Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2015 April 16; 2015: 1053–1056. doi:10.1109/ISBI.2015.7164052.

A STATISTICAL APPROACH TO INCORPORATE MULTIPLE ECG OR EEG RECORDINGS WITH ARTIFACTUAL VARIABILITY INTO INVERSE SOLUTIONS

J. Coll-Font¹, B. Erem², P. Štůví ek³, and D.H. Brooks¹

¹B-spiral group, Northeastern University, Boston (MA), USA

²Computational Radiology Lab, Boston Children's Hospital, Boston (MA), USA

³Charles University Hospital, Prague, Czech Republic

Abstract

Inverse methods for localization and characterization of cardiac and brain sources from ECG and EEG signals are notoriously ill-conditioned and thus sensitive to SNR in the measurements. Multiple recordings of the same underlying phenomenon are often available, but are contaminated by unmodeled correlated noise such as heart motion from respiration or superposition of atrial activation or on-going EEG in the case of inter-ictal spikes or evoked response in EEG. We address here the open question of how best to incorporate these multiple recordings, comparing standard ensemble averaging, a multichannel non-linear spline-based average designed to be less sensitive to timing variations from motion or modulation, and a probabilistic inverse incorporating a data-driven model of the noise correlation and using all recordings jointly. Results are tested on localizations of clinically recorded 120 lead ECGs during ventricular pacing.

Index Terms

Inverse ECG and EEG; ensemble average

1. INTRODUCTION

Inverse imaging in electroencephalography (EEG) or electrocardiography (ECG) attempts to reconstruct internal — hidden — electrical processes of the body using surface measurements — the scalp in the case of EEG, the torso for ECG. These non-invasive methods have are of both clinical and scientific interest, for example diagnosis and guidance of surgical interventions in epilepsy or cardiac ablations, or enhanced estimation of normal or abnormal function with evoked potentials or cardiac studies. In both EEG and ECG it is frequently possible to record multiple instances, or epochs, of the same putative underlying phenomenon, such as multiple inter-ictal spikes or multiple arrhythmic beats during pre-surgery studies. Because the inverse reconstruction of brain or cardiac potentials are both highly ill-posed problems, and thus sensitive to small perturbations due to noise, which can have tremendous impact on the estimated solutions, multiple epochs offer the promise of

reduction of noise effects, and perhaps less *a priori* regularization, leading to solutions which may be both more accurate and sufficiently stable.

However, it is not obvious how it to best take advantage of these multiple recordings. Standard approaches either pick the "best" single recordings to invert, reconstruct the hidden potentials using each epoch individually and then ensemble average the individual solutions, or ensemble average the recorded epochs and then apply the inverse method to the averaged recording. Because of the the sensitivity of reconstructions to the SNR of the "input" to the inverse procedure, and the need for regularization tailored to that SNR, these methods do not typically lead to equivalent results. More fundamentally, all these methods imply the assumption that the epochs are independent realizations of a stochastic process consisting of the same underlying signal contaminated by noise that is spatially and temporally uncorrelated. However in practice there are significantly correlated noise sources, in effect interference in the signal, stemming from uncontrolled "background" phenomena such as on-going brain activity in the EEG or spatial effects of respiration and atrial electrical activity in recordings during arrhythmic heartbeats.

As one attempt to overcome this interference, our group recently developed a non-linear averaging method for ECG that is more robust to this background noise [1]. This method only uses the geometric properties of the measurements in the high-dimensional measurements space, while ignoring their time stamps, to recover the denoised signal of interest, by approximating it with a Cartesian-product B-spline model [2].

From a statistical point of view, though, an optimal approach, if a statistical model of the interference is available, would be to directly use all the recordings jointly in the inverse estimate, accounting for the background noise and variability across epochs. A natural way of describing these type of problems is to use a probabilistic characterization of the signals of interest and the noise activity. Given this probabilistic model, it is possible to estimate the hidden potentials through Bayesian approaches or simply maximization of the posterior probability of the hidden potentials. In the simplest setting, under Gaussian noise assumptions, these methods are identical and are equivalent to Tikhonov regularization applied to the ensemble average of the input signal [3]. However when the interfering "noise" has structure, as it does in both these settings, a more sophisticated statistical model may lead to a new regularized solution method. Here we present a statistical model of the interfering effects of respiration and uncontrolled atrial activity, developed in the context of electrocardiography body surface measurements during artificially-induced arrhythmic beats (using a intracardiac catheter device), and compare inverse results to those obtained both with ensemble averaging and our non-linear spline average, extending an earlier study in which we simply compared the averaging methods alone [4].

2. METHODS

We start with the assumption that the forward problem describing the relationship between unknown sources and surface measurements has been solved, resulting in the relationship $y_t = Ax_t$, where y_t and x_t are the surface and hidden signals at time t and A encodes the forward solution. In this section we will discuss the two approaches to be compared: average first,

then "invert", or reconstruct with all the data at once. To do so, we will first introduce the pipeline for the averaging techniques and then we will describe the combined approach.

2.1. Pre-Averaging methods

Averaging before applying inverse methods relies on the assumption that the pre-processing step will increase SNR and less regularization will be needed in the inverse to obtain equivalent stability. The standard approach, ensemble average the measured potentials across epochs, is optimal only in the case when the realizations are only contaminated with AWGN. Another approach to averaging is to fit a more complex model to the potentials; here we use a spline-based model. In previous work, Erem *et al.* observed that the measured potentials can be represented as a curve embedded in a high dimensional vector space defined by the potentials at each electrode [5].¹ However, biological noise, such as respiration, alters the trajectory of the potentials such that samples at equivalent time instances of different epochs appear at different positions along the curve. Thus, any denoising method that relies on time stamps of the signal to average will be affected by this variability. To overcome this problem, our group developed a spline averaging method, which fits a Bspline model that is defined over a pseudo-time parameter that is independent of the time stamps of the samples. This model represents the measured signal at epoch r and time t ($y_{r,t}$) as:

$$y_{r,t} = \sum_{i=1}^K \mathbf{k}_{y_i} \mathbf{b}_i^T(\theta_{r,t}) + \mathbf{w}$$

$\theta_{r,t}$ =pseudo- time.

$\mathbf{k}_{y_i}$ =knot points. (1)

K =number of knot points.

$\mathbf{b}_i$ =B- spline basis.

$\mathbf{w}$ =AWGN.

The spline averaging optimizes over two sets of parameters, the knot points $\mathbf{k}_{y_i}$, which do not depend on the epoch and thus describe a "spline average" curve, and the pseudo-time $\theta_{r,t}$, by solving:

¹More recent work (as yet unpublished) has indicated that this variability is predominately due to both respiration and atrial activity.

$$\min_{\mathbf{k}, \theta} \sum_{r=1}^R \left\| \mathbf{y}_{r,t} - \sum_{i=1}^K \mathbf{k}_{xi} \mathbf{b}_i^T(\theta_{r,t}) \right\|_2^2 \quad (2)$$

by recursively solving for the knot points $\mathbf{k}$ and then updating the pseudo-time θ , similar to expectation maximization.

Similar to the inverse method applied to spline fits to recordings from individual beats in [1], since the knot points fully describe the curve and can be seen as potential distributions themselves, we reconstruct only at the knot points and then obtain the whole temporal sequence of the solution using the pseudo-times obtained in the averaging step. With a Tikhonov regularization method we have:

$$\min_{\mathbf{k}_{xi}} \|k_{yi} - Ak_{xi}\|_2^2 + \lambda \|Dk_{xi}\|_2^2 \quad (3)$$

where the regularization term Dk_{xi} is an gradient estimator λ is determined using the L-curve method.

2.2. Probabilistic Inverse

Our general approach here is to develop a probabilistic model from observations of the data to characterize the biological noise in the measurements. In the particular case of ECG, we observed that the spread of the samples around the ensemble average appeared to be approximately Laplacian distributed as shown in Figure 1, with a shape parameter that was consistent across all epochs but with a mean that shifted from one epoch to the next. We postulate that this shift in the mean could be produced by differences in respiration and atrial activity across different epochs. Further observations showed that the time average at each epoch had a non-trivial spatial distribution that could highly impact the inverse estimation if not taken into account. Thus we propose a probabilistic data model based on the Laplace distribution with mean given by the forward model of the potentials ($\mathbf{Ax}$) shared across epoch plus an unknown potential distribution $\mathbf{b}_r$ which is constant in time but varies for different epochs. Since the actual solutions cannot be observed we cannot develop a source signal model with the same approach. However due to the ill-posedness of the inverse problem we need a prior on the sources that is flexible but also regularizes to avoid overfitting. Here we used a prior distribution equivalent to the regularization term in the Tikhonov method applied for the spline average, a 0-mean Gaussian distribution with covariance matrix equal to the inverse volumetric derivative, thus favoring spatially smooth distributions. The resulting probability model is

$$p(x_t) = \frac{1}{(2\pi)^{d/2} |D|^{1/2}} e^{-\frac{1}{2\sigma^2} (x_t^T D x_t)} \quad (4)$$

$$p(y_{r,t} | x_t) = \frac{1}{2\sigma_l} e^{-\frac{1}{\sigma_l} (\|x_t - Ax_t + b_r\|_1)}$$

Thus, the posterior probability of the heart potentials $\mathbf{X}$ given all the measurements $\mathbf{Y}$ is:

$$p(\mathbf{X}|\mathbf{Y}) \propto \prod_{r=1}^R \prod_{t=1}^T p(y_{r,t}|x_t)p(x_t)$$

We find the MAP estimate by minimizing the negative log posterior:

$$\min_{x_t, b_r} L(\mathbf{x}|\mathbf{y}) = \sum_{r=1}^R \sum_{t=1}^T \|y_{r,t} - Ax_t + b_r\|_1 + \frac{\sigma^2}{\sigma_l} \|Dx_t\|_2^2 + ct. \quad (5)$$

To solve for the unknown parameters x_t and b_r , we implemented a tailored alternative directions method of multipliers (ADMM) [6] algorithm. ADMM is an iterative optimization method that sequentially solves for each unknown parameter in the objective function while making use of the augmented Lagrangian to account for constraints. In particular, this MAP problem posed in the ADMM perspective can be written as:

$$\min_{\mathbf{x}_t, \mathbf{b}_r} = \sum_{r=1}^R \sum_{t=1}^T \|z_{r,t}\|_1 + \lambda \|Dx_t\|_2^2 \quad (6)$$

$$st. z_{r,t} = y_{r,t} - Ax_t + b_r$$

where the augmented Lagrangian is:

$$\mathcal{L}(\mathbf{z}_{r,t}, \mathbf{x}_t, \mathbf{b}_r, \mu) = \sum_{r=1}^R \sum_{t=1}^T \|z_{r,t}\|_1 + \lambda \|Dx_t\|_2^2 + \mu^T (z_{r,t} - y_{r,t} + Ax_t - b_r) + \frac{\rho}{2} \|z_{r,t} - y_{r,t} + Ax_t - b_r\|_2^2 \quad (7)$$

This objective function is then optimized over x_t and b_r . The proportion between variances given by λ can be computed with the L-curve method. Thus, the ADMM algorithm is:

Until convergence:

1. $\mathbf{x}_t^{k+1} = \arg\min_{x_t} \mathcal{L}(\mathbf{z}_{r,t}^k, \mathbf{x}_t, \mathbf{b}_r^k, \mu^k)$
2. $\mathbf{z}_{r,t}^{k+1} = \arg\min_{z_{r,t}} \mathcal{L}(\mathbf{z}_{r,t}, \mathbf{x}_t^{k+1}, \mathbf{b}_r^k, \mu^k)$
3. $\mathbf{b}_r^{k+1} = \arg\min_{b_r} \mathcal{L}(\mathbf{z}_{r,t}^{k+1}, \mathbf{x}_t^{k+1}, \mathbf{b}_r, \mu^k)$
4. $\mu^{k+1} = \mu^k + (z_{r,t}^k - y_{r,t}^k + Ax_t^k - b_r^k)$

3. EXPERIMENTS AND RESULTS

3.1. Dataset and Experiments

As an initial comparison of the two distinct ways of approaching the problem we tested it using ECG data from a pacing experiment previously presented in [1]. The dataset were

120-electrode body surface ECG recordings obtained during ventricular pacing of three consenting volunteers with healthy ventricles in the Charles University Hospital of Prague in the Czech Republic in conjunction to clinical atrial ablations. The ventricles were paced multiple times at several locations using a catheter device whose position was recorded using the CARTO mapping system. The forward matrix was obtained based on X-Ray CT scans around the heart while the torso geometry was obtained by fitting a generic geometry. Details are in [1].

This data was then used to estimate the underlying heart potentials with 3 different variations of the inverse problem explained in Section 2: Ensemble average and spline averaging, both followed by an inverse solution, and the proposed probabilistic inverse using all epochs at once. We compared the estimated position earliest activated node in each solution to the CARTO-recorded pacing site coordinates as the Euclidean distance (in mm) between these two positions.

3.2. Results

Figure 2 shows the histogram of the error of each inverse method for all subjects and pacing locations. No clear systematic difference between the three methods is apparent; all peak around 32mm of error and have long tails to 80mm. In order to attempt to better observe the differences among the methods, in Figure 3 we show histograms of the error difference between the averaging methods (spline and ensemble) versus the probabilistic inverse, that is, $Error_{spl} = Error_{spline} - Error_{prob.}$ and $Error_{ens} = Error_{ensemble} - Error_{prob.}$. Thus negative values indicate worse performance of the probabilistic inverse with respect to the averaging based techniques, while positive values indicate improvement. Here we observe some differences across both methods and subjects. All subjects show a range of results, from more than 20mm improvement with the probabilistic method to equivalent loss in performance. However, for the first subject, the averaging methods tend to do better than the probabilistic model, while in the second this trend is reversed and in the third they do not show much difference. We also tried to segregate results according to the anatomical site of the pacing (left ventricular wall, right ventricular wall, and septum), to investigate if there was any systematic relationship, where the anatomical site of the pacing was categorized as in [1]. However while we did see some consistent differentiation within a given subject when results were categorized by pacing site location, again the results were inconsistent across subjects. For example, for the first subject the left ventricular wall pacings were less accurate with the probabilistic approach, while accuracy on the septal sites was generally the same, while for the second subject there was considerable improvement on the left ventricular wall while again results on the septum were similar.

4. DISCUSSION

At this point, with data available from only three subjects, we cannot reach a definitive conclusion about which approach is preferred. It may well be that we can learn to differentiate among datasets where one or the other method is likely to achieve the best results. However, we point out that the initial probabilistic model designed for this inverse ECG experiments was relatively simple. We believe there is considerable room for better

characterization of the noise. Finally the model presented here considers only spatial variation, and we are currently working to incorporate a temporal model similar to the spline approximation. For example, one may consider a linear approximation of the baseline drift at each epoch b_r instead of the constant used here, as well as better characterizations of the prior that based on biophysical constraints. In addition, we are currently working on an analogous model for inter-ictal spike based source localization. In any case, we believe our attempt to carefully consider and model the question of how best to leverage the availability of repeated recordings will lead in the future to successful attempts to optimize their use.

Acknowledgments

This work was supported in part by the National Institutes of Health grant P41GM103545.

REFERENCES

1. Erem B, Coll-Font J, Martínez-Orellana R, Stovicek P, Brooks Dana H. Using Transmural Regularization and Dynamic Modeling for Non-Invasive Cardiac Potential Imaging of Endocardial Pacing with Imprecise Thoracic Geometry. *IEEE TMI*. 2014;1–13.
2. Martínez-Orellana R, Erem B, Brooks DH. Time Invariant Multielectrode Averaging For Biomedical Signals. *ICASSP*. 2013
3. Kaipio, J.; Somersalo, E. Statistical and Computational Inverse Problems. New York: App. Math. Sci. Springer; 2010.
4. Coll-Font J, Erem B, Stovicek P, Brooks DH. Using a new time-independent average method for noninvasive cardiac potential imaging of endocardial pacing with imprecise thorax geometry. *CinC* 2014. 2014
5. Erem B, Stovicek P, Brooks DH. Manifold learning for analysis of low-order nonlinear dynamics in high-dimensional electrocardiographic signals. *ISBI*. 2012 May.:844–847.
6. Boyd S. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trends Mach. Learn.* 2010; 3(1):1–122.

Fig. 1.

Histogram of the displacement of the measurements from the mean for 4 different repetitions. The dashed lines indicate the mean of each epoch.

Fig. 2. Histogram of the localization error for all 3 subjects: Probabilistic inverse (blue bars), spline-average (green) and ensemble-average (red).

Fig. 3.

Histogram of the localization error variation of the probabilistic inverse with respect to spline average and ensemble average. The bar colors correspond to the comparison: Spline average - probabilistic (blue) and ensemble average - probabilistic (red). The vertical lines indicate the average variation w.r.t. spline (red) and ensemble (cyan). Each row represents a different subject.