Hierarchical Multiple-Model Bayesian Approach to Transmural Electrophysiological Imaging

Azar Rahimi, Jingjia Xu, and Linwei Wang

Rochester Institute of Technology, Rochester, NY, 14623, USA

Abstract. Noninvasive electrophysiological (EP) imaging of the heart aims to mathematically reconstruct the spatiotemporal dynamics of cardiac current sources from body-surface electrocardiographic (ECG) data. This ill-posed problem is often regularized by a fixed constraining model. However, this approach enforces the source distribution to follow a preassumed spatial structure that does not always match the varying spatiotemporal distribution of current sources. We propose a hierarchical Bayesian approach to transmural EP imaging that employs a continuous combination of multiple models, each reflecting a specific spatial property for current sources. Multiple models are incorporated as an Lp-norm prior for current sources, where p is an unknown hyperparameter with a prior probabilistic distribution. The current source estimation is obtained as an optimally weighted combination of solutions across all models, the weight being determined from the posterior distribution of p inferred from ECG data. The accuracy of our approach is assessed in a set of synthetic and real-data experiments on human heart-torso models. While the use of fixed models ($L1$- and $L2$-norm) only properly recovers sources with specific structures, our method delivers consistent performance in reconstructing sources with various extents and structures.

Keywords: Transmural electrophysiological imaging, Lp-norm regularization, multiple-model estimation, Bayesian inference.

1 Introduction

Noninvasive cardiac electrophysiological (EP) imaging aims to noninvasively reconstruct subject-specific cardiac current source dynamics from body-surface electrocardiographic (ECG) data. It has shown promising potential in diagnosis of cardiac dysfunctions such as arrhythmia [1]. However, cardiac EP imaging involves a severely ill-posed problem caused by two factors: 1) the limited number of ECG measurements compared to the large number of unknown cardiac sources; and 2) the lack of a unique transmural solution because, in a *quasi-static* electromagnetic field, different configurations of 3D sources may produce identical surface measurements [2].

Over the past three decades, different approaches have been developed for noninvasive cardiac EP imaging. It includes reconstruction of potential dynamics on the epicardium [3], activation front on both the epicardium and endocardium [4],

P. Golland et al. (Eds.): MICCAI 2014, Part II, LNCS 8674, pp. 538–545, 2014.

and intramural action potential dynamics into the depth of the ventricular wall [5,6]. To overcome the ill-posedness of the problem, these methods often impose a certain constraining model to regularize the problem. The most widely-used constraining model includes variants of $L2$-norm that imposes spatial and/or temporal smoothness on the solution [3,4,5,6]. Recently, $L1$-norm models are emerging to enforce the solution to be spatially sparse [7,1,8,9].

However, imposing a fixed prior model involves a potential risk of mismatch between the model and the actual data. In cardiac EP imaging, for example, the application of pre-defined models enforces a pre-assumed spatial structure on the source distribution that does not necessarily reflect the dynamic spatiotemporal property of cardiac sources. Cardiac current sources undergo a complex spatiotemporal change in each cardiac cycle. Their activity starts from a few sparsely distributed focal points, and then forms a sharp region of excitation wavefront with time-changing spatial structures. Furthermore, in a pathologic heart with increased heterogeneity in tissue property, the spatial structure of current sources are even harder to predict *a priori*. Therefore, a fixed model that is suitable for a specific condition of current sources may not be applicable to other scenarios. Rather, the model needs to be able to adapt to the spatiotemporal change in cardiac current sources in various conditions. This limitation of fixed-model regularization has been demonstrated in [10] using a fixed Lp-norm model with a predefined p value, where different values of p are needed in order to optimally reconstruct sources with different spatial extents and structures.

In this paper, we present a hierarchical Bayesian approach to cardiac EP imaging that employs a continuous combination of models, each reflecting a specific spatial property for transmural current sources. The 3D source estimation is then obtained as an optimally weighted combination of solutions across all models. Specifically, multiple models are incorporated as an Lp-norm prior for current sources, where the hierarchical structure allows the inclusion of p as an unknown hyperparameter with a prior probabilistic distribution. Compared to a pre-defined value for p, our model consists of a continuous set of models with different values of p. The joint posterior distribution of current sources and hyperparameter p is then calculated using a Markov chain Monte Carlo (MCMC) technique. This posterior distribution automatically assigns a weight to each model, where the weights are algorithmically determined by the data. Finally, the current source distribution is obtained as an integral over the hyperparameter p. This allows multiple models to contribute to the final solution based on their weights determined from the posterior distribution of p. In addition, the hierarchical Bayesian structure allows us to include additional parameters, including the variance of measurement noise and current source prior, to be automatically inferred from the data. This avoids the challenge of attempting to pre-define sub-optimal parameters that often arises in deterministic regularization.

In a set of synthetic and real-data experiments, we assess the performance of our method compared to fixed-model approaches ($L1$- and $L2$-norm models) in recovering 3D source distributions with different extents and structures. Experiment results report superiority of our approach in recovering sources with

different extents and structures while fixed-model approaches are only able to reconstruct specific types of source distributions. Finally, our hierarchical multiple-model Bayesian approach can be applied to a broader variety of inverse problems that face the challenge of model-data mismatch if a fixed model is used.

2 Methodology

2.1 Forward Measurement Model

The *quasi-static* electromagnetism [2] explains the relation between ECG measurements and cardiac current sources as:

$$\sigma_{blk}\nabla^2\phi_e(\mathbf{r}) = \nabla\cdot(-\mathbf{D}_{int}v(\mathbf{r})),\ \forall\mathbf{r}\in\Omega_h;\quad \sigma\nabla^2\phi(\mathbf{r}) = 0,\ \forall\mathbf{r}\in\Omega_{t/h} \tag{1}$$

Poisson's equation ((1) left) describes, on a bidomain heart model, the relation between the extracellular potential ϕ_e and current sources v within the heart volume Ω_h. $\mathbf{D}_{int}$ is the anisotropic intracellular conductivity tensor obtained by a fiber structure mapper [11]. σ_{blk} is the bulk conductivity in the myocardium and is assumed to be isotropic and homogeneous. Laplace's equation ((1) right) describes, on the monodomain torso model, the distribution of potential ϕ within the torso volume external to the heart $\Omega_{t/h}$, which is often assumed homogeneous with conductivity σ. Solving these equations numerically on a subject heart-torso model [5,6] gives a linear biophysical model: $\mathbf{\Phi} = \mathbf{H}\mathbf{v} + \mathbf{n}$, where $\mathbf{\Phi}$ represents the $m\times 1$ measurement vector, $\mathbf{v}$ is $n\times 1$ current sources, $\mathbf{H}$ denotes $m\times n$ transfer matrix, and $\mathbf{n}$ is $m\times 1$ measurement noise.

2.2 Hierarchical Multiple-model Bayesian Inference

We propose a hierarchical Bayesian interpretation of cardiac EP imaging with an *Lp*-norm prior for current sources, where p is treated as an unknown hyper-parameter randomly distributed between 1 and 2. In this way, we are able to incorporate a continuous set of models to constrain our reconstruction.

According to the Bayes' theorem, the posterior probability of current sources $\mathbf{v}$ and model parameters $\mathbf{\Theta}$ given ECG data $\mathbf{\Phi}$ is formulated as:

$$P(\mathbf{v},\mathbf{\Theta}\mid\mathbf{\Phi}) \propto P(\mathbf{\Phi}\mid\mathbf{v},\mathbf{\Theta})P(\mathbf{v}\mid\mathbf{\Theta})P(\mathbf{\Theta}) \tag{2}$$

where $P(\mathbf{\Phi}\mid\mathbf{v},\mathbf{\Theta})$ and $P(\mathbf{v}\mid\mathbf{\Theta})P(\mathbf{\Theta})$ denote the likelihood and prior terms.

Multiple-model Prior. Considering the *Lp*-norm prior for statistically independent current sources $\mathbf{v}$, $||\mathbf{v}||_p = (\sum_i (v_i)^p)^{1/p}$, we represent the source probability distribution using a generalized Gaussian distribution as:

$$P(\mathbf{v}\mid p,\delta) = w(p)^n\delta^{-n}exp(-c(p)\sum_i \mid \frac{v_i}{\delta}\mid^p) \tag{3}$$

$$w(p) = \frac{p\Gamma(3/p)^{1/2}}{2\Gamma(1/p)^{3/2}},\ \ c(p) = [\frac{\Gamma(3/p)}{\Gamma(1/p)}]^{p/2}$$

where p represents the order of the Lp-norm prior. Fixing the p value to 1 or 2 converts the source prior to an $L1$-norm or $L2$-norm model. Instead, we consider it to be an unknown hyperparameter in order to incorporate a continuous combination of multiple models into our approach. To give no prior preference over the value of p, we assume it to follow a uniform distribution between 1 and 2. δ denotes the variance (inverse of precision) of the current source prior. It play a similar role to the regularization parameter in a deterministic setting, and it controls the contribution of the source prior to the regularization. To avoid the need to pre-define a sub-optimal value for δ through empirical studies, we assume it to be an unknown hyperparameter with a uniform distribution.

Likelihood Term. Assuming zero-mean normal distribution for the measurement noise, and statistically independent measurements $\mathbf{\Phi}$, the likelihood term follows a Gaussian distribution as:

$$P(\mathbf{\Phi} \mid \mathbf{v}, \mathbf{\Sigma}) = \frac{1}{(2\pi)^{m/2}|\mathbf{\Sigma}|^{1/2}} exp(-\frac{1}{2}(\mathbf{\Phi} - \mathbf{H}\mathbf{v})^T \mathbf{\Sigma}^{-1}(\mathbf{\Phi} - \mathbf{H}\mathbf{v})) \qquad (4)$$

where $\mathbf{\Sigma}$ represents the noise covariance. It controls the contribution of the data-fitting term to the regularization. Here we also consider it to be an unknown hyperparameter with a uniform distribution. Substituting the likelihood and prior terms in (2), the joint posterior distribution (2) can be re-formulated as:

$$P(\mathbf{v}, \mathbf{\Theta} \mid \mathbf{\Phi}) \propto P(\mathbf{\Phi} \mid \mathbf{v}, \mathbf{\Sigma}) P(\mathbf{\Sigma}) P(\mathbf{v} \mid p) \mathbf{P}(p) P(\delta) \qquad (5)$$

where $\mathbf{\Theta} = \{\delta, p, \mathbf{\Sigma}\}$ denotes a vector of all model hyperparameters.

Bayesian Inference via MCMC. A full Bayesian analysis of this problem is obtained by sampling the joint posterior distribution (5) using a MCMC technique called slice sampling [12]. Slice sampling generates samples of a random variable by uniformly sampling from under the curve of its density function. Unlike Gibbs sampling that requires conditional distributions of model parameters, or Metropolis-Hastings scheme that requires an accurate selection of the proposal distribution for an efficient random walk, slice sampling enables us to directly sample the joint posterior distribution with minimum tuning required.

Finally, the posterior distribution of current sources $P(\mathbf{v}|\mathbf{\Phi})$ is calculated as an integral over hyperparameters $\mathbf{\Theta}$. The final solution is thus a combination of multiple models (different p values in the Lp-norm prior), where each model is weighted/modulated by the posterior probability of its occurrence determined by the data. The posterior mean of $\mathbf{v}$ is used as the estimate of current sources.

3 Results

Synthetic Experiments. To test the impact of the prior model on the accuracy of source construction, we conduct a set of synthetic experiments on an

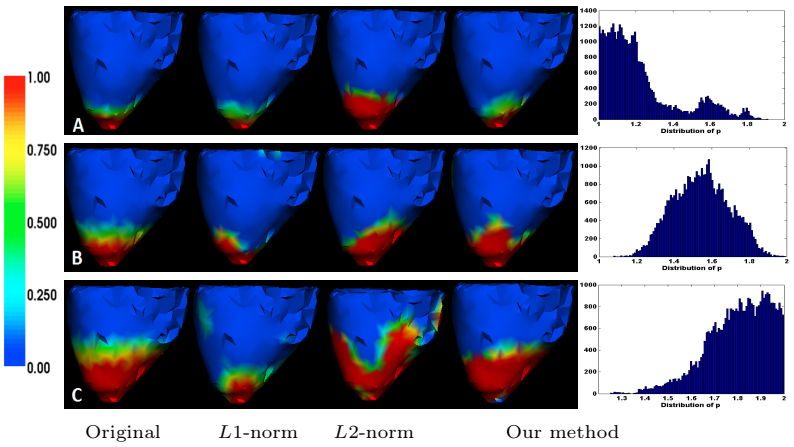

Fig. 1. Three examples of 3D source reconstruction using our method vs. $L1$- and $L2$-norm models. Active sources are centered at the apex of LV and cover 6%, 17%, and 40% of the myocardium in A, B and C, respectively. Posterior distributions of hyperparameter p for the three cases are also shown in the rightmost panel.

image-derived human heart-torso model. The torso surface is represented using 370 nodes. The heart volume is represented using 1019 nodes with $7mm$ resolution. In total, 40 different settings are considered, where active source region is located randomly in the myocardium with extents ranging from 5% to 40% of the ventricles. Settings of focal and extended sources simulate sparse source distribution at the beginning of the excitation and extended source distribution along diseased regions of the heart, respectively. Current sources within the active region are assigned with value 1 while the rest with value 0. 370-lead ECG data for each setting are simulated and corrupted with 20 dB Gaussian noise as input in our method. The 3D source reconstruction accuracy is measured in terms of correlation coefficient (CC) between the estimated and true sources. The performance of our approach is also compared to two fixed-model approaches: regularization using $L1$- and $L2$-norm by fixing the value of p to be 1 and 2.

Fig 1 presents 3 examples where the active source region is centered at left ventricle (LV) apex, and covers 6% , 17%, and 40% of LV (Fig 1A, 1B, and 1C), respectively. $L1$-norm detects the focal source (CC= 0.69), but its performance decreases when the region of active sources expands (CC= 0.60 and 0.40). In contrast, $L2$-norm provides an overly-smooth solution for the focal source region (CC= 0.52), but it obtains better estimation of larger source regions (CC= 0.65 and 0.68). Our method shows a performance similar to $L1$-norm (CC= 0.67) for the focal source region (Fig 1A), but it outperforms both fixed-model approaches in larger source regions with CC=0.70 and 0.74 (Fig 1B and 1C). Posterior distributions of hyperparameter p for these 3 cases (Fig 1) suggest that higher weights are assigned by data to p values closer to 1 for the focal source. This weight distribution for p shifts to higher values when the source region expands.

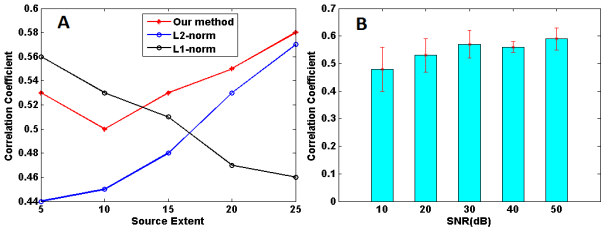

Fig. 2. A) Comparison of our method with $L1$- and $L2$-norm models in reconstructing active source regions with different extents ranging from 5% to 40% of the myocardium. B) Correlation coefficient obtained by our method for active sources in the presence of white Gaussian noises with different SNR levels.

As illustrated in Fig 2A, in these 40 different settings of source distributions, our method delivers consistent results for sources with various sizes while $L1$- and $L2$-norm models only satisfactorily capture specific type of source distributions. According to the paired student's t-test, the accuracy of our proposed method (CC= 0.52 ± 0.08) is significantly higher than $L1$ and $L2$-norm methods with CC= 0.50 ± 0.11 and 0.48 ± 0.13 ($p < 0.005$).

In addition, values of δ and $\mathbf{\Sigma}$ affect the accuracy of Bayesian inference by controlling the contribution of the prior and the data-fitting term. To verify the robustness of our approach to initial values of these parameters, we conduct experiments with different measurement noises (SNR levels ranging from 50 to 10 dB). For all experiments, initial values of δ and $\mathbf{\Sigma}$ are identical. As shown in Fig 2B, a slight change in CC value of different measurement noises demonstrates the ability of the hierarchical approach to automatically adapt hyperparameter values to deliver optimal performance. In these experiments, the final value of δ is mainly distributed over the range of 0.3 to 0.6, regardless of its initial value.

Real-Data Experiments. An important application of our method is to outline 3D ischemic regions during the ECG ST segment, where sources distribute along the boundary between ischemic and non-ischemic regions. Using 4 patient-specific human heart-torso models with prior ischemia [13], we assess the feasibility of our method in source imaging along the ischmic region. ECG data are recorded by 120-lead Dalhousie protocol [14]. Reference ischemic regions (Fig 3, Ref) are provided by cardiologists using LV 17-segment model [15].

As shown in Fig 3A, the center of ischemic region (black contour) in case 1 is located at segment 8. The sparse solution of $L1$-norm only detects part of the ischemic region border. The overly-diffused estimation of $L2$-norm does not reveal any information about the location or structure of the ischemic region. In comparison, our method is able to outline the distribution of current sources along the ischemic region (non-blue regions). In case 2, where the ischemic region center resides between segments 3, 4, 9, and 10, the three methods behave similar to case 1. In case 3, the ischemic region is centered at segment 10 and extends to the septum and inferior LV. Our method outperforms $L1$- and $L2$-norm estimations as they both include a false positive ischemic region at the right ventricle (RV). Case 4 has two ischemic regions centered at segment 1 and seg-

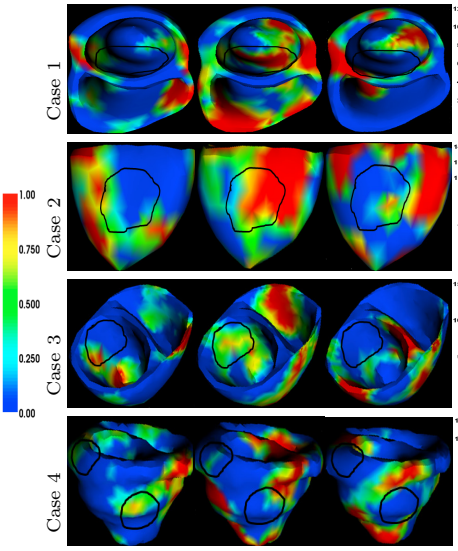

	Center	Segments	Size
Ref	8	1, 2, 3, 8, 9, 13, 14, 15	31%
Results	8,9	2, 3, 8, 9, 14, 15	27%

Ref	3/4/ 9/10	3,4,9,10	30%
Results	4	3, 4, 9, 10, 11, 15	36%

Ref	10/11	3, 4, 5, 9, 10, 11, 12, 15, 16	52%
Results	10	3, 4, 9, 10, 11, 16	45%

Ref	15	1, 9, 10, 11, 15, 17	14%
Results	10/15	9, 10, 11, 15, 17	12%

Fig. 3. Estimation of current source activity (non-blue color) along the ischemic region border (black contour) during the ECG ST-segment for 4 human subjects

ment 15. $L1$-norm behaves similar to other cases. $L2$-norm can only detect one of the ischemic regions, which is falsely extended to RV. Our method precisely detects the ischemic region located at segment 15, although it only shows some source activity around segment 1 without providing more information about the ischemia structure/extent. The results are quantified by thresholding the solution to obtain the region bordering the ischemic core, and then calculating the ischemic region in terms of location (using the 17-segment model of LV) and size (ratio of ischemic region volume to the total LV volume), as shown in Fig 3C.

Interestingly, as shown in Fig 3B, in cases 1 and 2 that have relatively compact ischemic regions (30%), the distributions of p is centered around 1.2. For case 3 that has a large ischemic region (52%), the distribution of p shifts toward higher values (~ 1.5). In the unique case 4, although the ischemic region is small (14%), the current source distribution is expected to be more diffused because of the presence of two separate ischemic regions. Accordingly, the p is more evenly distributed from 1.4 to 1.6. These results further demonstrate the capability and necessity of a multiple-model approach to automatically adjust the contribution of different models based on the source structure by inferring from the data.

4 Conclusions

This paper presents a hierarchical Bayesian approach to transmural EP imaging by incorporating a continuous combination of multiple constraining models. Our experiments show that the presented approach perform consistently for a

variety of source distributions by automatically adapting the model combination to different source structures. Currently, we focus on source resolutions around 5-7 mm. In future work, we will study the impact of different resolutions on the performance of our method. In addition, we will conduct real-data experiments on a broader spectrum of heart diseases where current sources are expected to exhibit more complex spatiotemporal structures.

Acknowledgement. This work is supported by the National Science Foundation CAREER Award ACI-1350374.

References

1. Ghosh, S., Rudy, Y.: Application of l1-norm regularization to epicardial potential solutions of the inverse electrocardiography problem. Annals of Biomedical Engineering 37(5), 692–699 (2009)
2. Plonsey, R.: Bioelectric phenonmena. McGraw Hill, New York (1969)
3. Pullan, A.J., Cheng, L.K., Nash, M.P., Bradley, C.P., Paterson, D.J.: Noninvasive electrical imaging of the heart: Theory and model development. Annals of Biomedical Engineering 29, 817–836 (2001)
4. Huiskamp, G., Greensite, F.: A new method for myocardial activation imaging. IEEE Transactions on Biomedical Engineering 44(6), 433–446 (1997)
5. Wang, L., Zhang, H., Wong, K., Liu, H., Shi, P.: Physiological-model-constrained noninvasive reconstruction of volumetric myocardial transmembrane potentials. IEEE Transactions on Biomedical Engineering 57(2), 296–315 (2010)
6. He, B., Wu, D.: Imaging and visualization of 3-d cardiac electric activity. IEEE Transactions on Information Technology in Biomedicine 5(3), 181–186 (2001)
7. Xu, J., Dehaghani, A.R., Gao, F., Wang, L.: Localization of sparse transmural excitation stimuli from surface mapping. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012, Part I. LNCS, vol. 7510, pp. 675–682. Springer, Heidelberg (2012)
8. Shou, G., Xia, L., Liu, F., Jiang, M., Crozier, S.: On epicardial potential reconstruction using regularization schemes with the l1-norm data term. Physics in Medicine and Biology 56(1), 57–72 (2011)
9. Wang, L., Qin, J., Wong, T.T., Heng, P.A.: Application of l1-norm regularization to epicardial potential reconstruction based on gradient projection. Physics in Medicine and Biology 56(19), 6291–6310 (2011)
10. Rahimi, A., Xu, J., Wang, L.: Lp-norm regularization in volumetric imaging of cardiac current sources. Computational and Mathematical Methods in Medicine (2013)
11. Nash, M.: Mechanics and Material Properties of the Heart using an Anatomically Accurate Mathematical Model. PhD thesis, Univ. of Auckland, New Zealand (1998)
12. Neal, R.M.: Slice sampling. The Annals of Statistics 31, 705–767 (2003)
13. Goldberger, A.L., et al.: Physiobank, physiotoolkit, and physionet components of a new research resource for complex physiological signals. Circulation 101, 215–220 (2000)
14. Title, L.M., Iles, S.E., Gardner, M.J., Penney, C.J., Clements, J.C., Horacek, B.M.: Quantitative assessment of myocardial ischemia by electrocardiographic and scintigraphic imaging. Journal of Electrocardiology 36(suppl.), 17–26 (2003)
15. Cerqueira, M.D., Weissman, N.J., Dilsizian, V., et al.: Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart. Circulation 105, 539–542 (2002)