

Automated Interictal Spike Detection and Source Localization in MEG using ICA and Spatio-Temporal Clustering

A. Ossadtchi¹, S. Baillet², J.C. Mosher³
D. Thyerlei⁴, W. Sutherling⁴, R.M. Leahy^{1,5}

Abstract

Objective: MEG dipole localization of epileptic spikes is useful in epilepsy surgery for mapping the extent of abnormal cortex and to focus intracranial electrodes. Visually analyzing large amounts of data produces fatigue and error. Most automated techniques are based on matching of interictal spike templates or predictive filtering of the data and do not explicitly include source localization as part of the analysis. This leads to poor sensitivity versus specificity characteristics. We describe a fully automated method that combines time-series analysis with source localization to detect clusters of focal neuronal current generators within the brain that produce interictal spike activity.

Methods: We first use an ICA (Independent Components Analysis) method to decompose the multichannel MEG data and identify those components that exhibit spike-like characteristics. From these detected spikes we then find those whose spatial topographies across the array are consistent with focal neural sources, and determine the foci of equivalent current dipoles and their associated time courses. We then perform a clustering of the localized dipoles based on distance metrics that takes into consideration both their locations and time courses. The final step of refinement consists of retaining only those clusters that are statistically significant. The average locations and time series from significant clusters comprise the final output of our method.

Results and Significance: Data were processed from four patients with partial focal epilepsy. In all three subjects for whom surgical resection was performed, clusters were found in the vicinity of the resected area.

Conclusion: The presented procedure is promising and likely to be useful to the physician as a more sensitive, automated and objective method to help in the localization of the interictal spike zone of intractable partial seizures. The final output can be visually verified by neurologists in terms of both the location and distribution of the dipole clusters and their associated time series. Due to the clinical relevance and demonstrated promise of this method, further investigation of this approach is warranted.

Key words: Interictal spike, Epilepsy, Automatic detection, Independent

¹ Signal and Image Processing Institute University of Southern California, Los Angeles, CA90089

² Cognitive Neuroscience and Brain Imaging Lab, CNRS, La Salpêtrière Hospital, Paris, France

³ Los Alamos National Laboratory, Los Alamos, NM87545

⁴ Epilepsy and Brain Mapping Program, Huntington Memorial Hospital, Pasadena, CA91105

⁵ corresponding author

This work supported by the National Institutes of Health under grants R01 EB002010 and R01 NS020806

1 Introduction

Published methods for automatic spike detection in interictal EEG and MEG data employ one or more of the following approaches: morphological analysis, template matching, predictive filtering, and independent component analysis (ICA). Methods from the first category use a morphological description of the epileptic spikes. Morphological information includes such characteristics of a waveform as sharpness, amplitude, duration and convexity (Faure, 1985). For example, (Gotman and Gloor, 1976) describe a method based on splitting the waveforms into a set of elementary half-waves. For each such half-wave, parameters describing its shape are calculated and a decision is made on the basis of comparing the estimated parameters to a set of previously determined values typical for interictal spikes. A thorough review of this and other morphological methods can be found in (Gotman, 1999, 1993).

Template matching approaches also use *a priori* information about the spike shape, which is embodied in the finite impulse response of a matched filter. A decision is made by comparing the output of the filter with a predefined threshold that controls the sensitivity vs. specificity characteristics of the method. For optimal performance, template matching methods require knowledge of the noise correlation matrix and a good agreement between the hypothesized and real spike shapes (Therrien, 1992).

The majority of the methods from these first two categories were developed for single channel data and do not take advantage of spatial structure in the measurements. Performance is also limited by the variability of interictal spike shapes between subjects as well as within a single subject (Flanagan et al., 2002). The fact that the methods were originally developed for use with EEG measurements makes it difficult to apply them to MEG data, which is characterized by a significantly lower SNR.

Methods using wavelet transforms, neural network and expert system architectures to detect interictal patterns can be also categorized as template matching approaches. One disadvantage of such methods is that features are extracted separately from each channel, which does not make good use of the spatial structure of the data. A review and comparison of these methods can be found in (Wilson and Emerson, 2002).

Use of single-channel or multichannel predictive filtering techniques for spike detection is based on an assumption of local stationarity of background activity (Diambra and Malta, 1999). Interictal spikes are then detected as a deviation from stationarity. Implicit in this approach is that non-stationary behavior comes only from interictal discharges. In reality, there are many different sources of non-stationary bursts in the data which can lead to high false

positive rates. Multichannel autoregressive models were used by (Franaszczuk and Bergey, 1999) for detecting dynamic changes in brain activity due to bursts of interictal and ictal activity. This method operates on intracranially recorded data with relatively high SNR and is capable of identifying spike trains rather than single spikes. As with the previous two classes, predictive filtering techniques for interictal spike detection were developed for EEG data and may not perform well with noisier MEG data.

Use of context for elimination of false positive detections was originally proposed by (Glover et al., 1989) who described a knowledge based system for detection of epileptic sharp waves in 16 channel EEG data. The context information extracted from electrooculogram, electrocardiogram and electromyogram is used to eliminate false positive detections not specific to abnormal epileptic activity. (Flanagan et al., 2002) proposed using sharp wave source location to eliminate events whose spatial pattern cannot be explained by a dipolar model. A source localization procedure is applied to the multichannel data around each detected sharp transient to obtain a set of dipole locations. A subset of these dipolar events is then retained through visual inspection and selection of those which form interesting clusters.

Several spike detection approaches based on ICA have recently been reported in applications to EEG recordings. (Kobayashi et al., 1999) perform both model based and real data demonstrations of the use of ICA to isolate epileptic discharges from multichannel EEG data. In this approach ICA is applied to spatio-temporal data and components resembling abnormal epileptic activity are selected by visual inspection; these components are then interpreted by an expert. (Kobayashi et al., 2002a,b) use ICA decomposition together with the RAP-MUSIC source localization approach (Mosher and Leahy, 1999) to detect potential epileptogenic regions. Rather than fitting a dipole to each independent component separately (Zhukov et al., 2000), (Kobayashi et al., 2002a) followed the multidimensional ICA paradigm (Cardoso, 1998) and defined an interictal subspace spanned by the columns of the estimated mixing matrix visually identified as corresponding to epileptic components. A dipole model is then fitted to this subspace. While this approach does make use of the spatial topography of the detected spikes, the method requires visual interpretation of independent components to identify the interictal subspace and manual cluster analysis to discard spurious sources.

In this paper we describe a novel method that combines several aspects of the approaches reviewed above but in such a way that the technique is completely automated. We first use ICA (Bell and Sejnowski, 1995) to decompose the spike-like and background components of the MEG signal into separate spatial topographies and associated time series. We then perform spike detection based on a subset of the most spiky independent components using a simple thresholding technique. We use source localization to retain only those spikes

whose topographies can be explained by a current dipole pattern with a fit of 95% or higher. Post-processing of the localized sources of spiky activity is based on the concept of spatio-temporal clustering. We group the localized dipoles into a set of spatially tight clusters with similar activation waveforms. To improve specificity we retain only those clusters that satisfy a statistical significance test. The final output of the algorithm is the set of significant clusters with their average dipole location and time series. We report on realistic simulations in which performance is evaluated by computing ROC (receiver operating characteristic) curves (Metz, 1986) for the method at each level of refinement. We demonstrate that each consecutive step improves the specificity vs. sensitivity characteristics of the method. We also present the application of our method to four clinical interictal MEG datasets and compare results with expert analysis of this data in which spikes are first visually identified and current dipoles then fitted to each identified spike. Preliminary reports on this method were presented in (Ossadtchi et al., 2001) and (Ossadtchi et al., 2002).

2 Method

2.1 The Data Model

We begin by describing our spatio-temporal model for MEG data, which is assumed to consist of spike activity, spatially and temporally correlated or coherent background activity, and additive white noise. The method for spike detection described below is based on identifying the spike-like components in this data.

We represent interictal MEG data at time t as an $[M \times 1]$ spatial vector $\mathbf{x}(t)$, which we model as a linear combination of focal source topographies corrupted by additive noise with non-stationary spatial structure:

$$\mathbf{x}(t) = \begin{bmatrix} \mathbf{a}_1, \dots, \mathbf{a}_K \end{bmatrix} \begin{bmatrix} e_1(t) \\ \vdots \\ e_K(t) \end{bmatrix} + \begin{bmatrix} \mathbf{b}_1, \dots, \mathbf{b}_L \end{bmatrix} \begin{bmatrix} p_1(t) \\ \vdots \\ p_L(t) \end{bmatrix} + \mathbf{n}(t) \quad (1)$$

where $e_k(t)$ is the time course of the k -th of K dipole components responsible for interictal spikes, with corresponding spatial topography \mathbf{a}_k (Mosher and Leahy, 1999, 1998). Time course $p_l(t)$ represents the l -th of L additional non-spike related spatially coherent sources (e.g. α , γ , δ waves, sharp sleep waves, eye-blink and cardiac artifacts) with corresponding spatial topography \mathbf{b}_l . The $[M \times 1]$ vector $\mathbf{n}(t)$ models additional incoherent noise present in the data.

We assume that each of the components responsible for interictal activity $e_k(t)$ has a heavy tailed marginal probability density function $f_k(e_k)$ at time t with the cumulative distribution function:

$$P(f_k(e_k) < x) = \frac{1}{1 + e^{-\alpha_k x}}, \quad \forall x \quad (2)$$

The form of this density is consistent with that assumed in the ICA method used for spike detection that we describe in Section 2.2.1. The corresponding spatial topography \mathbf{a}_k is assumed to be formed by a single current dipole representing focal neural activation at some unknown location in the brain.

We assume a spatially coherent background activity that exhibits correlations between sources. We model this using a multichannel autoregressive model (MAR) (Kaminski and Blinowska, 1991) in order to capture non-instantaneous interactions between sources. While this model is not explicitly used in our automated spike detection method, we do use this model to simulate data to explore the sensitivity and specificity characteristics of our approach as described in Section 3.1. The MAR model and its use in modeling spontaneous background brain activity is described in Appendix A.

2.2 Spike Detection, Localization and Clustering

2.2.1 Independent Components Analysis

Independent Components Analysis (ICA) is a method for finding statistically independent components in multisensor data. The original model for ICA assumes instantaneous mixing (Bell and Sejnowski, 1995), which holds well for the quasistatic electromagnetic properties of MEG data. We use ICA as a preprocessing step that allows us to perform spatial filtering of the data and separate interictal spikes from background activity.

Application of ICA to the data vector $\mathbf{x}(t)$ results in an unmixing matrix \mathbf{W} and a set of independent components $\mathbf{z}(t)$ related as:

$$\mathbf{z}(t) = \begin{bmatrix} z_1(t) \\ \vdots \\ z_M(t) \end{bmatrix} = \mathbf{W}\mathbf{x}(t) \quad (3)$$

where the unmixing matrix is selected to optimize some measure of statistical independence between the components of the vector $\mathbf{z}(t)$. A number of ICA methods, differing in the metrics used to measure statistical dependence, have

been described in the literature (Cardoso, 1989; Hyvarinen and Oja, 1997; Bell and Sejnowski, 1995). Here we use the Infomax method of Bell and Sejnowski for reasons described below.

Consider the case where (i) the dimension of $\mathbf{x}(t)$ is $M = K + L$ with K and L the number of spike and coherent background components respectively, (ii) each of these components is statistically independent, and (iii) the background noise process is negligible. Then the unmixing matrix W will approximately recover the components $e_k(t)$ and $p_l(t)$ in equation (1), within a permutation and scale factor. In practice the above conditions will not be met: the background components are correlated, $K + L$ may be less than M , and there is additional background noise. However, provided that the spike components are independent of the background components, and the spike SNR is reasonably large, then ICA will find an unmixing matrix that approximately separates the spike components from background components, with the remainder being due to background noise (Knuth, 99). As we describe below, this is sufficient for our purposes. Thus ICA applied to data conforming to the model in equation (1) will result in a set of independent components of which K can be attributed to interictal spike activity in the brain and L to spatially coherent electrical activity unrelated to the interictal spikes. In this setting the estimated components are no longer scaled copies of either $e_k(t)$ or $p_l(t)$, but rather linear combinations of either $e_k(t)$, $k = \{1, \dots, K\}$ or $p_l(t)$, $l = \{1, \dots, L\}$. We still benefit from such a decomposition since interictal activity time courses are sparse and thus their linear combinations will tend to preserve most of the spikes so that they are detectable in the independent components of $\mathbf{z}(t)$.

We use the Infomax method (Bell and Sejnowski, 1995) for ICA decomposition. The key idea of this approach is to find an unmixing matrix W for which the joint entropy of the quantities $y_i(t) = \phi(\mathbf{w}_i \mathbf{x}(t) + b_i)$, $i = 1 \dots M$, is maximized; $\phi()$ is a nonlinear scalar function, \mathbf{w}_i are the rows of the unmixing matrix W , and b_i are bias terms estimated from the data and are zero when the data is zero mean with a symmetric density. Using a natural gradient approach (Amari, 1998), the adaptation procedure for the $k + 1$ st iteration is (Bell and Sejnowski, 1995):

$$W_{k+1} = W_k + \mu E \left(\mathbf{I} + \phi(W \mathbf{x}(t)) \mathbf{x}(t)^T \right) W_k \quad (4)$$

where $\phi(\mathbf{y}(t)) = (\phi(y_1(t)), \dots, \phi(y_M(t)))$.

We have chosen this ICA procedure as it allows us to specify a probability density function for the interictal time series. Our analysis of the time courses of interictal paroxysmal activity allows us to conclude that the probability density function of this activity has higher kurtosis values than the normal background brain activity. High kurtosis corresponds to distributions with heavy tails, i.e. distributions for which outliers are more likely than for a

Gaussian distribution of the same variance. Similar observations regarding the non-Gaussian distribution of spike waveforms have been reported by others (Robinson et al., 2002; Kobayashi et al., 1999, 2002a). We assume the spike data follow the heavy tailed distribution in equation (2), which corresponds to the use of the nonlinear function: $\phi(x) = 1/(1 + e^{-\alpha x})$. Interestingly, the parameter α of the distribution used in ICA does not affect the results other than through a scaling of the unmixing matrix which does not affect the subsequent processing that we perform.

2.2.2 Spikyness Spectrum and Spike Detection

To select those independent components that are indicative of interictal activity we introduce a "spikyness index" which is applied to each of the components computed from equation (3):

$$I_m = \frac{\max |z_m(t)|}{\frac{1}{T} \int |z_m(t)| dt}, \quad m = 1, \dots, M. \quad (5)$$

where the maximum is computed over the entire observation interval T . This index does not necessarily detect only epileptic spikes since any component with a strong maximum-to-absolute average ratio will produce large values. However, since we later prune the components based on other spatial and temporal features, at this point in the processing we are primarily concerned with high sensitivity rather than specificity.

We define the spikyness spectrum as the rank ordered elements I_m with $I_m \geq I_{m+1}$, $m = 1, \dots, M - 1$. We can now see the effect of the unmixing performed by ICA. Fig. 1 shows a typical normalized spikyness spectrum for interictal MEG data from a patient with temporal lobe epilepsy computed on (i) the raw data, (ii) the independent components, and (iii) the right singular vectors of the principal component or singular value decomposition (SVD) of the spatial-temporal data matrix

The figure clearly shows that ICA concentrates spikyness into fewer components than either the raw data or principal components analysis.

To find candidate epileptic spikes and their sources we first select the subset \mathfrak{S} of independent components with the largest spikyness. We then create a new time course $d(t)$ by non-linearly mapping this set of independent components into a single time series according to:

$$d(t) = \max_{m \in \mathfrak{S}} |z_m(t)| \quad (6)$$

To detect spikes we then apply a simple thresholding rule. As a result of this

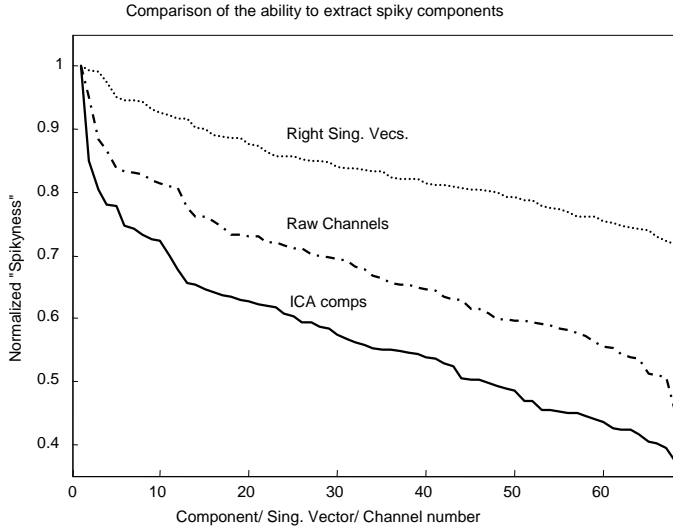


Fig. 1. Normalized "spikyness" spectrum for interictal MEG data computed from raw data and its decompositions using ICA and SVD. A decomposition that separates components containing spikes from those that do not will tend to exhibit a faster decrease in spikyness than the raw data; we refer to this characteristic as an increased degree of spikyness. Note that ICA increases the degree of spikyness compared to raw data while SVD actually reduces it.

spike detection procedure we obtain a vector of time markers $\boldsymbol{\theta} = [\theta_1, \dots, \theta_N]$ that point to the indices corresponding to the maximum of each detected spike.

2.2.3 Candidate spike localization and RAP-MUSIC

To improve the specificity of the detection procedure we select the subset of the detected spikes that can be localized as focal neuronal sources. To do this we fit an equivalent current dipole model to the data in the vicinity of each spike and accept the solution only if there is at least a 95% fit of a dipole model to the data.

We use the RAP-MUSIC algorithm (Mosher and Leahy, 1999) to localize the equivalent current dipoles from spatio-temporal data over a time window of $\pm\tau$, with $\tau = 16ms$, around the apex of each detected spike. The spatio-temporal data is first filtered by retaining only the temporal components and their corresponding spatial topographies from the set \mathfrak{S} of most spiky components i.e. we form a lower rank approximation $\hat{\boldsymbol{x}}(t)$ of the vector of measurements $\boldsymbol{x}(t)$ as

$$\hat{\boldsymbol{x}}(t) = \sum_{m \in \mathfrak{S}} \boldsymbol{v}_m z_m(t) \quad (7)$$

where \mathbf{v}_m is the m th column of the inverse of the unmixing matrix \mathbf{W} . For the i -th element of the vector $\boldsymbol{\theta}$ found during the initial spike detection step, we form an estimate of the signal subspace $\hat{\Phi}_s$ as the space spanned by the first r left singular vectors of the cropped spatial temporal matrix \hat{X}_c :

$$\hat{X}_c = [\hat{\mathbf{x}}((\theta_i - \tau)), \dots, \hat{\mathbf{x}}((\theta_i + \tau))] \quad (8)$$

The number of samples in the temporal window centered around the maximum of the spike is $2\tau f_s + 1$. The signal subspace is found from the SVD of \hat{X}_c as

$$\hat{X}_c \hat{X}_c^T = [\hat{\Phi}_s, \hat{\Phi}_e] \begin{bmatrix} \Lambda_s & 0 \\ 0 & \Lambda_e \end{bmatrix} [\hat{\Phi}_s, \hat{\Phi}_e]^T \quad (9)$$

where the columns of $\hat{\Phi}_s$ and $\hat{\Phi}_e$ span, respectively, the estimated signal and noise subspaces (Mosher and Leahy, 1999). The dimension of the signal subspace, $r = 4$, was selected to be one or two dimensions larger than the maximum expected number of sources in each spike interval, which in practice rarely exceeded two.

RAP MUSIC finds sources in an automated and recursive fashion by projecting the topography for candidate source locations against the estimated signal subspace. The method is robust to non-dipolar components in the data and can find multiple sources in a single signal subspace (Mosher and Leahy, 1999). Importantly, dipoles will only be found with this procedure if the signal subspace contains a strong dipolar topography. The algorithm is applied to data in the vicinity of each spike in turn. The resulting set of dipoles is then further processed using the clustering procedure described below.

2.2.4 Clustering

Associated with each localized dipole is a location vector $\hat{\boldsymbol{\rho}}_i = [\hat{\rho}_i^x \ \hat{\rho}_i^y \ \hat{\rho}_i^z]$ and its time course $\hat{\mathbf{e}}_i^T = [\hat{e}_i((\theta_i - \tau)), \dots, \hat{e}_i((\theta_i + \tau))]^T$. A single isolated source is of little diagnostic value and consequently we further reduce the number of detected sources through a clustering procedure. Only those sources that fall within one of a number of automatically determined clusters are retained in the final set of detected sources.

To perform clustering we first compute the matrix containing Euclidean distances between each pair of estimated location vectors, i.e. $d_{ij} = |\hat{\boldsymbol{\rho}}_i - \hat{\boldsymbol{\rho}}_j|$. To find the first cluster we chose row i_0 according to the following criterion:

$$i_0 = \arg \max_i \sum_j I(d_{ij} < \delta) \quad (10)$$

where $I()$ is the indicator function and δ is a user specified threshold parameter that determines the maximum radius of the clusters. The first cluster is then the set of dipoles $I_0 = \{j : d_{i_0j} < \delta\}$. The entries of the matrix corresponding to the dipoles from the identified cluster are removed and the procedure repeated until the remaining clusters contain less than N_{min} dipoles, where N_{min} is a user specified parameter.

Each spatial cluster is then subdivided into smaller clusters with similar waveforms of activation using the same procedure as above, but with the Euclidean distance metric applied to the time series rather than locations. This two-step procedure allows us to sort localized spikes into tight spatial clusters with similar temporal properties.

2.2.5 Cluster significance

One further refinement is applied before the procedure is complete. We reduce the number of clusters by removing those that do not exceed a user selected significance threshold. With a large number of sources placed at random within the head, clusters would naturally arise by chance. We therefore test for cluster significance under the null hypothesis of randomly located sources over the head volume. Since the dipoles in each cluster have time series that are similar to each other but distinct from the other clusters, we apply the test separately for each cluster with a Bonferoni correction for multiple hypothesis tests.

To test cluster significance we use as a statistic $\hat{\eta}$, the number of dipoles in the cluster. Under the null hypothesis all of the N_d detected dipoles are uniformly scattered throughout a spherical volume containing the brain. We compute the distribution of the statistic η , under the null hypothesis, equal to the maximum number of dipoles enclosed in a spherical volume with radius r equal to half the distance between the two furthest dipoles in the cluster being tested. The maximum is taken over all possible locations within the spherical brain volume. We assess cluster significance by computing the probability of the statistic η exceeding the observed value $\hat{\eta}$ under the null hypothesis. We derive an approximate distribution for the maximum statistic η under the null hypotheses in Appendix B.

The testing procedure is applied to each cluster in turn. To control for false positives caused by multiple hypothesis tests, we use a Bonferoni correction scheme in which we adjust the significance level to $\alpha = 0.01/N_c$ where N_c is the total number of detected clusters.

3 Simulation Studies

3.1 Simulation Description

To validate the method we performed a set of realistic simulations. Our goal was to obtain ROC (Metz, 1986) to understand the sensitivity and specificity characteristics of our approach, and to investigate whether each of the consecutive steps described above delivered improved accuracy.

We simulated spatio-temporal vectors of spontaneous interictal activity on the basis of real subject data. As described in equation (1), the spatial-temporal vector $\mathbf{x}(t)$ is comprised of three additive parts: focal interictal activity with independent topographies \mathbf{a}_k , spatially coherent background brain activity with topographies \mathbf{b}_l and spatially non-coherent sensor noise $\mathbf{n}(t)$. We simulated each of those three components separately.

We simulated interictal spikes on the basis of averaged waveforms of spikes detected in MEG data collected from a patient with temporal lobe epilepsy using the acquisition system described in Section 4.1. The time course of the averaged spike $s(t)$ is shown in Fig. 2a. The source location of the simulated interictal spike was chosen to coincide with that determined from the same subject dataset and to lie in the right temporal lobe, Fig. 2b. The spatial topography \mathbf{a}_{int} associated with this source was calculated using a spherical head model; this spherical model was also used in all subsequent processing.


We generated a pseudo-random sequence of interictal spikes using a 2-state (binary) Markov chain sequence $f(t)$. The activity at the sensors due to the spike generator was then computed as:

$$\mathbf{x}_{int}(t) = \mathbf{a}_{int} (s(t) * f(t)) \quad (11)$$


where $*$ denotes the convolution operator.

Spike detection in simulated spatially and temporally incoherent noise is relatively straightforward and does not reflect the difficulty of the true problem. Consequently, to perform a more realistic evaluation of our approach under known conditions, we used a more realistic model for background brain activity by training the MAR model described in Appendix A on spontaneous activity recorded from a normal subject at rest. Details of the training procedure and related statistical tests are also contained in the Appendix. The resulting spatially coherent background noise $\mathbf{x}_{bckg}(t)$ sampled from the MAR model is the second additive component in our data model.

The third component of our data model, spatially incoherent sensor noise



a)



b)

Fig. 2. Location and waveform of interictal spike used in computer simulation studies: a) Averaged MEG spike waveform from subject with temporal lobe epilepsy, b) Location of the spike generator for the simulated data.

$\mathbf{n}(t)$, was simulated as Gaussian white noise filtered with a low-pass filter ($f_c = 40Hz$). The simulated data were then computed as a sum of the three components described above:

$$\tilde{\mathbf{x}}(t) = \mathbf{x}_{int}(t) + \mathbf{x}_{bckg}(t) + \mathbf{n}(t). \quad (12)$$

3.2 Simulation Results

To evaluate our approach we applied spike detection and localization method described above to the simulated data and computed the true positive rate (i.e. the fraction of true spikes detected) vs. false positive rate (i.e. the rate of detected events with peaks further than 100 *ms* from the nearest true spike). We plot these performance values as ROC (receiver operating characteristic) curves (Metz, 1986). We evaluated the performance of our approach at three distinct stages:

1. ICA + spike detection based on the spikyness spectrum and thresholding
2. Spike detection based on dipole localization for each spike detected in Step 1
3. Clustering and significance testing of spikes detected and localized as dipoles in Step 2

To generate different operating points in Stage 1 we varied the threshold used to locate spikes from equation (6). To generate different points in Stages 2 and 3 we kept the spike threshold constant at a value which achieved better than .98 true positive rate in Stage 1 and varied the threshold for the RAP MUSIC method in the fit of the dipole model to the signal subspace. As this threshold is dropped, sensitivity will increase at the expense of specificity.

Three ROC curves corresponding to the three stages are shown in Fig. 3. The true positive rate is computed as the fraction of true spikes that are detected. The false positive rate is given as both the probability of a false positive at each sample point and also as the rate of false positives per true spike. The ideal ROC curve would show 100% sensitivity at the lowest possible false positive rate, i.e. the ideal curve rises rapidly towards a true positive fraction approaching unity. The results in Fig. 3 indicate that each stage of refinement substantially improves specificity without affecting the achievable sensitivity. The false positive rate at the end of Stage 3, expressed relative to the true spike rate, shows a worse case of approximately 30% chance of a false positive per true spike with a true positive rate exceeding 98%.

In this study we have attempted to perform a realistic simulation, but whether performance after Stages 2 and 3 for real data will approach the accuracy indicated in these ROC curves will require application to a wide range of clinical data sets. In Section 4 we report on our preliminary studies with clinical MEG data.




Fig. 3. Receiver Operating Characteristic (ROC) curves showing performance at three stages in the detection algorithm. Plotted for each stage of the processing are the true positive rate vs. false positive rate as described in the text. The results show clear improvement in sensitivity vs. specificity performance from Stage 1 to 2 and from Stage 2 to 3.

4 Results

We applied the spike detection method described above to spontaneous MEG data sets collected from four different subjects. The four patients had complicated clinical presentations and required invasive electrode recordings for localization of the seizure origin for surgical planning. Their intractable complex partial seizures could not be localized by standard non-invasive Video-EEG and imaging criteria. In each case there was a clinical hypothesis which justified proceeding to invasive electrodes due to the likelihood of finding an operable seizure focus.

Subject	Age (Gender)	Phase I	MEG	Phase II	Phase III	Phase IV (mos)	Automatic detection outcome
A	36M	LF	LTP	LSMA ^d , LH ^g	no	NA	LT[58],LP[33], C[73,66,78,56,31]
B	17M	RF	RF,RT	RF ^g	RF, RT	1(32)	RF[57,71],RT[82]
C	34F	RT	RTlat	RTlat ^g	RTant	1(18)	RTlat[26]
D	32M	LT,RH	RP, RT	RT ^g ,RT ^d	RT	3(9)	RTpost[35,35,51], RTmes[46]

Table 1

Summary of data for the 4 clinical subjects for which MEG data was analyzed in this study. A description of each column follows. Phase I: Scalp and sphenoidal inpatient Video-EEG telemetry; MEG: Magnetoencephalography and magnetic source imaging(MSI); Phase II: Intracranial monitoring with depth (d) or grid (g) electrodes; Phase III: Focal excisional surgery; Phase IV: Outcome classification with number of follow-up months (mos), follow-up outcome class: 1=seizure-free, 3=90 % seizure reduction; Automated detection outcome: Location of dipole clusters [number of spikes in each cluster]. Abbreviations: L=left, R=right, T=temporal, H=hippocampus, P = Parietal, F=frontal, SMA = somatosensory area, C=deep white matter, ant = anterior, lat=lateral, post=posterior, mes=mesial

4.1 MEG data acquisition

MEG was recorded using a whole cortex CTF neuromagnetometer (CTF Corporation, Port Coquitlam, Canada) in a magnetically shielded chamber (VakuumSchmelze, Hanau, Germany). The magnetometer has 68 sensor channels (1st derivative co-axial gradiometers, coil diameter 1.98 cm, baseline 5.0 cm) and 32 reference channels. The data sample rate was 250 Hz per channel with a band-pass of 1-70 Hz, a dynamic range of 16 bits and system noise of 5 to 7 fT/\sqrt{Hz} . For data analysis, we computed an approximate third order spatial gradient using the reference channels. The system accuracy was tested with a dry calibration phantom and localization error for a single dipole measured at $1.0 \pm 0.4mm$.

4.2 Comparison to the four phase clinical procedure

There are four parameters to be set in the automated procedure: the threshold for initial spike detection in the ICA time series, the threshold to use in dipole localization in RAP MUSIC, and thresholds for determining cluster size and significance. After adapting these parameters during initial studies with human data, we were able to process all four clinical data sets with the same

parameter settings. We compared the result of our automatic method to that of the four phases (Phase I - Phase IV) of the clinical procedure summarized in Table 1. For each patient the table gives abbreviated loci of the detected foci of abnormal activity during Phase I (excluding MEG) and of results obtained from MEG using standard analysis, i.e. visual selection of spikes followed by single dipole localization. The Phase II column reports locations of foci obtained from intracranial monitoring. The column corresponding to Phase III (resection) indicates the location of surgical resection if performed. Phase IV summarizes the results of post-surgical monitoring of the patient. The table shows the outcome class of each patient and number of months that have passed since surgery. The results of the automatic procedure are summarized in the right-most column. The table shows that regions of abnormal activity detected by the automatic procedure were similar to those on which the four phase clinical procedure were focused. The brain region in which resection was performed coincided with the location of one of the clusters found by our automatic procedure in all three cases in which a resection was performed.

4.3 Graphical comparison to the results of Phase I.

One of the steps in Phase I of the clinical procedure is manual analysis of spontaneous interictal MEG recordings. The goal of this step is to detect abnormal spike activity and determine its spatial origin. In Figs. 4, 5, 6, 7 we graphically compare the results of manual detection performed by a qualified examiner to that obtained by our method.




Fig. 4. Automated detection and clustering results for subject A, compared to the manual detection procedure. Results of the automated procedure are shown as circles with radii equal to 3 standard deviations of the distance of dipoles from the centroid for all sources in the cluster. Circles with thicker lines indicate probable epileptogenic clusters while thinner lines indicate non-epileptogenic clusters as determined by their location and averaged time courses. The number of spikes in each cluster is shown on the coronal sections. Manual detection results are shown as bright square dots.

The number of clusters for the four subjects using our procedure was A: 7, B: 3, C: 1, and D: 4. Visual inspection of the temporal properties of the cluster time-courses allowed us to attribute clusters localized in white matter regions in the center of the head in subject A to the normal sleep activity spread over the entire cortex. These clusters were excluded from further consideration. In Fig. (8) we show two clusters for subject C. Fig. (8)a shows distinct spike-like behavior in the temporal lobe. The second example, Fig. (8)b, also shows tight spatial clustering but inspection of the time series shows little structure so that this cluster, which is possibly due to muscle artifacts, can be rejected as a potentially epileptogenic region. In fact this cluster was rejected by our automated procedure because of the large variability in the time series, and was included here only to demonstrate the importance of considering both the spatial and temporal characteristics of putative clusters of localized sources.




Fig. 5. Automated and manual detection and clustering results for subject B. Details as for Fig. 4.

To quantitatively compare human observer performance with our automated procedure we performed a further analysis of this clinical data. For different values of the RAP-MUSIC correlation threshold we computed the proportion of spikes found by a qualified examiner which were also found by our automatic procedure. We counted spikes as matching if the distance between the time markers for the human and automated detector were less than 120 ms. The results of this analysis are shown in Fig. 9. For a typical RAP MUSIC threshold of 95%, the automated procedure found between 48% and 93% of the spikes found by the human examiner. It should be noted however, that the automated procedure finds far more sources, even after clustering and significance testing of the clusters, than is practical for a human examiner. As an example, in one 10 min record of spontaneous MEG data from a patient with partial epilepsy consisting of 150,000 samples and 68 channels, the manual examiner selected 13 spikes. The initial spike detection (Stage 1 in Section 3.2) produced 1342 spikes, including 12 of the 13 manually detected spikes. Stage 2 reduced this number to 376 dipoles keeping 11 out of 13 of the manually detected. Keeping only statistically significant clusters resulted in 91 dipoles split into two clusters. As shown in the earlier figures, many of these sources clustered into regions consistent with the foci of interictal epileptic activity. This indicates the potential power of the automated procedure, not only in reducing time required for manual analysis of data, but in finding much larger sets of plausible focal sources that can potentially be used by physicians in treatment planning.




Fig. 6. Automated and manual detection and clustering results for subject C. Details as for Fig. 4.

5 Discussion and conclusions

We have described a procedure for detection of significant clusters of dipolar sources with spike-like time courses in interictal MEG data. The method first finds candidate spikes in the MEG data, then attempts to localize current dipoles to account for each spike and retains only the sources which conform to this model. These dipoles are then clustered according to their spatial and temporal characteristics and finally the clusters tested for statistical significance to yield a set of significant clusters as candidate epileptogenic regions.

Once four user specified parameters are chosen, the procedure described above is a fully automated method for screening of spontaneous MEG data for interictal events. It supplies the physician with a more complete and potentially more meaningful dataset of source foci and time series than is feasible using a manual spike detection protocol.

Our simulation studies were based on a realistic model for coherent background noise with parameters determined from a normal subject at rest. ROC studies of the method were based on dipolar spike-like sources embedded in this coherent background activity. These studies show the benefit of the multistage





Fig. 7. Automated and manual detection and clustering results for subject D. Details as for Fig. 4.


refinement of the initially detected spikes in terms of substantially improved specificity with near constant maximum sensitivity.

In practice, the method was found to produce clusters in MEG data from patients with partial focal epilepsy who were candidates for surgical resection. In the 3 of 4 cases in which a resection was performed, the automated procedure found a significant cluster in the vicinity of the area that was subsequently resected. In these real cases, some spurious clusters remained, even after application of the significance test. However, most of these can be easily rejected by a neurologist since their centroids fall in deep white matter and the time series are inconsistent with epileptogenic activity. Sharp transients of sleep might potentially cause confusion, but these are usually either low amplitude or mid line in location and therefore likely to be excluded. Occipital partial epilepsy is uncommon and usually obvious clinically, thus the normal sharp transients from the occipital lobe should not cause a significant clinical problem.

In summary, the above procedure is promising and likely to be useful to the physician as a more sensitive, automated and objective method to help in the localization of the interictal spike zone of intractable partial seizures. The final



a)



b)

Fig. 8. Location in three orthogonal views of two apparent dipole clusters found for subject C and their corresponding time series. The bold line in the time series graph is the average time course over all dipoles. (a) This cluster in the temporal lobe shows clear spike activity in the time courses and represents a potentially epileptogenic region; (b) this cluster, while also spatially tight, shows a high degree of variability in the time courses and little evidence of spike activity and is believed to be due to muscle artifacts in the data. This apparent cluster was discarded by our automated procedure due to the variability in the time courses and is included here as a demonstration of the importance of considering both the spatial and temporal characteristics of detected sources.

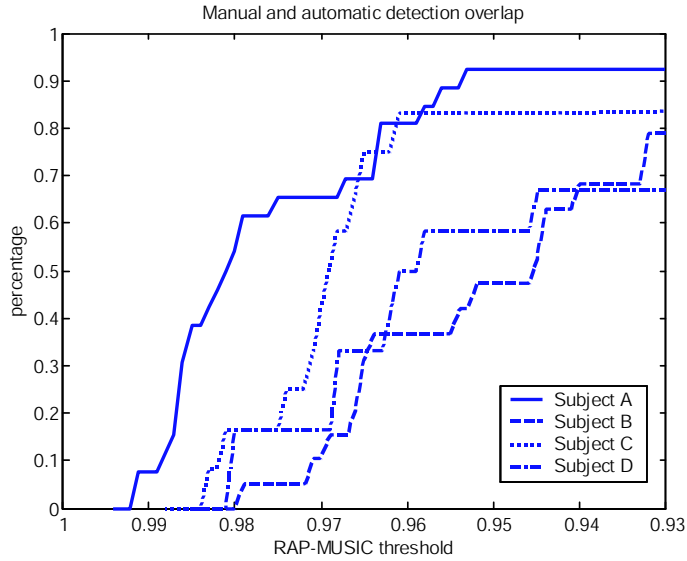


Fig. 9. Comparison of manual and automatic spike detection procedures. For each of the 4 subjects, we plot the fraction of the manually detected dipolar sources that are also found by the automated procedure. Different points on the curve are found using different acceptance thresholds during RAP MUSIC dipole localization. For the 4 subjects the following number of dipoles were found by the qualified reader: A: 13, B: 26, C: 12, D: 12

output can be visually verified by neurologists in terms of both the location and distribution of the dipole clusters and their associated time series. Due to the clinical relevance and demonstrated promise of this method, further investigation of this approach is warranted.

Appendix

A Multichannel Autoregressive (MAR) Model of Spontaneous Brain Activity

The MAR model (Franaszczuk and Bergey, 1999) is an extension of the autoregressive model characterized by a series of coefficient matrices A_k $k = 1, \dots, P$, where P is the order of the MAR model. The $(m, n)^{th}$ element of A_k represents the interaction from the n^{th} to the m^{th} source at time lag k . Formally, the time series generated by the model can be written as

$$\mathbf{p}(t) = - \sum_{k=1}^P A_k \mathbf{p}(t-k) + \boldsymbol{\epsilon}(t) \quad (\text{A.1})$$

where $\boldsymbol{\epsilon}(t)$ is a spatially and temporally white vector representing the "innovations" process that drives the model.

The MAR model can be written in terms of a matrix transfer function, i.e. in the z -domain:

$$\mathbf{P}(z) = (\mathbf{H}(z))^{-1} \mathbf{E}(z) \quad (\text{A.2})$$

where the matrix transfer function $\mathbf{H}(z) = \{h_{ij}(z)\}$ is defined as

$$\mathbf{H}(z) = 1 + \sum_{k=1}^P A_k z^k \quad (\text{A.3})$$

The degree of interaction between the sources of background activity at different frequencies can be computed using the directed transfer function (DTF) matrix $\Gamma = \{\gamma_{ij}^2\}$ (Kaminski and Blinowska, 1991):

$$\gamma_{ij}^2(f) = \frac{|h_{ij}(f)|^2}{\sum_{k=1}^L |h_{ik}(f)|^2} \quad (\text{A.4})$$

where $|h_{ij}(f)|^2$ is the cross spectral density between the i th and j th sources.

A matrix of scalar parameters reflecting interaction for a frequency band of interest Δ can be computed by integrating over that frequency band:

$$g_{ij} = \frac{\int_{\Delta} |h_{ij}(f)|^2}{\int_{\Delta} \sum_{k=1}^L |h_{ik}(f)|^2} \quad (\text{A.5})$$

The matrix $G = \{g_{ij}\}$ is called integrated directed transfer function (IDTF) matrix for the MAR model (Kaminski and Blinowska, 1991).

Here we use the MAR model to represent coherent background brain activity. We now describe how we trained the model using spontaneous MEG data from a normal subject at rest. We define a coherent background activity subspace as that spanned by the dominant left singular vectors of the SVD of the spatial-temporal data matrix \mathbf{X} collected from the subject using the setup described in 4.1 at the sample rate of 250 Hz. We retained only those singular vectors such that the subspace spanned by them captures 95% of the variance in the data; in the data used in this simulation, this corresponded to the first $r = 48$ singular vectors. From these vectors we constructed the matrix:

$$\mathbf{Y}_{bckg} = \begin{bmatrix} v_1^T \\ \vdots \\ v_r^T \end{bmatrix}. \quad (\text{A.6})$$

where v_i is the i th right singular vector of \mathbf{X} .


To simulate temporal activity in this subspace we use the MAR model described in equation (A.1). Parameters for MAR models of different orders were estimated from \mathbf{Y}_{bckg} by solving a system of normal equations (Kay, 1988). The estimated models were then cross-validated on independent segments of data (Golub et al., 1979). Using the estimated coefficient matrices for each model order, we computed the total prediction error across all channels, i.e. the trace of the prediction error correlation matrix. The prediction errors plotted in Fig. A.10 indicate an appropriate model order of $P = 4$. The IDTF matrix, equation (A.5), is shown in Fig. A.1b. This matrix indicates a strong degree of interaction between the different components in the model since if each component were independent, the IDTF matrix would be diagonal.

The use of the normal equations to solve for MAR parameters is optimal in a maximum likelihood sense only for a Gaussian innovation process $\boldsymbol{\epsilon}(t)$. To investigate the normality of the residuals we used the Kolmogorov-Smirnov test. The p-values for this test for each channel are shown in Fig. A.2b. The plot shows that the null hypothesis of normality cannot be rejected at the significance level $\alpha = 0.05$. We can therefore use normally distributed noise to drive our MAR model in order to simulate realistic spatially coherent background brain activity.


The MAR model and coherent background signal subspace were combined to generate coherent background activity according to:

$$\mathbf{x}_{bckg}(t) = \begin{bmatrix} \mathbf{u}_1, \dots, \mathbf{u}_r \end{bmatrix} \mathbf{y}_{bckg}(t) \quad (\text{A.7})$$

where \mathbf{u}_r are the first $r = 48$ left (spatial) singular vectors of the spatial-



a)



b)

Fig. A.1. a) Prediction Error variance vs. MAR model order for estimated parameters when applied to cross-validation data. Note that the cross validation prediction error indicates a distinct minimum at model order $P=4$; b) The Integrated Directed Transfer Matrix for the estimated 4th order MAR parameters; the figure shows a high degree of correlation between channels in the data.

temporal data matrix X and $\mathbf{y}_{bckg}(t)$ is the MAR process generated using the estimated MAR parameters and driven by white Gaussian noise.




Fig. A.2. Computed p-values for Kolmogorov-Smirnov test for normality applied to residual or prediction errors using cross-validation data. The results indicate that there is no evidence in this data to support use of a distribution other than Gaussian to drive the MAR models.

B Cluster Significance Computation

We describe the statistic used to test the null hypothesis as described in Section 2.2.5 that each dipole cluster arose by chance through a uniform random distribution of the dipoles within the brain volume. In order to compute an analytical expression for the significance level of an observation under the null hypothesis we divide the spherical head volume V into a 3-D grid with cubic voxels of small size $\delta \times \delta \times \delta$, yielding the total number of voxels $N_v = \lfloor \frac{V}{\delta^3} \rfloor$. For a particular placement of dipoles we then compute the number of dipoles in each such voxel, i.e.

$$N_i = \sum_{j=1}^{N_d} I_i(j), \quad i = 1, \dots, N_v, \quad j = 1, \dots, N_d \quad (\text{B.1})$$

where N_d is total number of detected dipoles and $I_i(j)$ is an indicator function equal to 1 if the j -th dipole falls in the i -th voxel.

We now compute the probability of observing more than $\hat{\eta}$ dipoles within a spherical patch of radius r under the null hypothesis. The probability of observing $\eta \leq \eta_0$ dipoles within a spherical patch of radius r located in a fixed position can be computed using the binomial cumulative distribution function

as

$$P(\eta \leq \eta_0) = \sum_{t=0}^{\eta_0} \binom{N_d}{t} \left(\frac{m}{N_v}\right)^t \left(1 - \frac{m}{N_v}\right)^{N_d-t} \quad (\text{B.2})$$

where $m = \frac{4}{3}\pi \left(\frac{r}{\delta}\right)^3$ is the number of voxels within the spherical volume.

To compute the maximum distribution over all possible locations of the patch we make the approximating assumption that the number of dipoles within the spherical volume is independent for different locations of this volume when the allowed sphere centers are spaced by distance r from each other. Under this assumption the significance level α for the maximum statistic can be computed as

$$\alpha = P(\max(\eta_1, \dots, \eta_{N_p}) > \hat{\eta}) = 1 - \prod_{i=1}^{N_p} p(\eta_i \leq \hat{\eta}) = 1 - P(\eta_i \leq \hat{\eta})^{N_p} \quad (\text{B.3})$$

where η_i is the number of dipoles within the spherical volume in the i -th position and $N_p = \frac{4}{3}\pi \left(\frac{R_h-r}{r}\right)^3$ is the number of possible locations of the patch within the sphere of radius R_h .

To assess the accuracy of the independence approximation we performed Monte Carlo simulations in which we compute the number of dipoles required to achieve a given significance level α for different cluster sizes. In Fig. B.1 we show results of computation using equation (B.3) for different levels of significance α . We also show the number of dipoles required to achieve $\alpha = 0.01$ based on 100 Monte Carlo trials. These results show that equation (B.3) provides a mildly conservative number compared to Monte Carlo trials and therefore can be used to control false positives with reasonable accuracy.

References

- Amari, S. I. (1998). Natural gradient works efficiently in learning. *Neural Computation*, 10 2:251–276.
- Bell, A. and Sejnowski, T. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7:1129–1159.
- Cardoso, J. (1989). Source separation using higher order moments. In *Proceedings of ICASSP*, pages 2109–2112.
- Cardoso, J. (1998). Multidimensional independent component analysis. In *Proceedings of ICASSP*.
- Diambra, L. and Malta, C. (1999). Nonlinear models for epileptic spikes. *Physical Review E*, 59:929–937.
- Faure, C. (1985). Attributed strings for recognition of epileptic transients in eeg. *Int. J. Biomed. Comput.*, 16:217–229.




Fig. B.1. Number of dipoles required to achieve a given significance level α as a function of cluster size for a fixed brain volume. The values are computed for total of $N_d = 1000$ dipoles using the approximate formula (B.3) for $\alpha = 0.01, 0.001$ and 0.0001 and using the Monte Carlo method for $\alpha = 0.01$.

- Flanagan, D., Agarwal, R., and Gotman, J. (2002). Computer-aided spatial classification of epileptic spikes. *Journal of Clinical Neurophysiology*, 19:125–135.
- Franaszczuk, P. and Bergey, G. (1999). An autoregressive method for the measurement of synchronization of interictal and ictal EEG signals. *Biol. Cybern.*, 81(1):3–9.
- Glover, J., Raghavan, N., Ktonas, P., and Jr., J. F. (1989). Context-based automated detection of epileptogenic sharp transients in the eeg: elimination of false positives. *IEEE Trans Biomed Eng*, 36:519–527.
- Golub, G., Heath, M., and Wahba, G. (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223.
- Gotman, J. (1993). *J Jr. Engel ed., Surgical treatment of the epilepsies*, chapter Computer Applications, pages 429–444. New York: Raven Press.
- Gotman, J. (1999). Automatic detection of seizures and spikes. *Journal of Clinical Neurophysiology*, 16:130–140.
- Gotman, J. and Gloor, P. (1976). Automatic recognition and quantification of interictal epileptic activity in the human scalp eeg. *Electroenceph. Clinical Neurophysiology*, 41:513–529.
- Hyvarinen, A. and Oja, E. (1997). A fast fixed-point algorithm for independent component analysis. *Neural Computation*, 9(7):1483–1492.
- Kaminski, M. and Blinowska, K. (1991). A new method of the description of the information flow in the brain structures. *Biological Cybernetics*, 65:203–210.
- Kay, S. (1988). *Modern Spectral Analysis*, chapter Computer Applications. Prentice Hall, Englewood Cliffs.

- Knuth, K. (99). A bayesian approach to source separation. In *Proceedings of the First International Workshop on Independent Component Analysis and Signal Separation: ICA99*, pages 283–288.
- Kobayashi, K., Akiyama, T., and Nakahori, T. (2002a). Systematic source estimation of spikes by a combination of independent component analysis and RAP-MUSIC, I: Principles and simulation study. *Clinical Neurophysiology*, 113:713–724.
- Kobayashi, K., Akiyama, T., and Nakahori, T. (2002b). Systematic source estimation of spikes by a combination of independent component analysis and RAP-MUSIC, II: Preliminary clinical application. *Clinical Neurophysiology*, 113:725–734.
- Kobayashi, K., James, C., Nakahori, T., Akiyama, T., and Gotman, J. (1999). Isolation of epileptiform discharges from unaveraged eeg by independent component analysis. *Clinical Neurophysiology*, 110(10):1755–63.
- Metz, C. (1986). ROC methodology in radiologic imaging. *Invest. Radiol.*, 21:720–733–276.
- Mosher, J. and Leahy, R. (1998). Recursive MUSIC: a framework for EEG and MEG source localization. *IEEE Trans. on Biomed. Eng.*, 45 11:1342–1355.
- Mosher, J. and Leahy, R. (Feb. 1999). Source localization using recursively applied and projected (RAP) music. *IEEE Trans. on Signal Processing*, 47 2:332–40.
- Ossadtchi, A., Mosher, J., Jerbi, K., Lopez, N., Baillet, S., Sutherling, W., and Leahy, R. (2001). Automated interictal spike detection and source localization in meg using independent component analysis ICA. In *American Epilepsy Society 55th Annual Meeting, Philadelphia*.
- Ossadtchi, A., Mosher, J., Lopez, N., Baillet, S., Sutherling, W., and Leahy, R. (2002). Automated detection of dipole clusters in interictal MEG data. In *13th International Conference on Biomagnetism, Jena, Germany*.
- Robinson, S., Virba, J., and Ishii, R. (2002). Finding epileptic loci by nonlinear parametrization of source waveforms. In *13th International Conference on Biomagnetism, Jena, Germany*.
- Therrien, C. (1992). *Discrete random signals and statistical signal processing*. Prentice Hall, Englewood Cliffs, NJ.
- Wilson, S. and Emerson, R. (2002). Spike detection: a review and comparison of algorithms. *Clinical Neurophysiology*, 113:1873–1883.
- Zhukov, L., Weinstein, D., and Johnson, C. (2000). Independent component analysis for eeg source localization. an algorithm that reduces complexity of localizing multiple neural sources. *IEEE Engineering in Medicine and Biology*, 19:87–96.