Published in final edited form as:

Med Image Comput Comput Assist Interv. 2013 ; 16(0 1): 98–105.

Detecting Epileptic Regions Based on Global Brain Connectivity Patterns

Andrew Sweet¹, Archana Venkataraman¹, Steven M. Stufflebeam², Hesheng Liu², Naoro Tanaka², Joseph Madsen³, and Polina Golland¹

¹MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA

²Athinoula A. Martinos Center for Biomedical Imaging, Boston, MA

³Boston Children's Hospital, Boston, MA

Abstract

We present a method to detect epileptic regions based on functional connectivity differences between individual epilepsy patients and a healthy population. Our model assumes that the global functional characteristics of these differences are shared across patients, but it allows for the epileptic regions to vary between individuals. We evaluate the detection performance against intracranial EEG observations and compare our approach with two baseline methods that use standard statistics. The baseline techniques are sensitive to the choice of thresholds, whereas our algorithm automatically estimates the appropriate model parameters and compares favorably with the best baseline results. This suggests the promise of our approach for pre-surgical planning in epilepsy.

1 Introduction

Focal epilepsy is a chronic neurological disorder, in which seizures are triggered by a few isolated regions before spreading to the rest of the brain [1]. In cases where anticonvulsant medication fails to mitigate these seizures, surgical resection of the epileptic regions may be prescribed. Accurate localization of these regions is crucial to minimize the size of the excision, and hence, to limit potential damage to brain function. For some patients, localization is achieved using intracranial electroencephalography (iEEG), in which electrodes are implanted directly onto the cortical surface. Unfortunately, iEEG is highly invasive and only provides limited coverage of the cortex.

Recently, it has been suggested that epilepsy is associated with functional disorganization during and between seizures [2]. Resting state functional MRI (rsfMRI) can help quantify this disorganization since temporal correlations in rsfMRI reflect the intrinsic functional connectivity of the brain [3]. rsfMRI is particularly attractive for epilepsy because it is non-invasive and provides full coverage of the cortex. Prior empirical studies have revealed abnormal functional connectivity in focal epilepsy patients [4], which may roughly correspond to epileptic regions [5]. However, these analyses focused on pre-defined brain networks and produced results that are sensitive to user-specified parameters. Here, we demonstrate a novel method that automatically identifies epileptic regions based on global functional connectivity patterns.

Most prior work in connectivity analysis is motivated by population studies and is ill suited to epilepsy. For example, univariate tests and random effects analysis are commonly used to identify statistical differences between a clinical population and normal controls [6]. In contrast to population studies, we cannot assume that the abnormal regions are common across patients. Furthermore, connectivity analysis typically yields discriminative

connections and provides little insight into the associated region properties. Therefore, even patient-specific connectivity analysis is not suitable for this application [7]. One solution is to aggregate population differences across connections into information about regions [8]. However, this approach still assumes a consistent set of abnormal regions for the clinical population. In contrast, our method detects abnormal regions within a heterogeneous patient group.

We demonstrate our algorithm on a case study of six epilepsy patients. Our results correspond well with the epileptic regions localized via iEEG.

2 Extracting Diseased Regions from Connectivity

Fig. 1 illustrates our assumptions about the relationship between the diseased brain regions and the observed abnormalities in functional networks. Our model operates on a parcellation of the brain into regions that are consistently defined across subjects. In this work, we subdivide the cortical surface into 50–100mm² patches, which are comparable in size to the coverage of a single iEEG electrode. Optimizing the parcellation to maximize detection accuracy is a non-trivial problem that we leave for future work.

We assume that epileptic regions are the foci of abnormal neural communications in the brain. Hence, they are associated with the greatest deviations from the functional connectivity template of a control population. Below we formalize the random variables in our model and summarize the corresponding inference algorithm to fit the model to the data. We then describe how to evaluate the the detection performance of our method.

Diseased Regions

The binary vector $R^m = [R_1^m, \dots, R_N^m]$ indicates the state, healthy ($R_i^m = 0$) or diseased ($R_i^m = 1$), for each region $i \in \{1, \dots, N\}$ in patient m , $\forall m = 1, \dots, M$. We assume an *i.i.d.* Bernoulli prior for R_i^m with the unknown parameter π^r shared across regions and patients, i.e., $P(R_i^m = 1) = \pi^r$.

Latent Connectivity

The labels R^m imply a graph of abnormal functional connectivity, which emanates from diseased regions based on a simple set of rules: (1) a connection between two diseased regions is always abnormal, (2) a connection between two healthy regions is always healthy, and (3) a connection between a healthy and a diseased region is abnormal with probability η .

We use latent functional connectivity variables F_{ij} and $\bar{F}_{ij}^m$ to model the neural synchrony between regions i and j in the control population and in patient m , respectively. Formally, the latent functional connectivity template F_{ij} of the control population is a tri-state random variable drawn from a multinomial distribution with parameter π^f . These states represent little or no functional co-activation ($F_{ij} = 0$), positive functional synchrony ($F_{ij} = 1$), and negative functional synchrony ($F_{ij} = -1$). For notational convenience, we represent F_{ij} as a length-three indicator vector $[F_{ij,-1} \ F_{ij,0} \ F_{ij,1}]$ with exactly one of its elements equal to one, i.e., $P(F_{ijk} = 1) = \pi_k^f$.

Ideally, $\bar{F}_{ij}^m \neq F_{ij}$ for abnormal connections and $\bar{F}_{ij}^m = F_{ij}$ for healthy connections. To account for noise and subject variability, we assume that the latent connectivity can deviate from the above rules with probability ε .

$$P(\overline{F}_{ij}^m | F_{ij}, R_i, R_j) = \begin{cases} (1-\varepsilon)^{F_{ij}^T \overline{F}_{ij}^m (\frac{\varepsilon}{2})^{1-F_{ij}^T \overline{F}_{ij}^m}}, & R_i^m = R_j^m = 0, \\ \varepsilon^{F_{ij}^T \overline{F}_{ij}^m (\frac{1-\varepsilon}{2})^{1-F_{ij}^T \overline{F}_{ij}^m}}, & R_i^m = R_j^m = 1, \\ \varepsilon_1^{F_{ij}^T \overline{F}_{ij}^m (\frac{1-\varepsilon_1}{2})^{1-F_{ij}^T \overline{F}_{ij}^m}}, & R_i^m \neq R_j^m, \end{cases} \quad (1)$$

such that $\varepsilon_1 = \eta\varepsilon + (1 - \eta)(1 - \varepsilon)$. The first condition in Eq. (1) states that if both regions are healthy ($R_i^m = R_j^m = 0$), then the edge $\langle i, j \rangle$ is healthy and the functional connectivity of patient m is equal to that of the control population with probability $1 - \varepsilon$, and it differs with probability ε . The second term is similarly obtained by replacing ε with $1 - \varepsilon$. The probability ε_1 in the third condition reflects the coupling between η and ε when the region labels differ.

Although we specify separate region and connectivity variables for each patient, the parameters $\{\pi^r, \eta, \varepsilon\}$ associated with the disease are shared across the patient group. Under this assumption, the characteristics of change are common across patients, but the disease is localized to a different subset of regions in each individual. Our model can also be applied to a single patient.

Data Likelihood

The rsfMRI correlation B_{ij}^l is a noisy observation of the functional connectivity template F_{ij} , i.e., $P(B_{ij}^l | F_{ij} = 1; \{\mu, \sigma^2\}) = \mathcal{N}(B_{ij}^l; \mu_k, \sigma_k^2)$, where $\mathcal{N}(\cdot; \mu, \sigma^2)$ is a Gaussian distribution with mean μ and variance σ^2 . We fix $\mu_0 = 0$ to center the parameter estimates. The likelihood of $\overline{B}_{ij}^m$ has the same functional form and parameter values, but uses the latent functional connectivity $\overline{F}_{ij}^m$ of patient m instead of the control template F_{ij} .

Approximate Inference

We combine the prior and likelihood terms to obtain the full probability distribution for the generative model in Fig. 1(c). Our goal is to estimate the region labels $\{R^m\}$ from the observed rsfMRI correlations $\{B, \overline{B}\}$. To improve robustness of the estimation, we marginalize out the latent functional connectivity $\overline{F}^m$ for all patients $m = 1, \dots, M$.

The resulting expressions are heavily coupled across patients and across pairwise connections. Therefore, we use a fully factorized variational approximation (mean-field) to the posterior distribution of the remaining latent variables $\{R^m, \overline{F}\}$ for maximum likelihood estimation of the parameters $\Theta = \{\pi, \varepsilon, \eta, \mu, \sigma^2\}$. We emphasize that both the posterior distributions of the latent variables and the non-random model parameters Θ are estimated directly from the data.

The marginal posterior probability $\hat{p}_i^m = P(R_i^m = 1 | B, \overline{B}; \hat{\Theta})$ quantifies how likely region i in patient m is to be diseased given the observed connectivity data $\{B, \overline{B}\}$ and the parameter estimates $\hat{\Theta}$.

Baseline Methods

Our generative framework automatically infers the region labels based on global connectivity patterns. To evaluate the accuracy and stability of our approach, we consider two baseline methods that also translate connection information into region properties.

The first method counts the number of connections that differ from a control population. Formally, we quantify the deviation associated with connection $\langle i, j \rangle$ in patient m via the z -

statistic $z_{ij}^m = (\bar{B}_{ij}^m - m_{ij}) / s_{ij}$, where m_{ij} and s_{ij} are the mean and standard deviation of the corresponding rsfMRI correlations $\{B_{ij}^l: l=1 \dots L\}$ within the healthy population. The *connectivity* statistic summarizes the deviations associated with region i in patient m as the proportion of significantly different connections: $\hat{z}_i^m(\alpha) = \frac{1}{N-1} \sum_{j \neq i} \mathbb{1}(|z_{ij}^m| > \alpha)$, where α is a user-specified significance threshold and $\mathbb{1}(\cdot)$ is a function that is equal to 1 if its argument is true and is 0 otherwise. The absolute value accommodates both positive and negative correlation differences.

The second method computes the degree d_i^m of region i in patient m by counting the number of connections with rsfMRI correlation above a user-specified threshold β , i.e.,

$d_i^m = \sum_{j \neq i} \mathbb{1}(\bar{B}_{ij}^m > \beta)$. The associated *degree* statistic $\bar{z}_i^m(\beta)$ quantifies how abnormal the degree of a node is relative to the null distribution estimated from the normal population. This statistic is closely related to the approach of [5], which, to the best of our knowledge, is the only existing method to localize epileptic regions based on whole-brain rsfMRI connectivity analysis.

Model Evaluation

The iEEG electrode labels indicate one of three possible scenarios: abnormal activity as a seizure begins (*ictal*), abnormal activity between seizures (*interictal*), and no abnormal activity. Our goal is to identify the ictal areas that correspond to epileptic regions while simultaneously avoiding detection in the areas of normal activity.

As the iEEG electrodes lie on the surface of the brain, the cortical origin of the abnormal activity measured at each electrode is uncertain. This uncertainty is exacerbated by potential misalignment and significant brain shift due to the required craniotomy. Therefore, quantifying the agreement between the electrode labels and the diseased regions identified by the methods is not necessarily helpful. Instead, we qualitatively evaluate the performance of each method by visual comparison. We deem successful detections to be those that overlap with or are immediately adjacent to the ictal areas. We emphasize that the electrode grids cover only a fraction of the cortical surface; we cannot draw conclusions about any detections outside of this coverage.

3 Experimental Results

We have performed extensive simulations on synthetic data, in which the observations are sampled from our model. The experiments demonstrate that our inference algorithm recovers the true region labels for a wide range of model initializations. We omit these results here and instead focus on the clinical findings.

Data and Pre-processing

We illustrate our method on a clinical study of six focal epilepsy patients. For each patient, the data includes an anatomical scan (MPRAGE, TE=3.44ms, FOV=256mm $\times$ 256mm, res=1mm³), a between-seizure rsfMRI scan (EPI, 152–456 vols, TR=5s, TE=30ms, res=2mm³), a CT volume acquired after iEEG electrode implantation (res=0.5 $\times$ 0.5 $\times$ 2.5–5mm), and the iEEG electrode labels. Anatomical and rsfMRI scans were acquired for 38 control subjects using the same imaging protocols.

We uniformly subdivide the Freesurfer cortical surface template [9] into $N = 1153$ regions and non-linearly register the resulting parcellation to the MNI152 template [10]. Our rsfMRI processing pipeline includes motion correction via rigid registration, slice timing correction and spatial normalization to the MNI152 template. We then spatially smooth each volume

using a 6mm Gaussian kernel, temporally low-pass filter the time courses and remove global contributions from the white matter, ventricles and the whole brain. The rsfMRI observation B_{ij}^l is computed as the Pearson correlation coefficient between the mean time courses of region i and region j in subject l .

The CT volume is rigidly registered to the MPRAGE volume using FSL [10]. While the registration is mostly accurate, the electrodes appear within the cortical surface due to brain shift during implantation. We correct for brain shift by projecting each electrode center onto its closest vertex on the smoothed pial surface [9]. For evaluation, we project the baseline and model results onto the smoothed pial surface by associating each vertex with the maximum value along its normal, up to 20mm inside the cortex. This is because abnormal activity measured on the brain surface can originate from within the cortical ribbon.

Detecting Epileptic Regions

Fig. 2 visualizes the marginal posterior probability $\hat{p}_i^m$ obtained by our method, with the degree statistic $\bar{z}_i^m$ as a baseline result. We set the correlation threshold at $\beta = 0.5$, which yields qualitatively similar results to those presented in [5]. We display only the regions for which $\bar{z}_i^m > 2$, which roughly corresponds to an uncorrected p-value of $p < 0.05$.

Our algorithm identifies a much richer set of abnormal regions than the degree statistic $\bar{z}_i^m$. This translates to better detection of the ictal areas in Patient 1, Patient 4 and Patient 6. Our model also localizes many of the interictal areas in Patient 3, Patient 4 and Patient 5. While not strictly epileptic, interictal regions can develop after surgery to trigger seizures in the future [1]. In general, our algorithm avoids regions with normal activity. It also detects regions where there is no electrode coverage. Although we can only speculate about whether these regions are epileptic, they may be good candidate locations for electrode placement in pre-surgical planning. Patient 6 is the most difficult case as evidenced by the widespread electrode coverage. Here, both baseline methods fail to detect the frontal areas; instead, they favor regions with no electrode coverage. In contrast, our model correctly identifies all the ictal regions.

Optimizing the threshold α of the connectivity statistic against the iEEG electrode labels leads to a similar detection accuracy to that of our method but with more false detections in the areas of normal activity (not shown). The optimal parameter value varies across subjects, which highlights the challenge of using standard hypothesis testing in this application.

Parameter Sweep

Fig. 3 illustrates how varying the user-defined threshold parameters in the baseline methods affects the results. Although the connectivity statistic $\hat{z}_i^m$ can identify the ictal regions, the results are sensitive to the threshold α . In contrast, the degree statistic $\bar{z}_i^m$ marginally detects the ictal regions for any threshold β . Lowering β would include more regions, but it does not make sense to use such a low correlation threshold to identify functional connections. In contrast to the baseline methods, our algorithm is completely data-driven and automatically selects the appropriate parameter values. Empirically, we observe a sensitivity to threshold values for all six patients.

4 Conclusions

We have proposed a novel generative framework for epilepsy based on rsfMRI correlations. Our model assumes that epileptic foci induce a network of abnormal functional connectivity in the brain. The resulting algorithm consistently detects regions in the immediate vicinity of the ictal spiking areas, as localized by iEEG. Future directions include applying the method

to a larger patient cohort and evaluating the effects of region size on the detection performance. Small patches can better localize the epileptic regions but are susceptible to inter-subject variability and registration errors. Conversely, large patches mitigate these issues but may smooth away focal effects. Overall, our results illustrate the promise of our approach for pre-surgical planning in epilepsy.

Acknowledgments

This work was supported in part by the National Alliance for Medical Image Computing (NIH NIBIB NIMIC U54-EB005149), the Neuroimaging Analysis Center (NIH NCRR NAC P41-RR13218 and NIH NIBIB NAC P41-EB-015902), the NSF CAREER Grant 0642971 and the MIT Lincoln Lab Collaboration.

References

1. Lüders H, et al. The Epileptogenic Zone: General Principles. *Epileptic Disorders: International Epilepsy Journal with Videotape*. 2006; 8(suppl 2):S1–S9. [PubMed: 17012067]
2. Kramer M, et al. Epilepsy as a Disorder of Cortical Network Organization. *The Neuroscientist*. 2012; 18(4):360–372. [PubMed: 22235060]
3. Fox M, Raichle M. Spontaneous Fluctuations in Brain Activity Observed with Functional Magnetic Resonance Imaging. *Nature*. 2007; 8:700–711.
4. Luo C, et al. Disrupted Functional Brain Connectivity in Partial Epilepsy: a Resting-State fMRI Study. *PloS one*. 2011; 7(1):e28196. [PubMed: 22242146]
5. Stufflebeam S, et al. Localization of Focal Epileptic Discharges using Functional Connectivity Magnetic Resonance Imaging. *Journal of Neurosurgery*. 2011:1–5.
6. Greicius MD, et al. Resting-State Functional Connectivity in Major Depression: Abnormally Increased Contributions from Subgenual Cingulate Cortex and Thalamus. *Biological Psychiatry*. 2007; 62:429–437. [PubMed: 17210143]
7. Varoquaux, G.; Baronnet, F.; Kleinschmidt, A.; Fillard, P.; Thirion, B. Detection of Brain Functional-connectivity Difference in Post-stroke Patients using Group-level Covariance Modeling. In: Jiang, T.; Navab, N.; Pluim, JPW.; Viergever, MA., editors. *MICCAI 2010, Part I. LNCS*. Vol. 6361. Springer; Heidelberg: 2010. p. 200-208.
8. Venkataraman, A.; Kubicki, M.; Golland, P. From Brain Connectivity Models to Identifying Foci of a Neurological Disorder. In: Ayache, N.; Delingette, H.; Golland, P.; Mori, K., editors. *MICCAI 2012, Part I. LNCS*. Vol. 7510. Springer; Heidelberg: 2012. p. 715-722.
9. Fischl B. FreeSurfer. *NeuroImage*. 2012; 62:774–781. [PubMed: 22248573]
10. Smith S, et al. Advances in Functional and Structural MR Image Analysis and Implementation as FSL. *NeuroImage*. 2004; 23(51):208–219.

Fig. 1.

(a) Latent network model of diseased regions in a heterogeneous population. The red nodes are diseased regions and are unique for each patient; red edges correspond to abnormal functional connections emanating from the diseased regions. (b) Electrode arrays are placed on the surface of each patient's brain. Red circles denote electrodes that exhibit abnormal activity. (c) Graphical representation of our generative model. Vector R^m specifies diseased regions in patient m . F_{ij} and $\bar{F}_{ij}^m$ represent the latent functional connectivity between regions i and j in the control population and in the m^{th} patient, respectively. B_{ij}^l and $\bar{B}_{ij}^m$ are the observed time course correlations in control subject l and epilepsy patient m , respectively. Boxes denote non-random parameters; circles indicate random variables; shaded variables are observed.

Fig. 2.

Marginal posterior probability $\hat{p}_i^m$ inferred by our algorithm (left) and the degree statistic $\bar{z}_i^m$ (right), projected to the smoothed pial surface of each patient. The correlation threshold is set to $\beta = 0.5$. iEEG electrodes are shown as circles with colors denoting expert labels: normal activity (black), interictal abnormal activity (yellow) and ictal abnormal activity (red). Only views with electrode coverage are shown.

Fig. 3.

Region localization for Patient 1 when sweeping the significance threshold α for the connectivity statistic $\hat{z}_i^m$ (top left) and the correlation threshold β for the degree statistic $\bar{z}_i^m$ (bottom). The model results are presented for comparison (top right). iEEG electrodes are shown as circles with colors denoting expert labels: normal activity (black), interictal abnormal activity (yellow) and ictal abnormal activity (red).