performed using FSL (fMRIB Software Library) version 5.0.7 (Oxford, United Kingdom, www.fmrib.ox.ac.uk/fsl). Functional connectivity between the 47 ROIs was estimated by placing a 6-mm spherical seed in Montreal Neurological Institute (MNI) space at the location of each of the 47 ROIs. Each patient's fMRI BOLD image was registered to the patient's high-resolution structural image using FLIRT (FMRIB's Linear Image Registration Tool) (Jenkinson et al., 2002; Greve and Fischl, 2009), and the high-resolution structural was registered to the standard MNI space using FNIRT (FMRIB's Non-linear Image Registration Tool) (Andersson et al., 2007). Functional connectivity between each pair of nodes was computed as the partial Pearson correlation between the averaged regional time-series. This provided us with a 47×47 correlation matrix. An edge was then considered as included in the connectivity network if the correlation between the regions exceeded a given threshold. The threshold was chosen so that the average number of neighbors for each region was approximately 5, yielding a connectivity structure close to a three-dimensional lattice. The resulting network was used as the neighborhood matrix S in the specification of the MRF prior (Equation 5) on γ and also in the ICAR prior (Equation 4) on the slab portion of the prior on $\mu_{k,j}$. The estimated functional connectivity matrix and resulting neighborhood matrix S are shown in **Figure 1**. We observe several known connectivity relationships, including functional connectivity between regions in the brainstem (midbrain, pons); between the primary and associative visual cortices; between the cerebellar hemispheres and vermis; and between ipsilateral and contralateral ROIs (Quigley et al., 2003).

3.2. Biomarker Selection and Clustering

In our approach to model fitting we consider a grid of values of K to find the number of states K yielding the best model fit

FIGURE 1 | Spatial connectivity network between ROIs constructed from resting-state fMRI imaging: **(A)** partial Pearson correlation matrix, **(B)** neighborhood matrix S of binarized edges.

that also provides improved clinical interpretability. For the study of this paper, model fit for each value of K for $K = 2, \ldots, 6$ was assessed using the deviance information criterion (DIC) of Spiegelhalter et al. (2002). We found that $K = 2$ clusters allowed for a parsimonious model permitting meaningful clinical characterization of high- and low-risk patients, with minimal to no further improvement in the DIC for larger values of K. This result was confirmed through model comparison using the posterior Bayes factor (Aitkin, 1991), with a posterior Bayes factor greater than 1 from comparisons of the $K = 2$ model to $K = 3, \ldots, 6$ models. Results we report here are based on the combined posterior output from two MCMC chains, with each chain initialized with different numbers of discriminatory ROIs and number of subjects in each subgroup. Other initial values were set as $\boldsymbol{\mu}_k^{(0)} = \mathbf{0}, \sigma_k^{(0)} = 1 \, \forall k, \sigma_0^{(0)} = 1, \boldsymbol{\beta}^{(0)} = \mathbf{0}$. We ran each MCMC chain for 100,000 iterations, with the first 50,000 sweeps discarded as burn-in.

As discussed in section 2.3.3, the hyperparameter e of the MRF prior (Equation 5) regulates the prior sparsity whereas f induces smoothness, with higher values of f yielding a higher prior probability that a region is selected given that its neighbors are selected. The choice of e and f has been discussed by Li and Zhang (2010) and Stingo et al. (2013). It is known that with distributions as in Equation (5) a phase transition boundary problem can be encountered, where the number of selected regions increases sharply for small changes in f (Li and Zhang, 2010). Here we set the sparsity parameter to $e = -4.5$, corresponding to a lower bound on the prior probability of selection of 1%. As for the prior smoothness, f, a plot of the prior over a grid of values $f \in \{0.1, 0.2, 0.3, \ldots, 0.9\}$ revealed that the phase transition

starts at a prior smoothness of $f = 0.2$ and becomes severe at around $f = 0.4$. As suggested by Li and Zhang (2010), the prior smoothness parameter f was therefore set to a value far from the phase transition boundary. Here we present results for two values, $f = 0.01$ and $f = 0.1$, representing different levels of small-to-moderate effect of the prior information on connectivity. As for the other hyperparameter settings, we placed a vague prior on the mixing parameters $\boldsymbol{\pi}$, that is, $\alpha_k = 1 \, \forall k$, and fixed the prior shape and scale parameters of the inverse gamma priors on σ_k and σ_0 to be non-informative with $a_k = 2$ and $b_k = 1 \, \forall k$, and $a_0 = 2$ and $b_0 = 1$. We also set the unscaled variance of the ICAR prior to $c_k = 5$, and the prior mean and covariance of $\boldsymbol{\beta}$ to $\boldsymbol{m}_\beta = \mathbf{0}$ and $V_\beta = 5\mathbb{I}$, respectively. Age of the patient at surgery, epilepsy duration, and history of generalized tonic clonic seizures were controlled for as baseline covariates in the logistic likelihood.

Convergence of each MCMC chain was assessed using two independent tests: the Raftery-Lewis diagnostic (Raftery and Lewis, 1992) and the Geweke test (Geweke, 1991). In addition, convergence of the multiple chains was assessed using the Gelman-Rubin potential scale reduction factor, based on the implementation in the R package "coda" (Raftery and Lewis, 1992). Convergence diagnostics indicated convergence to the stationary distribution (results reported in the Supplementary Material). Agreement between MCMC chains was assessed through the Pearson correlation between the marginal posterior probabilities of ROI selection and cluster allocation of each pair of chains.

For posterior inference, our primary interest is in the estimation of the discriminatory regions, the latent states, and
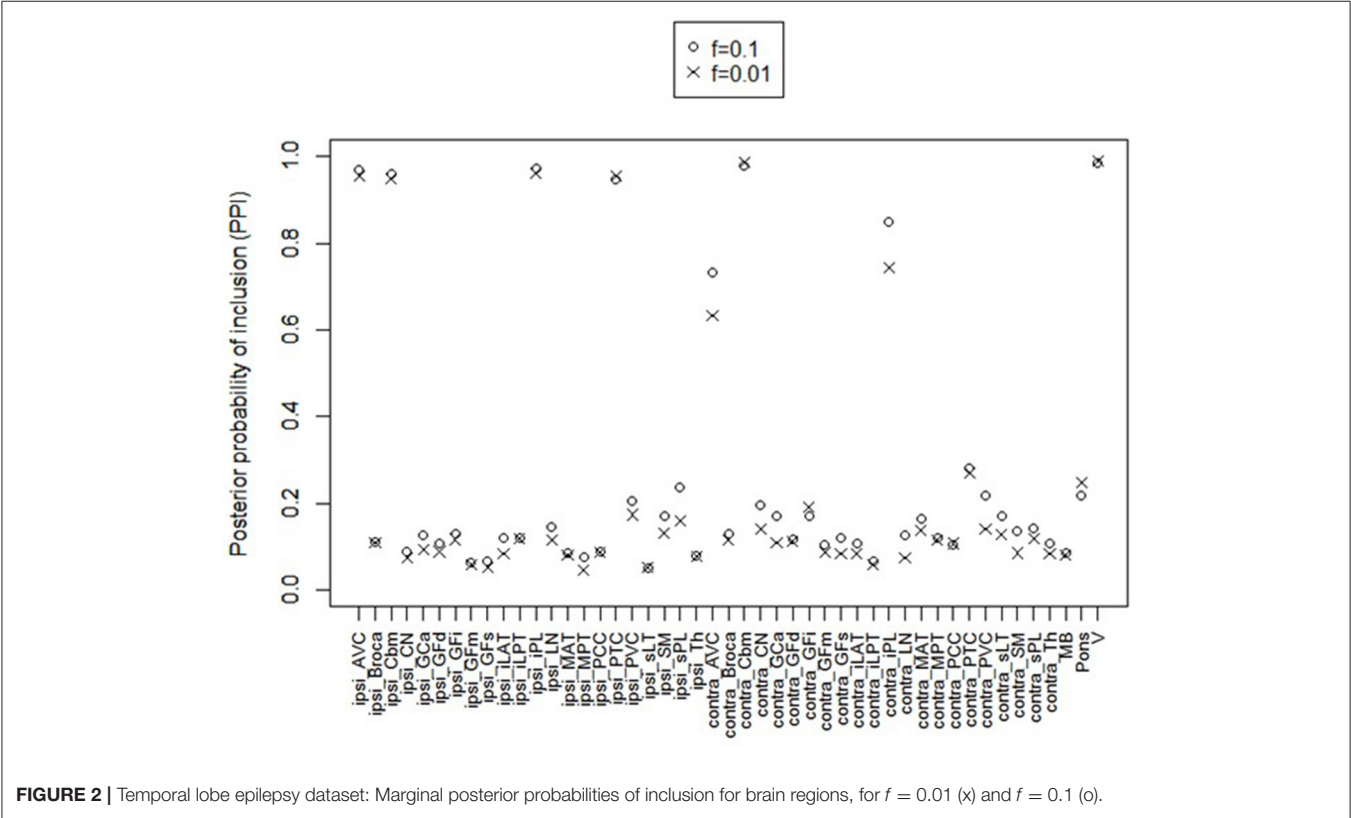
their association with the binary clinical outcome, as captured by the parameters γ, η, and β, respectively. Trace plots for these parameters showed good mixing for all chains (figures not shown). **Figure 2** shows the marginal posterior probabilities of inclusion (PPIs) for each of the 47 brain regions, with different graphical symbols for the settings of $f = 0.01$ (x) and $f = 0.1$ (o). Based on this plot, a selection of the discriminatory regions can be done by thresholding the PPIs. For example, the median model (Barbieri and Berger, 2004) selects the same subset of 8 ROIs under both $f = 0.01$ and $f = 0.1$. The selected brain regions are listed in **Table 1**, and graphically depicted in **Figure 3**. To examine the sensitivity of the selected regions to the formulation of the network prior, we additionally ran the model under a neighborhood matrix S defined by simple Euclidean distance. Selected discriminatory regions were robust to the formulation of the network, with the exception of the contralateral associative visual cortex, which had a marginal PPI of 0.303 ($f = 0.1$) and 0.311 ($f = 0.01$) under a network defined by spatial neighbors. This decrease in posterior probability is an effect of the MRF prior, due to the functional connectivity present between the ipsilateral and contralateral associative visual cortex in **Figure 1B** which is not captured based on spatial distance.
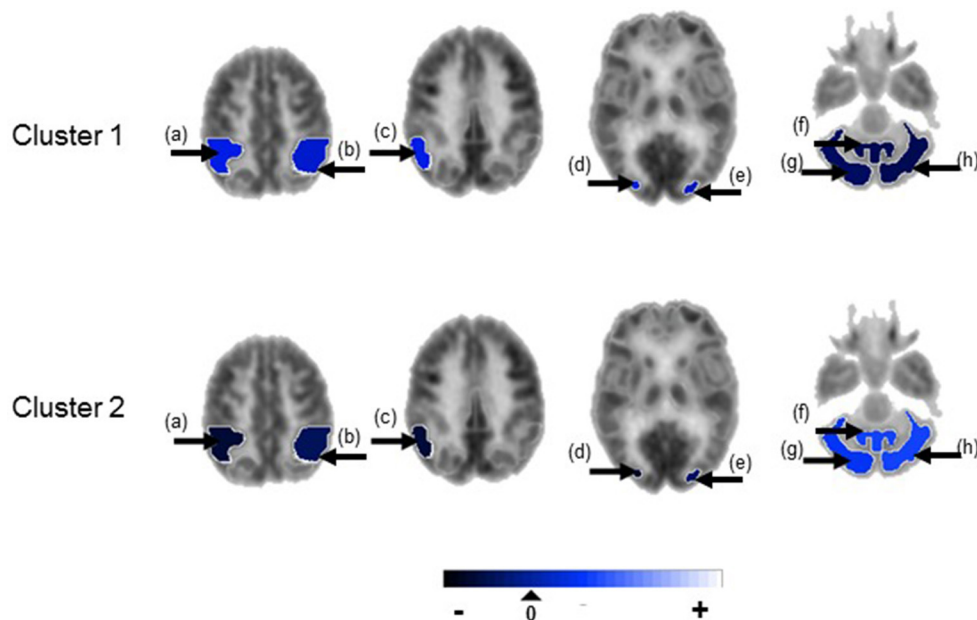
   Figure 4 shows the marginal posterior probabilities of sample allocations for each of the 19 MTLE-HS patients. A classification of the subjects into two subgroups can be obtained, for example, by assigning subjects according to the posterior mode of η. For interpretation of the two subgroups, one can examine the PET metabolic activities characterizing the subjects. These are

shown in **Figure 5** for the selected brain regions. Furthermore, posterior inference for the β parameters is summarized in **Table 2**. These results suggest that the two subgroups identify patients at different levels of risk for post-operative seizure recurrence, with one subgroup having a $e^{\beta} = 5.2$ times greater odds of persistent post-operative seizures 1 year after surgery (**Table 2**). This corresponds to a 90% posterior probability of an odds ratio >1 for post-surgical seizure freedom between the two identified subgroups (**Table 2**). **Figure 5** reveals, in particular, that the subgroup with greater odds of post-operative seizure recurrence (Cluster 2) is characterized by lower levels of interictal glucose metabolism in the bilateral associative
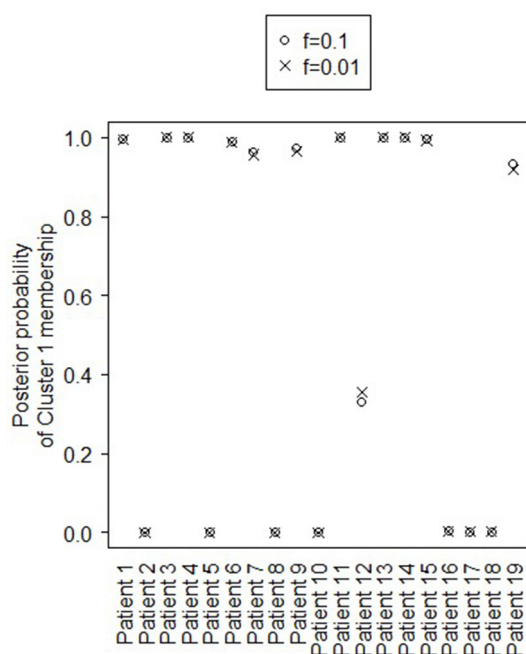
TABLE 1 | Temporal lobe epilepsy dataset: Selected brain regions and corresponding marginal posterior probabilities of inclusion (PPI).

ROI	PPI	
	$f = 0.01$	$f = 0.1$
Ipsilateral inferior parietal lobule	0.961	0.973
Ipsilateral parietotemporal cortex	0.955	0.948
Ipsilateral associative visual cortex	0.956	0.969
Contralateral inferior parietal lobule	0.742	0.850
Contralateral associative visual cortex	0.632	0.732
Contralateral cerebellar hemisphere	0.988	0.979
Ipsilateral cerebellar hemisphere	0.950	0.961
Cerebellar vermis	0.989	0.984

FIGURE 2 | Temporal lobe epilepsy dataset: Marginal posterior probabilities of inclusion for brain regions, for $f = 0.01$ (x) and $f = 0.1$ (o).

FIGURE 3 | Temporal lobe epilepsy dataset: Mean random effect (μ_k) of PET metabolic activity for ROIs with PPI greater than 0.5, shown on axial sections. (a) Ipsilateral inferior parietal lobule, (b) contralateral inferior parietal lobule, (c) ipsilateral parieto-temporal cortex, (d) ipsilateral associative visual cortex, (e) contralateral associative visual cortex, (f) cerebellar vermis, (g) ipsilateral cerebellar hemisphere, (h) contralateral cerebellar hemisphere. Non-selected ROIs are shown in grayscale.

FIGURE 4 | Temporal lobe epilepsy dataset: Marginal posterior probabilities of cluster allocation, for $f = 0.01$ (x) and $f = 0.1$ (o).

visual cortices, ipsilateral parieto-temporal cortex, and bilateral inferior parietal cortices, as well as higher levels of interictal glucose metabolism in the bilateral cerebellar hemispheres and cerebellar vermis. Our identification of these metabolic patterns may suggest extratemporal gliosis, as well as increased baseline levels of cortical excitability, in patients at higher risk for post-operative seizure recurrence. We provide further comment on the neurological significance of these findings in the Discussion.

3.3. Prediction Results

In addition to the identification of subgroups of subjects, characterized by latent pathologic conditions differentially associated to the outcome of interest, and the selection of imaging biomarkers that characterize the pathologic states of the subjects, our modeling approach allows a probabilistic estimate of an individual patient's risk of post-operative seizure recurrence. Probabilistic assessment of outcome risk may aid pre-surgical decision-making, by facilitating identification of patients with greater probability of seizure recurrence following anterior temporal lobe resection. Such information may potentially be weighed against the known risks of surgery (e.g., infection, bleeding, reactions to general anesthesia) to stratify patients according to predicted outcome. Here, we assessed prediction performance via importance-sampling cross-validation.

Figure 6 shows the receiver operating characteristic (ROC) curve, a plot of the false positive rates vs. the true positive rates, obtained for a grid of threshold values (0:0.05:1) on the estimated posterior predictive probabilities. The area under the curve (AUC) was 0.91. The optimal threshold, selected to maximize the Youden's index (Hiden and Glasziou, 1996), for imbalanced class sizes, resulted in an 84% predictive accuracy,
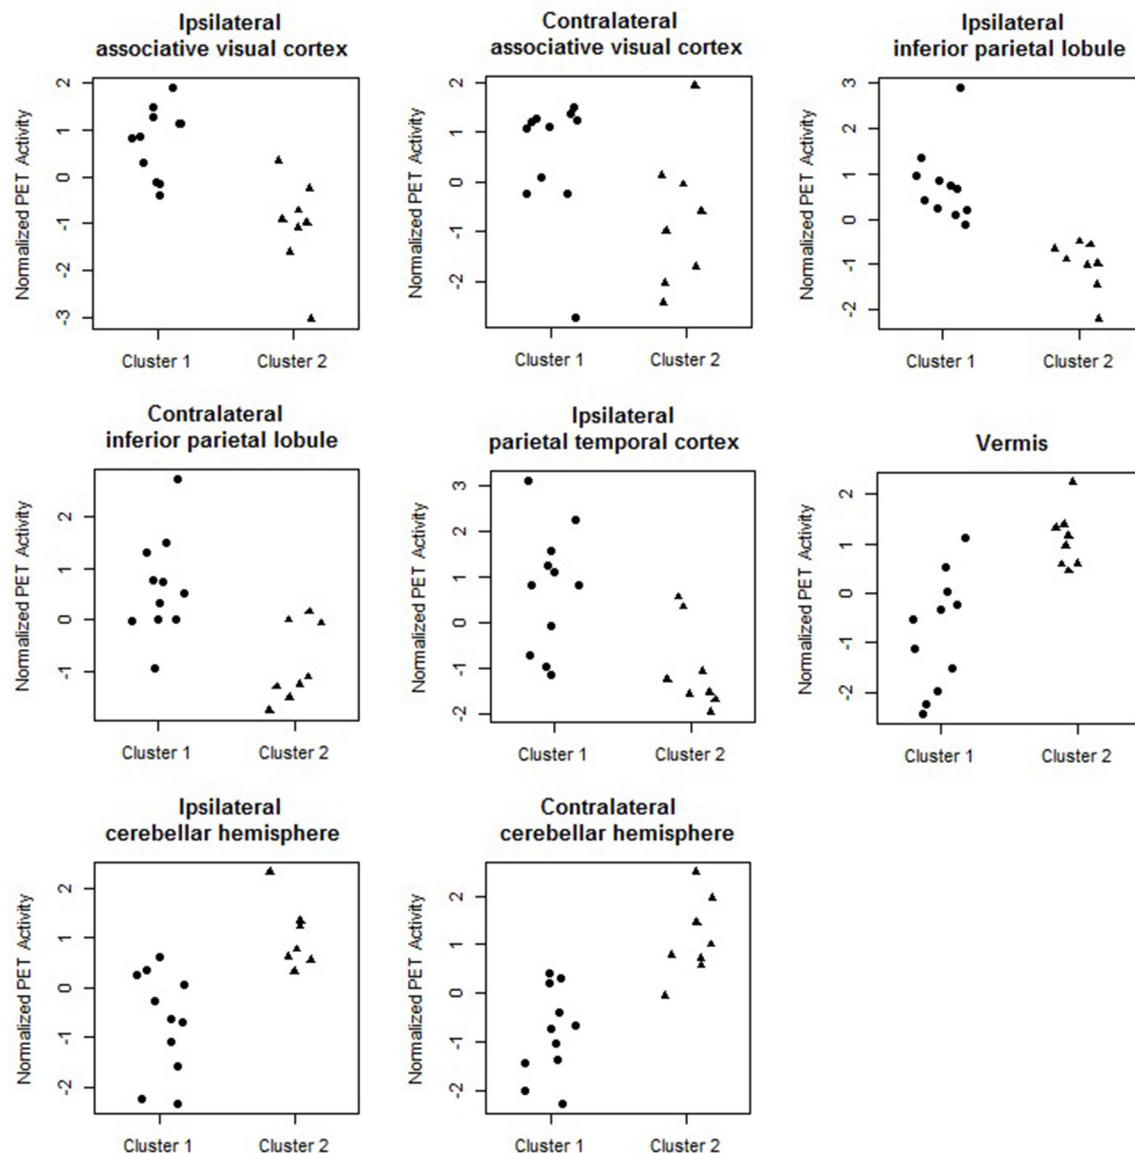
FIGURE 5 | Temporal lobe epilepsy dataset: Distribution of PET metabolic activity in the selected regions for the identified subgroups, for $f = 0.1$.

with correct prediction of post-surgical outcome in 16/19 patients, including 10/12 seizure-free patients and 6/7 non seizure free patients.
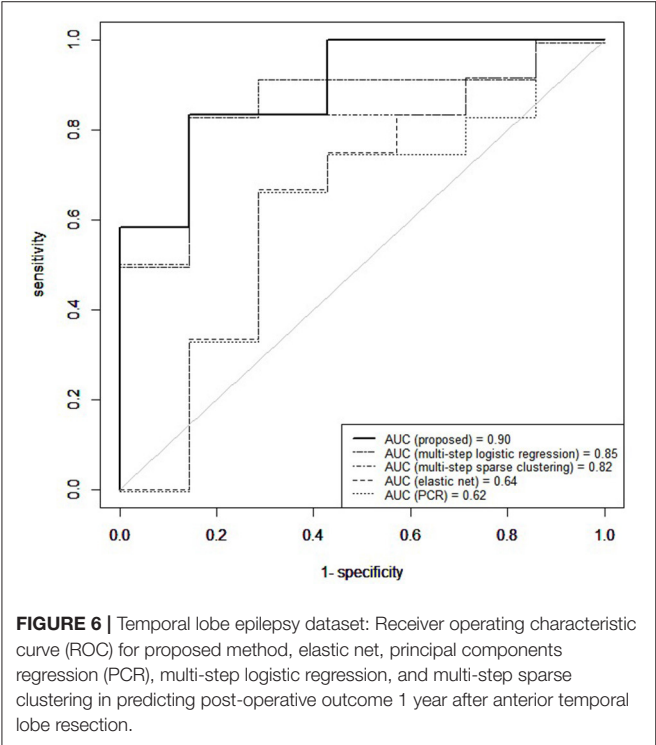
Our prediction results compared favorably to those we obtained on the same data with other analogous methods which predict binary outcomes from an identified underlying latent state. In particular, we compare to three multi-step approaches commonly used in prediction for their simplicity and computational speed. In the first approach, principal components was used to reduce the data to the top eight principal components, collectively explaining 85% of the variance in the data. The reduced principal components of $\mathbf{X}$ were then used as predictors within Bayesian logistic regression.

Predictive accuracy was assessed through the importance-sampling cross-validation prediction approach of Gelfand (1996). In the second approach, a multistep logistic regression approach was used, similarly to what has been done in neuroimaging studies (Versace et al., 2014). In this approach, a filtering approach was performed by calculating permutation p-values for each region and retaining regions with small p-values. Using this reduced subset of regions, patients were clustered using k-means. Bayesian logistic regression was fitted to predict post-surgical outcome from latent class membership, and importance sampling cross-validation used to assess predictive accuracy. In the third comparison, a multi-step version of our approach was used, in which sparse cluster analysis was separated from the

TABLE 2 | Temporal lobe epilepsy dataset: (a) Posterior mean of β; (b) 95% credible interval (CI) for β; and (c) posterior probability of odds ratio >1, e.g., $\mathbb{P}[e^{\beta_j} > 1|X, Y] = \mathbb{P}[\beta_j > 0|X, Y]$, shown for proposed approach ($f = 0.1$), multi-step logistic approach, and multi-step sparse clustering approach.

	Proposed Method			Multistep logistic regression			Multistep sparse clustering		
	(a)	(b)	(c)	(a)	(b)	(c)	(a)	(b)	(c)
β_0	−0.377	(−1.588, 3.587)	0.43	−0.225	(−0.8, 1.665)	0.41	−0.169	(−0.746, 1.718)	0.43
β_1	0.368	(0.275, 0.685)	0.99	0.250	(0.183, 0.48)	0.99	0.262	(0.196, 0.486)	0.99
β_2	−3.726	(−4.639, −0.76)	0.01	−1.25	(−1.742, 0.347)	0.06	−1.247	(−1.737, 0.358)	0.06
β_3	−0.082	(−0.117, 0.032)	0.09	−0.069	(−0.094, 0.011)	0.05	−0.079	(−0.107, 0.014)	0.05
β_4	1.649	(0.832, 4.364)	0.90	0.939	(0.453, 2.552)	0.88	0.41	(−0.113, 2.140)	0.68

Here $\beta_1 \equiv$ Epilepsy duration, $\beta_2 \equiv$ History of GTC, $\beta_3 \equiv$ Age at surgery, $\beta_4 \equiv$ Cluster 1 (v. Cluster 2).
Odds are with respect to seizure freedom.

FIGURE 6 | Temporal lobe epilepsy dataset: Receiver operating characteristic curve (ROC) for proposed method, elastic net, principal components regression (PCR), multi-step logistic regression, and multi-step sparse clustering in predicting post-operative outcome 1 year after anterior temporal lobe resection.

outcome model. In particular, a greedy forward search algorithm was used for simultaneous variable selection and clustering (Raftery and Dean, 2006). Patients were clustered based on the selected variables through a Gaussian mixture model (Fraley et al., 2012) and Bayesian logistic regression then used to predict post-surgical outcome from latent class membership, with predictive accuracy assessed through importance-sampling cross-validation. Prediction results using our unified approach attained superior predictive performance compared to multi-step approaches (**Figure 6**). Multi-step logistic regression and multi-step sparse clustering approaches attained higher predictive accuracy than PCR. We also compared to methods such as elastic net (Zou and Hastie, 2005), ridge regression (Hoerl and Kennard, 1970), and the Least Absolute Shrinkage and Selection Operator (LASSO) method of Tibshirani (1996) that, in particular, do not condition on latent states, but rather use

the X data as the covariates. Penalized regression approaches that did not condition on a latent state performed poorly in data with underlying latent states (see Supplementary Material). Additionally, in the Supplementary Material, we conduct a full comparison study among competing methods on synthetic data to evaluate results for both prediction and biomarker selection.

4. DISCUSSION

Our results have identified a subgroup of temporal lobe epilepsy patients with 5.8 times greater odds of post-operative seizure recurrence after anterior temporal lobe resection. These patients were characterized by lower levels of interictal metabolism in regions near the ipsilateral parieto-temporal-occipital junction. Lower interictal metabolism in peritemporal regions may suggest structural abnormalities such as gliosis or neuronal loss in these regions, alternatively or in combination with functional abnormality involving a widespread epileptogenic network which extends beyond the temporal lobe. Evidence for such a subgroup has been suggested by previous research, which found limited improvement in seizure outcomes in patients with electrocorticographical (ECoG) evidence of extratemporal involvement of inferior parietal cortex (Aghakhani et al., 2004). The implication of extratemporal brain structures in patients with poorer postsurgical outcomes supports the presence of latent pathologies in patients with epilepsy. Other ECoG studies have also suggested the presence of latent pathology in epilepsy involving spread of the epileptogenic focus and the possible creation of secondary foci (Rougier, 1990; D'Ambrosio et al., 2005). Therefore, lower interictal metabolism in this subset of patients may suggest a subtype of MTLE-HS with parietal involvement, which may lead to post-operative seizure generation if not resected. The involvement of posterior parietal regions in this subset of patients may result from connectivity to other regions clinically involved in MTLE. Structural connectivity exists between the presubiculum and the posterior parietal cortex through the cingulum, for example, and functional connectivity between these regions also exists through the default mode network (Buckner et al., 2008). Pulvinar atrophy has also been found in TLE patients with persistent post-operative seizures (Keller

et al., 2015), so connectivity of posterior parietal regions to the pulvinar nucleus may also play a role in posterior parietal involvement.

Patients at high risk for post-operative seizure recurrence were also characterized by higher levels of interictal glucose metabolism in the cerebellum. The cerebellum's role in inhibiting seizures has been investigated since the early 1940's, following the discovery that cerebellar stimulation may result in seizure modification or even termination (Moruzzi, 1950). Recent technological advances in techniques for cerebellar stimulation have led to renewed interest in the role of cerebellar stimulation in seizure inhibition, with a 41% seizure rate reduction achieved through cerebellar stimulation (Velasco et al., 2005). Direct optogenetic stimulation of the cerebellar Purkinje cells has been found to be sufficient to reduce the duration of seizures in temporal lobe epilepsy (Krook-Magnuson et al., 2014). It is postulated that the mechanism of cerebellar stimulation in seizure inhibition may be through increased inhibitory efferent output from the Purkinje cells to the deep cerebellar nuclei, resulting in increased inhibitory cerebellar output to the thalamocortical projections and thus decreased contralateral cortical excitability (Fountas et al., 2010). Likewise, the cerebral cortex exhibits feedback to the contralateral cerebellar hemispheres through corticopontocerebellar tracts. In our study, we found that the subgroup of MTLE-HS patients at high risk for post-operative seizure recurrence was characterized by higher levels of interictal glucose metabolism in the bilateral cerebellar hemispheres and cerebellar vermis, with slightly larger marginal posterior probability of discriminating high- vs. low-risk patients in the contralateral than the ipsilateral cerebellar hemisphere. Higher interictal glucose metabolism in the cerebellum may be caused by pre-operatively increased baseline levels of cortical excitability in high-risk patients, resulting in increased activity of corticopontocerebellar white matter tracts and increased crossed cerebellar metabolism. The localization of this phenomenon may be similar to that of cerebellar diaschisis, in which supratentorial lesions such as stroke may cause disruption of corticopontocerebellar tracts and therefore contralateral cerebellar hypometabolism. In the case of epilepsy, in which there is over- rather than underactivity of the cortex, overstimulation of the corticopontocerebellar tracts may lead to contralateral cerebellar hypermetabolism. Inhibitory outflow from the Purkinje cells may then result in hypometabolic activity in areas such as the inferior parietal lobule, congruent with the functional abnormality observed in the temporo-parieto-occipital junction as described above. Our observation of bilaterally increased glucose metabolism in the cerebellum suggest bilaterally increased cortical excitability in patients at high risk for post-operative seizure recurrence, with slightly higher cortical excitability ipsilaterally. The greater contralateral cerebellar involvement observed here is also consistent with our observation of ipsilaterally involved temporo-parieto-occipital regions due to crossed cerebello-cortical connections.

In addition to enhancing understanding of the pathophysiology behind post-operative seizure recurrence, our finding that patients at high risk for epilepsy surgery failure are

characterized by lower PET metabolism in peritemporal regions and higher cerebellar metabolism, provides a marker for patients where epilepsy surgery is at high risk for failure. These patients may be better candidates for neuromodulatory treatments for medication-refractory epilepsy, such as direct cortical stimulation, as is being used in responsive neurostimulation (RNS) at regions of seizure onset (Geller et al., 2017). We show that TLE patients at high risk for anterior temporal lobe resection failure have abnormal pre-surgical brain metabolic activity compared to those patients who attain post-surgical seizure freedom, suggesting a difference in the underlying brain networks of the two groups. The approach proposed here provides a method which may potentially allow for pre-surgical differentiation between patients with abnormal underlying brain activity.

In this paper we have developed a general integrative modeling framework to characterize the association between a set of image predictors and an individual clinical outcome that simultaneously (a) identifies subgroups of patients characterized by latent pathologies differentially associated to the outcome of interest, (b) identifies discriminatory brain regions across subjects, and (c) uses prior connectivity information from external data to inform the selection of biomarkers. Our Bayesian measurement error model provides a modeling approach for the prediction of post-surgical treatment response from imaging data which explicitly accounts for the unobserved disease state. As described in section 2.3.5, our model provides an approach in which a new prospective surgery candidate can come in, be scanned with PET imaging, assigned to a latent risk group, and evaluated for their probability of achieving seizure freedom if operated upon. By accounting for heterogeneity in the unobserved state, while allowing for incorporation of external prior information, we have obtained accurate prediction in data where surrogate measures, such as neuroimaging data, are observed. We have shown that our approach achieves superior predictive performance compared to commonly used approaches, such as principal components regression, ROI-based clustering, and ROI-based sparse regression, and additionally leads to accurate inference with respect to identification of latent states and variable selection.

We have used the proposed method to analyze data we have available from the University of California, Los Angeles Seizure Disorder Center, where the interest was in predicting the post-surgical outcome among MTLE-HS patients from pre-operative FDG-PET imaging. In the analysis, we have used resting-state fMRI imaging to inform the prior model. Our analysis has identified several discriminatory ROIs, together with a subgroup of patients at higher risk of post-operative seizures recurrence. Pre-surgical identification of regions pathophysiologically involved in post-operative seizure recurrence may assist in targeting these regions for interruption. Here, patients at higher risk were characterized by lower levels of interictal glucose metabolism in the bilateral associative visual cortices, ipsilateral parietotemporal cortex, and bilateral inferior parietal lobules, and higher levels of interictal glucose metabolism in the bilateral cerebellar hemispheres and cerebellar vermis. Cross-validated prediction of post-operative

seizure freedom has achieved an AUC of 0.91 and 84% predictive accuracy, showing superior predictive performance compared to methods which do not condition on latent states. One caution in interpreting the results of this study is the moderate statistical power due to limited sample size. Future corroboration on larger samples is needed prior to use in clinical practice. Pre-surgical identification of patients at high risk of not benefiting from surgery may improve treatment planning for these patients, including the potential avoidance of surgery risks in cases with low probability of benefit.

In our study, we have utilized standard PET ROIs obtained from quantitative assessment software used in clinical practice, where PET activity in each region of interest is computed by averaging within the ROI. Similar ROI-based approaches are utilized within the standard preprocessing protocol of NeuroQ to aid clinical interpretability, and have demonstrated clinical utility in neurological disorders such as Parkinson's disease (Akdemir et al., 2014), tinnitus (Smith et al., 2007), and epilepsy (Kerr et al., 2013). However, it is important to note that voxel-based data allow for a finer-grained approach to biomarker selection and may be of interest in future applications of our methodology. Use of other well-known atlases to segment PET data, such as the Automated Anatomic Labeling (AAL) atlas, may also be useful for comparing to other studies. Rigid registration and the use of PET-to-PET registration is also susceptible to PET signal variations, with hippocampal atrophy in TLE potentially contributing further to decreased registration accuracy as well as partial voluming effects. Further improvements in predictive accuracy may be seen with alternative pre-processing methods, including registration to high-resolution structural imaging and partial volume correction.

Future applications of our method to pre-operative mapping may wish to investigate finer parcellations of the brain, to better delineate the epileptogenic zone and more directly aid pre-operative mapping. Given the routine use of fMRI and EEG in the management of patients with epilepsy, it might also be possible to extend our general model formulation to the identification of spatial fMRI markers of disease outcome while taking advantage of the temporal resolution of EEG data to construct prior connectivity networks. Finally, even though the motivating example for our proposed model has come from the prediction of post-surgical outcomes in epilepsy surgery, data from other neurological disorders may also be analyzed. In such cases, it may be of interest to extend the treatment outcome to a multinomial likelihood, with larger sample sizes needed if such analysis is desired.

ETHICS STATEMENT

In this paper we analyze data that were collected as part of a clinical study between 2007 and 2012. No new data were generated for this manuscript. The data were provided to us in computer form and were recorded in such a manner that subjects could not be identified directly or indirectly through identifiers/codes. The protocol that generated the data is now closed to enrollment.

AUTHOR CONTRIBUTIONS

SC, MG, ZH, JS, and MV contributed to the design and analysis of the work; HY, SD, and JS contributed to the data acquisition; SC, MG, and MV wrote the paper; all authors revised the manuscript critically for important intellectual content.

FUNDING

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fnins.2017.00669/full#supplementary-material

REFERENCES

Aghakhani, Y., Rosati, A., Dubeau, F., Olivier, A., and Andermann, F. (2004). Patients with temporoparietal ictal symptoms and inferomesial EEG do not benefit from anterior temporal resection. *Epilepsia* 45, 230–236. doi: 10.1111/j.0013-9580.2004.43003.x

Aitkin, M. (1991). Posterior Bayes factors. *J. R. Stat. Soc. Ser. B (Methodol.).* 53, 111–142.

Akdemir, Ü. Ö., Tokçaer, A. B., Karakus, A., and Kapucu, L. Ö. (2014). Brain 18F-FDG PET imaging in the differential diagnosis of Parkinsonism. *Clin. Nucl. Med.* 39, e220–e226. doi: 10.1097/RLU.0000000000000315

Alpert, N. M., Reilhac, A., Chio, T. C., and Selesnick, I. (2006). Optimization of dynamic measurement of receptor kinetics by wavelet denoising. *Neuroimage* 30, 444–451. doi: 10.1016/j.neuroimage.2005.09.031

Andersson, J. L. R., Jenkinson, M., and Smith, S. (2007). *Non-linear Registration Aka Spatial Normalisation FMRIB Technical Report TR07JA2*. Technical report, FMRIB Centre, Oxford, UK.

Arribas, J., Calhoun, V. D., and Adali, T. (2010). Automatic Bayesian classification of healthy controls, bipolar disorder, and schizophrenia using intrinsic connectivity maps from fMRI data. *IEEE Trans. Biomed. Eng.* 57, 2850–2860. doi: 10.1109/TBME.2010.2080679

Baccalá, L. A., Alvarenga, M. Y., Sameshima, K., Jorge, C. L., and Castro, L. H. (2004). Graph theoretical characterization and tracking of the effective neural connectivity during episodes of mesial temporal epileptic seizure. *J. Integr. Neurosci.* 3, 379–395. doi: 10.1142/S0219635204000610

Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2014). *Hierarchical Modeling and Analysis for Spatial Data*. Boca Raton, FL: Chapman and Hall/CRC.

Barbieri, M. M., and Berger, J. O. (2004). Optimal predictive model selection. *Ann. Stat.* 32, 870–897. doi: 10.1214/009053604000000238

Bonilha, L., Rorden, C., Castellano, G., Cendes, F., and Li, L. M. (2005). Voxel-based morphometry of the thalamus in patients with refractory medial temporal lobe epilepsy. *Neuroimage* 25, 1016–1021. doi: 10.1016/j.neuroimage.2004.11.050

Bowman, F., Zhang, L., Derado, G., and Chen, S. (2012). Determining functional connectivity using fMRI data with diffusion-based anatomical weighting. *NeuroImage* 62, 1769–1779. doi: 10.1016/j.neuroimage.2012.05.032

Buckner, R. L., Andrews-Hanna, J. R., and Schacter, D. L. (2008). The brain's default network: anatomy, function, and relevance to disease. *Ann. N.Y. Acad. Sci.* 1124, 1–38. doi: 10.1196/annals.1440.011

Burge, J., Lane, T., Link, H., Qiu, S., and Clark, V. P. (2009). Discrete dynamic Bayesian network analysis of fMRI data. *Hum. Brain Mapp.* 30, 122–137. doi: 10.1002/hbm.20490

Carroll, R. J., Ruppert, D., Stefanski, L. A., and Crainiceanu, C. M. (2006). *Measurement Error in Nonlinear Models: A Modern Perspective*. Boca Raton, FL: CRC Press.

Centeno, M., and Carmichael, D. W. (2014). Network connectivity in epilepsy: resting state fMRI and EEG–fMRI contributions. *Front. Neurol.* 5:93. doi: 10.3389/fneur.2014.00093

Chiang, S., Cassese, A., Guindani, M., Vannucci, M., Yeh, H. J., Haneef, Z., et al. (2016). Time-dependence of graph theory metrics in functional connectivity analysis. *NeuroImage* 125, 601–615. doi: 10.1016/j.neuroimage.2015.10.070

Chiang, S., and Haneef, Z. (2014). Graph theory findings in the pathophysiology of temporal lobe epilepsy. *Clin. Neurophysiol.* 125, 1295–1305. doi: 10.1016/j.clinph.2014.04.004

Chiang, S., Haneef, Z., Stern, J. M., and Engel, J. (2017a). Use of resting-state fMRI in planning epilepsy surgery. *Neurol. India* 65, S25–S33. doi: 10.4103/neuroindia.NI_823_16

Chiang, S., Vankov, E. R., Yeh, H. J., Guindani, M., Vannucci, M., Haneef, Z., et al. (2017b). Temporal and spectral characteristics of dynamic functional connectivity between resting-state networks reveal information beyond static connectivity.

D'Ambrosio, R., Fender, J. S., Fairbanks, J. P., Simon, E. A., Born, D. E., Doyle, D. L., et al. (2005). Progression from frontal–parietal to mesial–temporal epilepsy after fluid percussion injury in the rat. *Brain* 128, 174–188. doi: 10.1093/brain/awh337

de Tisi, J., Bell, G. S., Peacock, J. L., McEvoy, A. W., Harkness, W. F., Sander, J. W., et al. (2011). The long-term outcome of adult epilepsy surgery, patterns of seizure remission, and relapse: a cohort study. *Lancet* 378, 1388–1395. doi: 10.1016/S0140-6736(11)60890-8

Derado, G., Bowman, F. D., Zhang, L., and Alzheimer's Disease Neuroimaging Initiative Investigators (2013). Predicting brain activity using a Bayesian spatial model. *Stat. Methods Med. Res.* 22, 382–397. doi: 10.1177/0962280212448972

Dupont, S., Semah, F., Clémenceau, S., Adam, C., Baulac, M., and Samson, Y. (2000). Accurate prediction of postoperative outcome in mesial temporal lobe epilepsy: a study using positron emission tomography with 18-Fluorodeoxyglucose. *Arch. Neurol.* 57, 1331–1336. doi: 10.1001/archneur.57.9.1331

Engel, J., McDermott, M. P., Wiebe, S., Langfitt, J. T., Stern, J. M., Dewar, S., et al. (2012). Early surgical therapy for drug-resistant temporal lobe epilepsy: a randomized trial. *JAMA* 307, 922–930. doi: 10.1001/jama.2012.220

Ercoli, L., Small, G., Siddarth, P., Kepe, V., Huang, S.-C., Miller, K., et al. (2012). Assessment of dementia risk in aging adults using both FDG-PET and FDDNP-PET imaging. *Int. J. Geriatr. Psychiatry* 27, 1017–1027. doi: 10.1002/gps.2816

Fountas, K. N., Kapsalaki, E., and Hadjigeorgiou, G. (2010). Cerebellar stimulation in the management of medically intractable epilepsy: a systematic and critical review. *Neurosurg. Focus* 29:E8. doi: 10.3171/2010.5.FOCUS10111

Fraley, C., Raftery, A., Murphy, T., and Scrucca, L. (2012). *mclust Version 4 for R: Normal Mixture Modeling for Model-based Clustering, Classification, and Density Estimation*. Technical Report No. 597, Department of Statistics, University of Washington.

Franaszczuk, P. J., and Bergey, G. K. (1998). Application of the directed transfer function method to mesial and lateral onset temporal lobe seizures. *Brain Topogr.* 11, 13–21. doi: 10.1023/A:1022262318579

Franaszczuk, P. J., Bergey, G. K., and Kamiński, M. J. (1994). Analysis of mesial temporal seizure onset and propagation using the directed transfer function method. *Electroencephalogr. Clin. Neurophysiol.* 91, 413–427. doi: 10.1016/0013-4694(94)90163-5

Gelfand, A. E. (1996). "Model determination using sampling-based methods," in *Markov Chain Monte Carlo in Practice*, Chap. 9, eds W. R. Gilks, S. Richardson, and D. Spiegelhalter (London: Chapman and Hall), 145–161.

Geller, E. B., Skarpaas, T. L., Gross, R. E., Goodman, R. R., Barkley, G. L., Bazil, C. W., et al. (2017). Brain-responsive neurostimulation in patients with medically intractable mesial temporal lobe epilepsy. *Epilepsia* 58, 994–1004. doi: 10.1111/epi.13740

Geweke, J. (1991). *Evaluating the Accuracy of Sampling-based Approaches to the Calculation of Posterior Moments*, vol. 196. Minneapolis, MN: Federal Reserve Bank of Minneapolis, Research Department Minneapolis.

Goldsmith, J., Huang, L., and Crainiceanu, C. M. (2014). Smooth scalar-on-image regression via spatial Bayesian variable selection. *J. Comput. Graph. Stat.* 23, 46–64. doi: 10.1080/10618600.2012.743437

Greve, D. N., and Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage* 48, 63–72. doi: 10.1016/j.neuroimage.2009.06.060

Haynes, J.-D., and Rees, G. (2005). Predicting the stream of consciousness from activity in human visual cortex. *Curr. Biol.* 15, 1301–1307. doi: 10.1016/j.cub.2005.06.026

Hiden, J., and Glasziou, P. (1996). Regret graphs, diagnostic uncertainty and Youden's Index. *Stat. Med.* 15, 969–986.

Hinne, M., Ambrogioni, L., Janssen, R. J., Heskes, T., and van Gerven, M. A. (2014). Structurally-informed bayesian functional connectivity analysis. *Neuroimage* 86, 294–305. doi: 10.1016/j.neuroimage.2013.09.075

Hoerl, A. E., and Kennard, R. W. (1970). Ridge regression: biased estimation for nonorthogonal problems. *Technometrics* 12, 55–67. doi: 10.1080/00401706.1970.10488634

Hoff, P. D. (2006). Model-based subspace clustering. *Bayesian Anal.* 1, 321–344. doi: 10.1214/06-BA111

Honey, C., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.-P., Meuli, R., et al. (2009). Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci. U.S.A.* 106, 2035–2040. doi: 10.1073/pnas.0811168106

Insel, T., and Cuthbert, B. (2015). Brain disorders? Precisely. *Science* 348, 499–500. doi: 10.1126/science.aab2358

Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17, 825–841. doi: 10.1006/nimg.2002.1132

Jiang, C.-R., Aston, J. A., and Wang, J.-L. (2009). Smoothing dynamic positron emission tomography time courses using functional principal components. *NeuroImage* 47, 184–193. doi: 10.1016/j.neuroimage.2009.03.051

Jiang, H., and Ogden, R. T. (2008). Mixture modeling for dynamic PET data. *Stat. Sin.* 18, 1341–1356.

Jirsch, J., Urrestarazu, E., LeVan, P., Olivier, A., Dubeau, F., and Gotman, J. (2006). High-frequency oscillations during human focal seizures. *Brain* 129, 1593–1608. doi: 10.1093/brain/awl085

Jorge, J., van der Zwaag, W., and Figueiredo, P. (2014). EEG-fMRI integration for the study of human brain function. *Neuroimage* 102, 24–34. doi: 10.1016/j.neuroimage.2013.05.114

Keller, S. S., Richardson, M. P., Schoene-Bake, J.-C., O'Muircheartaigh, J., Elkommos, S., Kreilkamp, B., et al. (2015). Thalamotemporal alteration and postoperative seizures in temporal lobe epilepsy. *Ann. Neurol.* 77, 760–774. doi: 10.1002/ana.24376

Kerr, W. T., Nguyen, S. T., Cho, A. Y., Lau, E. P., Silverman, D. H., Douglas, P. K., et al. (2013). Computer-aided diagnosis and localization of lateralized temporal lobe epilepsy using interictal FDG-PET. *Front. Neurol.* 4:31. doi: 10.3389/fneur.2013.00031

Kramer, M. A., Kolaczyk, E. D., and Kirsch, H. E. (2008). Emergent network topology at seizure onset in humans. *Epilepsy Res.* 79, 173–186. doi: 10.1016/j.eplepsyres.2008.02.002

Krook-Magnuson, E., Szabo, G. G., Armstrong, C., Oijala, M., and Soltesz, I. (2014). Cerebellar directed optogenetic intervention inhibits spontaneous hippocampal seizures in a mouse model of temporal lobe epilepsy. *eNeuro* 1, 1–15. doi: 10.1523/ENEURO.0005-14.2014

LaConte, S., Strother, S., Cherkassky, V., Anderson, J., and Hu, X. (2005). Support vector machines for temporal classification of block design fMRI data. *NeuroImage* 26, 317–329. doi: 10.1016/j.neuroimage.2005.01.048

Li, F., and Zhang, N. R. (2010). Bayesian variable selection in structured high-dimensional covariate spaces with applications in genomics. *J. Am. Stat. Assoc.* 105, 1202–1214. doi: 10.1198/jasa.2010.tm08177

Li, F., Zhang, T., Wang, Q., Gonzalez, M. Z., Maresh, E. L., Coan, J. A., et al. (2015). Spatial Bayesian variable selection and grouping for high-dimensional scalar-on-image regression. *Ann. Appl. Stat.* 9, 687–713. doi: 10.1214/15-AOAS818

Lin, T. W., de Aburto, M. A. K., Dahlbom, M., Huang, L. L., Marvi, M. M., Tang, M., et al. (2007). Predicting seizure-free status for temporal lobe epilepsy patients undergoing surgery: prognostic value of quantifying maximal metabolic asymmetry extending over a specified proportion of the temporal lobe. *J. Nucl. Med.* 48, 776–782. doi: 10.2967/jnumed.106.034249

Lüders, H., Engel, J., and Munari, C. (1993). "General principles," in *Surgical Treatment of the Epilepsies, 2nd Edn.* eds J. Engel Jr. (New York, NY: Raven Press), 609–622.

Luders, H. O. (2008). *Textbook of Epilepsy Surgery*. Boca Raton, FL: CRC Press.

Ma, S., Calhoun, Vince D anfd Phlypo, R., and Adalı, T. (2014). Dynamic changes of spatial functional network connectivity in healthy individuals and schizophrenia patients using independent vector analysis. *Neuroimage* 90, 196–206. doi: 10.1016/j.neuroimage.2013.12.063

McCallum, R., Dusing, R., Sarosiek, I., Cocjin, J., Forster, J., and Lin, Z. (2010). Mechanisms of symptomatic improvement after gastric electrical stimulation in gastroparetic patients. *Neurogastroenterol. Motil.* 22:161. doi: 10.1111/j.1365-2982.2009.01389.x

McDonald, C. R., Hagler, D. J., Ahmadi, M. E., Tecoma, E., Iragui, V., Gharapetian, L., et al. (2008). Regional neocortical thinning in mesial temporal lobe epilepsy. *Epilepsia* 49, 794–803. doi: 10.1111/j.1528-1167.2008.01539.x

Meunier, D., Lambiotte, R., and Bullmore, E. T. (2010). Modular and hierarchically modular organization of brain networks. *Front. Neurosci.* 4:200. doi: 10.3389/fnins.2010.00200

Millet, P., Ibáñez, V., Delforge, J., Pappata, S., and Guimón, J. (2000). Wavelet analysis of dynamic PET data: application to the parametric imaging of benzodiazepine receptor concentration. *Neuroimage* 11, 458–472. doi: 10.1006/nimg.2000.0563

Mitchell, T., Hutchinson, R., Niculescu, R., Pereira, F., Wang, X., Just, M., et al. (2004). Learning to decode cognitive states from brain images. *Mach. Learn.* 57, 145–175. doi: 10.1023/B:MACH.0000035475.85309.1b

Moruzzi, G. (1950). Effects at different frequencies of cerebellar stimulation upon postural tonus and myotatic reflexes. *Electroencephalogr. Clin. Neurophysiol.* 2, 463–469. doi: 10.1016/0013-4694(50)90083-6

Mueller, S. G., Laxer, K. D., Barakos, J., Cheong, I., Finlay, D., Garcia, P., et al. (2010). Involvement of the thalamocortical network in tle with and without mesiotemporal sclerosis. *Epilepsia* 51, 1436–1445. doi: 10.1111/j.1528-1167.2009.02413.x

Munsell, B. C., Wee, C.-Y., Keller, S. S., Weber, B., Elger, C., da Silva, L. A. T., et al. (2015). Evaluation of machine learning algorithms for treatment outcome prediction in patients with epilepsy based on structural

connectome data. *NeuroImage* 118, 219–230. doi: 10.1016/j.neuroimage.2015.06.008

Njiwa, J. Y., Gray, K., Costes, N., Mauguiere, F., Ryvlin, P., and Hammers, A. (2015). Advanced [18 F] FDG and [11 C] flumazenil PET analysis for individual outcome prediction after temporal lobe epilepsy surgery for hippocampal sclerosis. *NeuroImage* 7, 122–131. doi: 10.1016/j.nicl.2014.11.013

O'Sullivan, F. (2006). Locally constrained mixture representation of dynamic imaging data from PET and MR studies. *Biostatistics* 7, 318–338. doi: 10.1093/biostatistics/kxj010

Polson, N. G., Scott, J. G., and Windle, J. (2013). Bayesian inference for logistic models using Polya-Gamma latent variables. *J. Am. Stat. Assoc.* 108, 1339–1349. doi: 10.1080/01621459.2013.829001

Quigley, M., Cordes, D., Turski, P., Moritz, C., Haughton, V., Seth, R., et al. (2003). Role of the corpus callosum in functional connectivity. *Am. J. Neuroradiol.* 24, 208–212.

Raftery, A. E., and Dean, N. (2006). Variable selection for model-based clustering. *J. Am. Stat. Assoc.* 101, 168–178. doi: 10.1198/016214506000000113

Raftery, A. E., and Lewis, S. (1992). How many iterations in the Gibbs sampler. *Bayesian Stat.* 4, 763–773.

Raftery, A. E., and Lewis, S. M. (1992). One long run with diagnostics: implementation strategies for Markov chain Monte Carlo. *Stat. Sci.* 7, 493–497. doi: 10.1214/ss/1177011143

Richardson, S., and Gilks, W. (1993). Conditional independence models for epidemiological studies with covariate measurement error. *Stat. Med.* 12, 1703–1722. doi: 10.1002/sim.4780121806

Rosenow, F., and Lüders, H. (2001). Presurgical evaluation of epilepsy. *Brain* 124, 1683–1700. doi: 10.1093/brain/124.9.1683

Rougier, A. (1990). The epileptic focus versus the pathological focus. *Acta Neurochir. Suppl. (Wien)* 50, 1–5.

Ryvlin, P., and Kahane, P. (2005). The hidden causes of surgery-resistant temporal lobe epilepsy: extratemporal or temporal plus? *Curr. Opin. Neurol.* 18, 125–127. doi: 10.1097/01.wco.0000162852.22026.6f

Sandok, E. K., O'brien, T. J., Jack, C. R., and So, E. L. (2000). Significance of cerebellar atrophy in intractable temporal lobe epilepsy: a quantitative MRI study. *Epilepsia* 41, 1315–1320. doi: 10.1111/j.1528-1157.2000.tb04611.x

Savitsky, T., Vannucci, M., and Sha, N. (2011). Variable selection for nonparametric Gaussian process priors: models and computational strategies. *Stat. Sci.* 26, 130–149. doi: 10.1214/11-STS354

Silverman, D. H., and Melega, W. P. (2004). "Molecular imaging of biological processes with PET: evaluating biologic bases of cerebral function," in *PET molecular imaging and its biological applications*, ed M. E. Phelps (New York, NY: Springer-Verlag), 509–583.

Silverman, D. H., Mosconi, L., Ercoli, L., Chen, W., and Small, G. W. (2008). Positron emission tomography scans obtained for the evaluation of cognitive dysfunction. *Semin. Nucl. Med.* 38, 251–261. doi: 10.1053/j.semnuclmed.2008.02.006

Smith, J. A., Mennemeier, M., Bartel, T., Chelette, K. C., Kimbrell, T., Triggs, W., et al. (2007). Repetitive transcranial magnetic stimulation for tinnitus: a pilot study. *Laryngoscope* 117, 529–534. doi: 10.1097/MLG.0b013e31802f4154

Smith, M., and Fahrmeir, L. (2007). Spatial Bayesian variable selection with application to functional magnetic resonance imaging. *J. Am. Stat. Assoc.* 102, 417–431. doi: 10.1198/016214506000001031

Sokoloff, L., Reivich, M., Kennedy, C., Des Rosiers, M., Patlak, C. S., Pettigrew, K., et al. (1977). The [14c] deoxyglucose method for the measurement of local cerebral glucose utilization: theory, procedure, and normal values in the conscious and anesthetized albino rat1. *J. Neurochem.* 28, 897–916. doi: 10.1111/j.1471-4159.1977.tb10649.x

Spencer, S., Berg, A., Vickrey, B., Sperling, M., Bazil, C., Shinnar, S., et al. (2005). Predicting long-term seizure outcome after resective epilepsy surgery: the multicenter study. *Neurology* 65, 912–918. doi: 10.1212/01.wnl.0000176055.45774.71

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *J. R. Stat. Soc. B (Stat. Methodol.)* 64, 583–639. doi: 10.1111/1467-9868.00353

Stingo, F. C., Guindani, M., Vannucci, M., and Calhoun, V. D. (2013). An integrative Bayesian modeling approach to imaging genetics. *J. Am. Stat. Assoc.* 108, 876–891. doi: 10.1080/01621459.2013.804409

Tai, Y.-C., Lin, K. P., Hoh, C. K., Huang, S. H., and Hoffman, E. J. (1997). Utilization of 3-D elastic transformation in the registration of chest X-ray CT and whole body PET. *IEEE Trans. Nucl. Sci.* 44, 1606–1612. doi: 10.1109/23.632740

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B (Methodol.)* 58, 267–288.

Torosyan, N., and Silverman, D. H. (2012). Neuronuclear imaging in the evaluation of dementia and mild decline in cognition. *Semin. Nucl. Med.* 42, 415–422.

van Gerven, M. A. J., Cseke, B., de Lange, F., and Heskes, T. (2010). Efficient Bayesian multivariate fMRI analysis using a sparsifying spatio-temporal prior. *NeuroImage* 50, 150–161. doi: 10.1016/j.neuroimage.2009.11.064

Velasco, F., Carrillo-Ruiz, J. D., Brito, F., Velasco, M., Velasco, A. L., Marquez, I., et al. (2005). Double-blind, randomized controlled pilot study of bilateral cerebellar stimulation for treatment of intractable motor seizures. *Epilepsia* 46, 1071–1081. doi: 10.1111/j.1528-1167.2005.70504.x

Versace, F., Engelmann, J. M., Robinson, J. D., Jackson, E. F., Green, C. E., Lam, C. Y., et al. (2014). Prequit fMRI responses to pleasant cues and cigarette-related cues predict smoking cessation outcome. *Nicotine Tobacco Res.* 16, 697–708. doi: 10.1093/ntr/ntt214

Wahl, R. L., Jacene, H., Kasamon, Y., and Lodge, M. A. (2009). From recist to percist: evolving considerations for pet response criteria in solid tumors. *J. Nucl. Med.* 50(Suppl. 1), 122S–150S. doi: 10.2967/jnumed.108.057307

Wiebe, S., Blume, W. T., Girvin, J. P., and Eliasziw, M. (2001). A randomized, controlled trial of surgery for temporal lobe epilepsy. *New Eng. J. Med.* 345, 311–318. doi: 10.1056/NEJM200108023450501

Wieser, H.-G. (2004). Ilae commission report: mesial temporal lobe epilepsy with hippocampal sclerosis. *Epilepsia* 45, 695–714. doi: 10.1111/j.0013-9580.2004.09004.x

Willmann, O., Wennberg, R., May, T., Woermann, F., and Pohlmann-Eden, B. (2007). The contribution of 18F-FDG PET in preoperative epilepsy surgery evaluation for patients with temporal lobe epilepsy: a meta-analysis. *Seizure* 16, 509–520. doi: 10.1016/j.seizure.2007.04.001

Worrell, G. A., Gardner, A. B., Stead, S. M., Hu, S., Goerss, S., Cascino, G. J., et al. (2008). High-frequency oscillations in human temporal lobe: simultaneous microwire and clinical macroelectrode recordings. *Brain* 131, 928–937. doi: 10.1093/brain/awn006

Worrell, G. A., Parish, L., Cranstoun, S. D., Jonas, R., Baltuch, G., and Litt, B. (2004). High-frequency oscillations and seizure generation in neocortical epilepsy. *Brain* 127, 1496–1506. doi: 10.1093/brain/awh149

Zhang, L., Guindani, M., Versace, F., and Vannucci, M. (2014). A spatio-temporal nonparametric Bayesian variable selection model of fMRI data for clustering correlated time courses. *NeuroImage* 95, 162–175. doi: 10.1016/j.neuroimage.2014.03.024

Zou, H., and Hastie, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. B (Stat. Methodol.)* 67, 301–320. doi: 10.1111/j.1467-9868.2005.00503.x