



Automated Noninvasive Seizure Detection and Localization Using Switching Markov Models and Convolutional Neural Networks

Jeff Craley¹(✉) , Emily Johnson² , Christophe Jouny² ,
and Archana Venkataraman¹ 

¹ Department of Electrical and Computer Engineering, Johns Hopkins University,
Baltimore, USA

jcraley2@jhu.edu

² Department of Neurology, Johns Hopkins Medical Institute, Baltimore, USA

Abstract. We introduce a novel switching Markov model for combined epileptic seizure detection and localization from scalp electroencephalography (EEG). Using a hierarchy of Markov chains to fuse multichannel information, our model detects seizure onset, localizes the seizure focus, and tracks seizure activity as it spreads across the cortex. This model-based seizure tracking and localization is complemented by a nonparametric EEG likelihood using convolutional neural networks. We learn our model with an expectation-maximization algorithm that uses loopy belief propagation for approximate inference. We validate our model using leave one patient out cross validation on EEG acquired from two hospitals. Detection is evaluated on the publicly available Children’s Hospital Boston dataset. We validate both the detection and localization performance on a focal epilepsy dataset collected at Johns Hopkins Hospital. To the best of our knowledge, our model is the first to perform automated localization from scalp EEG across a heterogeneous patient cohort.

1 Introduction

Epilepsy is one of the most common neurological disorders, and 20–40% of patients are medically refractory and do not respond anti-epileptic drugs [3]. When refractory epilepsy is *focal*, i.e. originating from a single seizure onset zone [7], surgical resection of this area may be the only treatment available. Scalp EEG is the first modality used to localize the seizure onset zone. While scalp EEG is non-invasive and easy to acquire, it is plagued by artifacts, such as muscle and eye movements, which may completely obscure the seizure characteristics. In addition, visual inspection of EEG recordings is time consuming and requires extensive training due to the inherent difficulty in identifying seizure activity.

Automated seizure localization methods fall into three general categories: spike detection, source localization, and signal decomposition. Spike detectors

identify epileptiform activity between seizures. Channels containing this activity are noted as potential onset areas. However, the accuracy of these algorithms is hard to evaluate, as inter-rater agreement between clinicians for interictal epileptiform activity is low [10]. Source localization methods solve an inverse problem to identify the location within the brain the seizure activity originates from. However, these methods require manual identification of the seizure interval beforehand. Additionally, inverse methods require more expensive imaging, such as MRI, for accurate coregistration [8]. Decomposition methods, such as canonical decomposition, are used to localize the seizure in the sensor space but also require annotated onsets [9]. These methods are difficult to use in practice and do not provide much information beyond clinical review.

The work of [1,2] proposes an interesting alternative for seizure detection based on a coupled hidden Markov model (CHMM). The coupling between EEG channels acts as a spatio-temporal regularizer for the estimated seizure activity. While the CHMM achieves better detection accuracy, its ability to localize seizure activity requires heuristic evaluation of the output posterior probabilities.

We propose a novel Regime-Switching Markov Model for Propagation and Localization (R-SMMPL) and demonstrate it on both seizure detection and onset zone localization. Our model decouples detection, propagation, and localization into three interacting sets of variables. A switching variable controls the dynamic regime of the system, acting as a seizure onset and offset detector. In response to changes in this switching variable, we use a modified CHMM [1] to track the spread of seizure activity when seizures are detected. Our formulation includes a set of hierarchically linked location variables which allows us to tie onset location distributions between multiple recordings. We equip the R-SMMPL with nonparametric likelihoods using convolutional neural networks (CNNs) as in [2].

Unlike prior work, the R-SMMPL allows us to pool information across multiple seizure recordings into an onset zone hypothesis for each patient. Furthermore, it can easily incorporate expert information about the seizure onset times and locations. We validate our model on Johns Hopkins Hospital (JHH) dataset containing exclusively focal seizures and the Children’s Hospital Boston (CHB) dataset of unspecified seizure types. To our knowledge, the R-SMMPL is the first unified framework for both seizure detection and localization from scalp EEG.

2 R-SMMPL Formulation

Figure 1 shows our graphical model (left) and variable descriptions (right). The plate notation in Fig. 1 describes how multiple seizure recordings are aggregated for a single patient. Bold variables represent collections across time and, if applicable, EEG channel. Figure 3 illustrates the temporal evolution of a seizure as represented by the inner plate of Fig. 1. Seizure propagation pathways are defined by the graph \mathcal{S} , shown by the channel nodes and blue lines in Fig. 3. Notice that we have coupled neighboring and contralateral channels, as these are the most common propagation patterns observed in EEG seizure recordings.

Let N be the total number of patients, J^n be the number of recordings belonging to patient n , and let superscript nj denote recording j in patient n . M is the

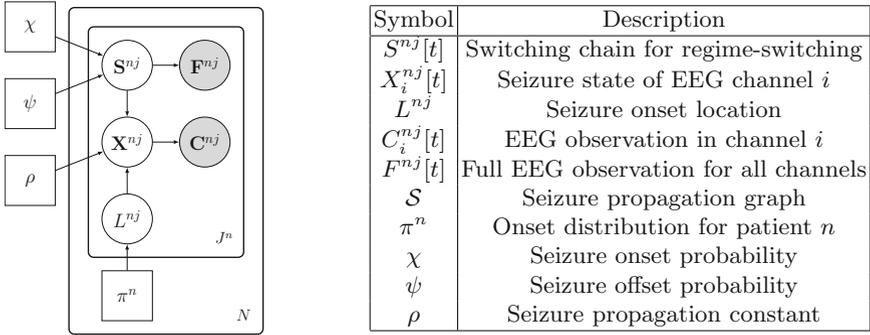


Fig. 1. Left: R-SMMPL plate model. Squares denote parameters while circles indicate random variables. Observed variables are shaded gray. **Right:** Variable descriptions. $\mathbf{S}^{nj} \triangleq \{S^{nj}[t]\}_{t=0}^T$, $\mathbf{X}^{nj} \triangleq \{X_i^{nj}[t]\}_{t=0, i=1}^{T, M}$, and similarly for \mathbf{F}^{nj} and \mathbf{C}^{nj} , respectively.

number of EEG channels (typically 18–20) and T is the recording duration. The switching chain $S^{nj}[t]$ tracks the overall state of the system as a seizure occurs and progresses. The chains $X_i^{nj}[t]$ track the spread of seizure activity through EEG channel i . Each recording has an onset location $L^{nj} \in \{1, 2, \dots, M\}$. Emission variables $F^{nj}[t]$ and $C_i^{nj}[t]$ are observed from the switching chain S^{nj} and the individual CHMM chains $X_i^{nj}[t]$, respectively. The joint distribution is:

$$P(\mathbf{L}, \mathbf{S}, \mathbf{X}, \mathbf{F}, \mathbf{C}) = \prod_{n=1}^N \prod_{j=1}^{J^n} P(L^{nj}) \prod_{t=1}^T P(S^{nj}[t] | S^{nj}[t-1]) P(F^{nj}[t] | S^{nj}[t]) \prod_{i=1}^M P(X_i^{nj}[t] | X_i^{nj}[t-1], L^{nj}, S^{nj}[t], X_{neS(i)}^{nj}[t-1]) P(C_i^{nj}[t] | X_i^{nj}[t])$$

Localization: For each patient, a multinomial location parameter π^n represents the probability that a seizure from patient n will exhibit onset in a particular EEG channel. For each recording, the onset location L^{nj} is drawn from π^n .

Regime-Switching and Propagation: The variables $S^{nj}[t]$ progress through five states: pre-seizure baseline, seizure onset, seizure spreading, seizure offset, and post-seizure baseline. The variables $X_i^{nj}[t]$ are binary and denote either normal ($X_i^{nj}[t] = 0$) or seizure ($X_i^{nj}[t] = 1$) in channel i at time t . Each recording begins in pre-seizure baseline with all channels exhibiting normal EEG activity.

Seizure onset and spread are shown on the right side of Fig. 3. At each time step, there is probability χ of a seizure occurring, represented by the switching chain $S^{nj}[t]$ transitioning into the onset state. At onset, chain $X_{L^{nj}}^{nj}[t]$

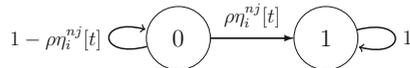


Fig. 2. Transition diagram for $X_i^{nj}[t]$ when $S^{nj}[t]$ is in the spreading state.

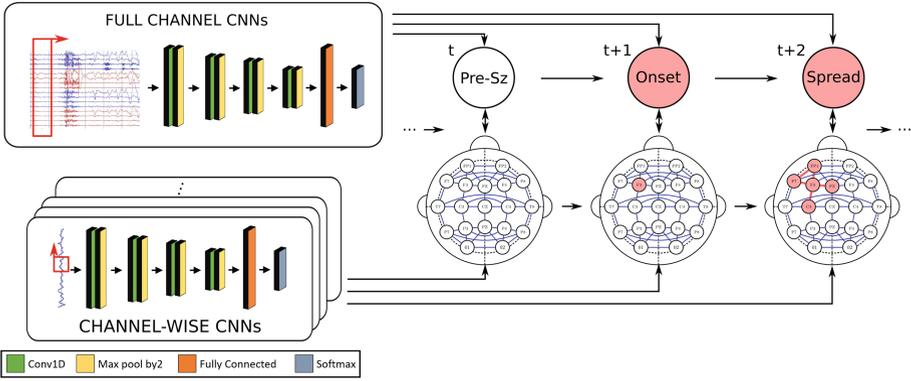


Fig. 3. Model schematic. The left side depicts the CNNs used for likelihood scoring prior to inference. The orientations of the kernels and convolutions are shown in red. At right the system is shown at seizure onset. Channel nodes and blue connections define the propagation graph \mathcal{S} . The seizure switching chain is shown above, where seizure activity is shown in red while normal activity is white. During spreading, seizure propagates through \mathbf{X}^{nj} (below) along the blue propagation pathways. (Color figure online)

enters the seizure state, representing abnormal activity at the seizure onset zone. The switching chain $S^{nj}[t]$ then immediately transitions to the spreading state.

During spreading, seizure activity is allowed to spread through the seizure propagation graph \mathcal{S} defined in Fig. 3. This spreading is governed by the probabilities in the transition diagram in Fig. 2. The probability that $X_i^{nj}[t]$ enters the seizure state (1) from the non-seizure state (0) at time t is proportionate to the number of possible ways a seizure can spread to channel i in \mathcal{S} . Let $\eta_i^{nj}[t] \triangleq \sum_{j \in ne\mathcal{S}} X_j^{nj}[t]$ be the number of neighbors in \mathcal{S} that are in the seizure state at time t . The probability $X_i^{nj}[t]$ enters the seizure state at time $t + 1$ is $\rho \eta_i^{nj}[t]$, where ρ is the parameter that governs how quickly the seizure spreads.

During the seizure, the probability of seizure offset at any time is ψ . When the switching chain $S^{nj}[t]$ enters an offset state, all EEG channels $X_i^{nj}[t]$ return to normal activity. This offset is immediately followed by a post-seizure baseline state for the remainder of the recording where no seizure activity is observed.

CNN Likelihood: Implemented in PyTorch, each CNN contains four convolution and pool layers as shown in Fig. 3. Convolution layers use eight kernels of five samples with two sample zero padding and LeakyReLU activation. Max pooling with a kernel size of two was used to halve the size of the representation at each layer. Softmax classification was performed on the concatenated result of the final pooling. All individual CNNs for $P(C_i^{nj}[t] | X_i^{nj}[t])$ were trained for 60 epochs; those for all channels $P(F^{nj}[t] | S^{nj}[t])$ were trained for 100 epochs using Adam, batches of 32 samples, a learning rate of 0.5, and cross entropy loss.

Algorithm 1. Approximate inference using loopy belief propagation.

```

1: function APPROXIMATE INFERENCE( $\mathbf{F}^{nj}, \mathbf{C}^{nj}, \chi, \psi, \rho, \pi^n$ )
2:   Pass the location variable,  $L^{nj}$ , to the  $\mathbf{X}^{nj}$  chains
3:   for Two Iterations do
4:     Forward-backward algorithm on  $\mathbf{S}^{nj}$  chain to update  $\gamma_S^{nj}[t]$ 
5:     Pass detection messages from  $\mathbf{S}^{nj}$  chain down to  $\mathbf{X}^{nj}$  chains
6:     Approximate forward-backward on  $\mathbf{X}^{nj}$  to update  $\gamma_{X_i}^{nj}[t]$ ,  $\xi_i^{nj}[t]$ , and  $\phi_i^{nj}[t]$ .
7:     Pass the  $\mathbf{X}^{nj}$  messages upward to the  $\mathbf{S}^{nj}$  chain
8:   end for
9:   Pass the  $\mathbf{X}^{nj}$  messages to  $L^{nj}$  to perform localization and update  $\tau^{nj}$ 
10:  return  $\gamma_S^{nj}[t]$ ,  $\gamma_{X_i}^{nj}[t]$ ,  $\xi_i^{nj}[t]$ ,  $\phi_i^{nj}[t]$ , and  $\tau^{nj}$ 
11: end function

```

By construction, the CNN outputs the posterior probability $P(X_i^{nj}[t] | C_i^{nj}[t])$ and $P(S^{nj}[t] | F^{nj}[t])$. Therefore, to obtain the likelihood factor, the discriminative CNN outputs are rescaled using Bayes rule, e.g. $P(C_i^{nj}[t] | X_i^{nj}[t]) \approx \frac{P(X_i^{nj}[t] | C_i[t])P(C_i^{nj}[t])}{\hat{P}(X_i^{nj}[t])} \propto \frac{P(X_i^{nj}[t] | C_i^{nj}[t])}{\hat{P}(X_i^{nj}[t])}$. We only require the likelihood up to a constant factor for inference and thus drop the $P(C_i[t])$ term. $P(X_i^{nj}[t])$ is approximated by the empirical distribution of seizure in the dataset, i.e. $\hat{P}(X_i^{nj} = 1) = \frac{\# \text{ seizure windows}}{\# \text{ windows}}$, $\hat{P}(X_i^{nj} = 0) = 1 - \hat{P}(X_i^{nj} = 1)$ as in [2].

3 Inference and Learning

The hierarchical and coupled nature of our R-SMMPL renders exact inference intractable. Therefore, we rely on loopy belief propagation [6] for approximate inference. Loopy belief propagation is a general class of algorithms where local marginal beliefs are passed as messages between neighboring random variables. These messages between represent the current local beliefs. By multiplying and summing these messages we can find posterior marginal beliefs of random variables in our model. The marginals needed for learning our model are defined below. Our message passing schedule is detailed in Algorithm 1. While loopy belief propagation provides no convergence guarantees, we observe this procedure to yield robust marginals, with little change after further message passing.

We use an expectation-maximization (EM) type algorithm for fitting the R-SMMPL to data. Our model contains three unknown transition parameters: the seizure onset probability χ , the offset probability ψ , and the spreading rate ρ . In addition we learn the onset distribution π^n for each patient. The updates are derived by setting the first derivative of the expected log-likelihood of the joint distribution with respect to the parameter of interest to zero.

The parameters χ and ψ are updated by dividing the total number of onset and offset transitions by the expected number of timesteps spent in the pre-seizure and seizure spreading state, respectively. Let the singleton posterior marginals of $S^{nj}[t]$ be defined as $\gamma_S^{nj}[t](k) \triangleq P(S^{nj}[t] = k | \mathbf{F}^{nj}, \mathbf{C}^{nj})$ where

$k = 0$ represents pre-seizure baseline and $k = 2$ is the spreading state.

$$\chi = \frac{\sum_{n=1}^N J^n}{\sum_{n=1}^N \sum_{j=1}^{J^n} \sum_{t=0}^T \gamma_S^{nj}[t](0)}, \quad \psi = \frac{\sum_{n=1}^N J^n}{\sum_{n=1}^N \sum_{j=1}^{J^n} \sum_{t=0}^T \gamma_S^{nj}[t](2)}$$

The spreading parameter ρ is updated by calculating the ratio of times channels enter a seizure state versus times channels remain in non-seizure states, weighted by the number of neighboring channels exhibiting seizure activity. Here we require the singleton and pairwise marginals $\gamma_{X_i}^{nj}[t](k) \triangleq P(X_i^{nj}[t] = k \mid \mathbf{F}^{nj}, \mathbf{C}^{nj})$ and $\xi_{X_i}^{nj}[t](k, l) \triangleq P(X_i^{nj}[t] = l, X_i^{nj}[t-1] = k \mid \mathbf{F}^{nj}, \mathbf{C}^{nj})$. In addition, let the expected number of neighbors in the seizure state be $\phi_i^{nj}[t] = E[\eta_i^{nj}[t] \mid \mathbf{F}^{nj}, \mathbf{C}^{nj}]$.

$$\rho = \frac{\sum_{n=1}^N \sum_{j=1}^{J^n} \sum_{i=1}^M \sum_{t=1}^T \xi_{X_i}^{nj}[t](0, 1)}{\sum_{n=1}^N \sum_{j=1}^{J^n} \sum_{i=1}^M \sum_{t=0}^T \phi_i^{nj}[t] \gamma_{X_i}^{nj}[t](0)}$$

Let $\tau^{nj}[t](i) \triangleq P(L^{nj} = i \mid \mathbf{F}^{nj}, \mathbf{C}^{nj})$ be the posterior belief of onset in channel i . For patient n , we update π^n via $\pi^n(i) \propto \sum_{j=1}^{J^n} \tau^{nj}(i)$. This update pools the expected beliefs regarding onset location across all of a patient's recordings.

4 Experimental Results

CHB Dataset: Our first dataset consists of publicly available seizure recordings acquired at Children's Hospital Boston (CHB) [4]. We selected 185 seizures from 24 patients for our experiment. Clinical annotations for CHB include onset and offset times. The type of seizure, general or focal, and potential onset localization are not provided. Recordings in this dataset were made at 256 Hz in a longitudinal montage using the standard 10/20 electrode placement system [5].

JHH Dataset: Our second dataset consists of focal seizure recordings from the epilepsy monitoring unit of Johns Hopkins Hospital. Expert clinical annotations from this hospital include rough onset and offset times as well as consensus of rough onset zone localizations, allowing us to evaluate both detection and localization. The dataset includes 88 seizure recordings from 15 patients. Recordings were sampled at 200 Hz using 10/20 electrode placement.

Preprocessing: We extracted seizure recordings with up to 10 min of baseline before and after the seizure annotations. Channels were normalized to mean zero and variance one. High- and low-pass filters were applied at 1.6 Hz and 50 Hz to remove DC offsets and noise. A notch filter at 60 Hz was applied to remove any possible power line contamination. One second windows with 250 ms overlap were extracted from all EEG channels. Test and train sets were separated using leave one patient out cross validation. Unlike studies which train patient-specific detectors, our evaluation focuses on generalizability to unseen patients.

Baseline Comparisons: We compare detection accuracies to the discriminative CNNs trained on the individual EEG channels (I-CNN) and those trained on all channels (S-CNN). These baselines let us assess the effect of the propagation model in fusing information across time and channels. We also compare to the CHMM model in [1] using the same I-CNN for likelihood scoring. The CHMM was shown to outperform standard machine learning classifiers in [1].

Seizure Detection: Table 1 reports the detection performance of our R-SMMPL and baseline methods. We evaluate each algorithm’s performance in terms of true positive rate (TPR), true negative rate (TNR), area under the curve (AUC), precision (P), and F1 score on a frame-wise basis. Performance was evaluated based on how well the methods detected the entire seizure interval. This evaluation is more stringent than prior work, which flags a single correct detection.

The R-SMMPL outperforms all baselines in TPR, P, and F1 scores. We observe that higher TPR comes at the cost of more false positives, reflected by lower TNR. Both the R-SMMPL and CHMM outperform the CNN baselines, illustrating the positive effect of data fusion through the use of spatio-temporal models. The main difference is that the CHMM provides localization information only via heuristic analysis, whereas the R-SMMPL provides it automatically.

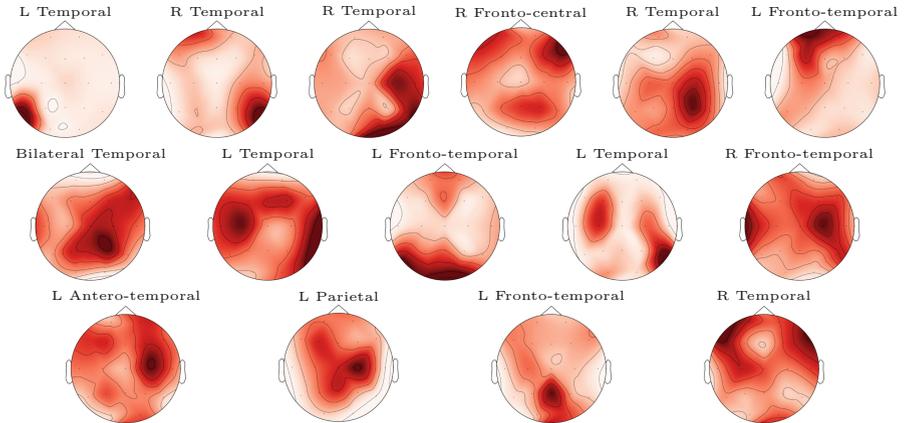


Fig. 4. Localization results from the JHH dataset. Posterior distributions over onset locations for each patient are shown with clinician provided onset diagnoses above.

Localization Results: We evaluate the localizing ability of our model when provided with a rough seizure onset time. We stipulate that the switching variable should remain in pre-seizure baseline until the clinician annotated onset. The R-SMMPL is free to switch on any time after this point.

Figure 4 shows the estimated location distribution π^n for each patient, along with the clinically diagnosed onset location. In this figure, red regions represent areas our model assigns high probability for the onset location. In the top

Table 1. Detection results for both datasets

Trial	JHH dataset					CHB dataset				
	TPR	TNR	AUC	P	F1	TPR	TNR	AUC	P	F1
R-SMMPL	0.62	0.84	<u>0.84</u>	0.65	0.53	0.67	0.94	0.86	0.58	0.58
CHMM	0.46	0.96	0.86	0.65	0.51	0.59	0.96	0.85	0.57	0.54
S-CNN	0.34	0.92	0.76	0.41	0.32	0.48	0.95	0.84	0.48	0.44
I-CNN	0.28	0.92	0.77	0.33	0.27	0.30	0.95	0.82	0.34	0.29

row we show cases where our algorithm reported a primary mode in agreement with clinical consensus of the seizure onset zone. The second row shows cases in which the secondary modes agrees with the clinical annotations. By pooling seizure localization information across all of each patient’s recordings, our model identifies likely seizure onset zones in agreement with clinical consensus in 11 of 15 patients. The R-SMMPL misidentifies the seizure onset location in just four patients. In summary, not only does the R-SMMPL *automatically* detect and track the seizure, but also leverages multiple seizure presentations to create an onset zone hypothesis for each patient. These results demonstrate the promise of R-SMMPL for clinical evaluation of epilepsy.

5 Conclusion

We have presented R-SMMPL, the first unified framework that provides clinically relevant detection and localization information from scalp EEG. R-SMMPL combines a probabilistic graphical model of seizure propagation with deep learning for data driven likelihood scoring. We derive an inference and learning procedure for the model and demonstrate its detection and localization abilities on wholly unseen patients, mirroring clinical conditions. Our methodology for automatic seizure onset zone localization by tracking seizure propagation is the first of its kind. In the future, we plan to integrate clinically informative seizure semiology into the prior distribution for the seizure onset. We also plan to reweight onset posteriors when pooling individual recordings to automatically distinguish between noisy recordings and those with better onset location evidence.

Acknowledgment. This work was supported by a JHMI Synergy Award (Venkataraman/Johnson) and NSF CAREER 1845430 (Venkataraman).

References

1. Craley, J., Johnson, E., Venkataraman, A.: A novel method for epileptic seizure detection using coupled hidden markov models. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11072, pp. 482–489. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00931-1_55

2. Craley, J., Johnson, E., Venkataraman, A.: Integrating convolutional neural networks and probabilistic graphical modeling for epileptic seizure detection in multi-channel EEG. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) IPMI 2019. LNCS, vol. 11492, pp. 291–303. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20351-1_22
3. French, J.A.: Refractory epilepsy: clinical overview. *Epilepsia* **48**, 3–7 (2007)
4. Goldberger, A.L., et al.: Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *Circulation* **101**(23), e215–e220 (2000)
5. Jurcak, V., et al.: 10/20, 10/10, and 10/5 systems revisited: their validity as relative head-surface-based positioning systems. *Neuroimage* **34**(4), 1600–1611 (2007)
6. Kschischang, F.R., et al.: Factor graphs and the sum-product algorithm. *IEEE Trans. Inf. Theory* **47**(2), 498–519 (2001)
7. Lüders, H.O., et al.: The epileptogenic zone: general principles. *Epileptic Disord.* **8**(2), 1–9 (2006)
8. Plummer, C., Harvey, A.S., Cook, M.: EEG source localization in focal epilepsy: where are we now? *Epilepsia* **49**(2), 201–218 (2008)
9. Vos, D., et al.: Canonical decomposition of ictal scalp eeg reliably detects the seizure onset zone. *NeuroImage* **37**(3), 844–854 (2007)
10. Wilson, S.B., Emerson, R.: Spike detection: a review and comparison of algorithms. *Clin. Neurophysiol.* **113**(12), 1873–1881 (2002)