A Patient-Specific Equivalent Dipole Model

Gabriel Cardoso, Geneviève Robin, Andony Arrieula, Mark Potse, Michel Haïssaguerre, Eric Moulines, Rémi Dubois

A Patient-Specific Equivalent Dipole Model

Gabriel Cardoso[1,4,6], Geneviève Robin[2], Andony Arrieula[3,4,5], Mark Potse[5,4,3], Michel Haïssaguerre[4,6,7], Eric Moulines[1], Rémi Dubois[4,6]

[1] CMAP Ecole Polytechnique, Palaiseau, France
[2] Université d'Evry, Evry, France
[3] CARMEN Research Team, Inria Bordeaux – Sud-Ouest, Talence, France
[4] IHU Liryc, fondation Bordeaux Université, Pessac, France
[5] Univ Bordeaux, IMB, UMR 5251, Talence, France
[6] Univ Bordeaux, CRCTB U4045, INSERM, Bordeaux, France
[7] Bordeaux University Hospital (CHU), Electrophysiology and Ablation Unit, Pessac, France

Abstract

Sophisticated models for the electrocardiographic inverse problem are available, but their reliance on imaging data and large numbers of electrodes limit their use. Simple models such as the equivalent dipole model (EDM) therefore remain relevant. We developed a probabilistic approach to the equivalent unbounded uniform single dipole problem and developed a natural extension to the bounded nonuniform case that relies on a patient-specific statistical inference of the propagation mechanism between the location of the dipole and the electrode locations. The two models were tested on data simulated with a detailed heart-torso model with four different activation sequences and three different sets of tissue characteristics. We observed a throughout enhancement of the ability to reconstruct the ECG of the patient-specific model when compared to the uniform unbounded dipole model.

1. Introduction

Modeling the electrical activity of the heart as a single, time-varying electric current dipole dates back to Einthoven et al.[1]. Einthoven's triangle [2, volume 1 section 10.3.6] assumes a single fixed dipole in an unbounded uniform body. Information from leads I, II, III of the standard ECG allows measurement of the individual components of such a dipole. However, the assumption of an unconfined uniform body significantly limits the interpretability of the resulting dipole parameters. A current dipole with magnitude $\boldsymbol{p}_t \in \mathbb{R}^3$ and position $\boldsymbol{x}_p \in \mathbb{R}^3$ in a conductor with an arbitrary bounded volume, generates a potential

$$y_t^i = \boldsymbol{c}(\boldsymbol{x}_i, \boldsymbol{x}_p) \cdot \boldsymbol{p}_t, \quad \boldsymbol{c} : \mathbb{R}^3 \times \mathbb{R}^3 \to \mathbb{R}^3 \tag{1}$$

at position $\boldsymbol{x}_i \in \mathbb{R}^3$, where $i \in \{1, ..., K\}$ is the index of the measurements [2, Chapter 10]. This is the main ground for the development of vectorcardiography and Frank's lead system, whose goal is to measure each of the components of $\boldsymbol{p}_t$ in an orthogonal basis in the equivalent electrical space. Frank [3] established a method for estimating an orthogonal lead systems, by inducing a current dipole in three directions in an artificial torso and measuring the resulting potential at the surface. This led to the creation of a 3d lead system called the Frank's lead system defined by F_x, F_y, F_z leads that represents an orthogonal basis of the equivalent electrical space.

The vectorcardiogram obtained using this lead system has been successfully applied to a number of problems [4].

An important limitation of this approach is that it is not patient-specific, being experimentally derived from an artificial torso. Therefore we developed a new patient-specific EDM by using Bayesian statistics to infer the function y_t^i for each patient. We tested the method on simulated ECGs produced with a detailed human heart and torso model, with different activation origins and different conduction properties.

2. Equivalent Unbounded dipole model

Our goal is to estimate the posterior distribution of the parameters given the observed data.

The following notation will be used throughout the rest of this paper: If a_t^i is a scalar variable indexed by $i \in \{1, ..., K\}$ and $t \in \{1, ..., T\}$, we write $\boldsymbol{A} = (a_1^1, a_1^2, ...a_T^K)$. Assuming propagation through an unbounded homogeneous medium, the torso potentials are

$$y_t^i | \boldsymbol{p}_t \sim \mathcal{N}\left(\boldsymbol{c}_{\text{HOM}}(\boldsymbol{x}_i) \cdot \boldsymbol{p}_t, \sigma_{i,t}^2\right), \boldsymbol{c}_{\text{HOM}}(\boldsymbol{x}_i) = \boldsymbol{x}_i / \|\boldsymbol{x}_i\|^3 \tag{2}$$

where p_t represents the dipole vector of a current dipole placed at $\mathbf{0}$.

Under this model, the log-likelihood of the observations is

$$\ell(\boldsymbol{Y}|\boldsymbol{P}) \propto -\sum_{t=1}^{T}\sum_{i=1}^{K}\frac{(y_t^i - \boldsymbol{p}_t \cdot \boldsymbol{c}_{\text{HOM}}(\boldsymbol{x}_i))^2}{\sigma_{i,t}^2} \tag{3}$$

If the measurements $(y_t^i)_{t=1,\dots,T}^{i=1,\dots,K}$ start at the beginning of the QRS complex of the ECG then it is reasonable to expect $\boldsymbol{p}_t$ to be small at the beginning, since it corresponds to the onset of the activation of the ventricles. There is no *a priori* reason to expect that the components of the dipole are positive or negative, so a Gaussian distribution centered on $\mathbf{0}$ is a natural candidate for the prior distribution $\boldsymbol{p}_0$.

Most of the common sampling rates in the ECG domain (sampling rate > 250 Hz) are higher than the frequency of the dipole curve $\boldsymbol{p}_t$. Therefore, we expect a small deviation between $\boldsymbol{p}_t$ and $\boldsymbol{p}_{t+1}$, i.e., we expect $\|\boldsymbol{p}_{t+1} - \boldsymbol{p}_t\|$ to be close to 0. This leads us to define the prior distribution as

$$\boldsymbol{p}_{t+1}|\boldsymbol{p}_t = \mathcal{N}(\boldsymbol{p}_t, \sigma_p^2 I), \quad \boldsymbol{p}_0 = \mathcal{N}(\mathbf{0}, I). \tag{4}$$

The log probability of the posterior can be derived through Bayes' equation and it is negative if the observation $\boldsymbol{Y}$ is given, up to an arbitrary normalizing constant, by:

$$\sum_{t=1}^{T}\left[\sum_{i=1}^{K}\frac{(y_t^i - \boldsymbol{p}_t \cdot \boldsymbol{c}_{\text{HOM}}(\boldsymbol{x}_i))^2}{\sigma_{i,t}^2} + \frac{\|\boldsymbol{p}_t - \boldsymbol{p}_{t-1}\|^2}{\sigma_p^2}\right] + \|\boldsymbol{p}_0\|^2 \tag{5}$$

This setting is the well known case of Bayesian linear regression [5] with Gaussian measurements noise and state disturbances. It can be shown that

$$\begin{aligned}\boldsymbol{P}|\boldsymbol{Y} \sim \mathcal{N}(&\boldsymbol{\mu}_P = \Sigma_P \mathcal{M}^T \Sigma_m^{-1} \boldsymbol{Y}, \\ &\Sigma_P = (R_{\text{reg}} + \mathcal{M}^T \Sigma_m^{-1} \mathcal{M})^{-1})\end{aligned} \tag{6}$$

when Σ_P is invertible, which is the case in our settings.

3. Our model

To make the EDL model patient-specific we consider $\boldsymbol{c}$ as a random function that must be estimated for each patient. To develop a Bayesian approach, we need to define a prior over a function space and a likelihood. The main idea is to use the model developed in section 2 to define a Gaussian process as a prior distribution over a function space. For $\boldsymbol{f} : \mathbb{R}^3 \to \mathbb{R}^3$:

$$\begin{aligned}&y_t^i|\boldsymbol{f}, \boldsymbol{p}_t, \boldsymbol{x}_i \sim \mathcal{N}\left(\boldsymbol{c}_{\text{INH}}(\boldsymbol{x}_i) \cdot \boldsymbol{p}_t, \sigma_{i,t}^2\right) \\ &\boldsymbol{c}_{\text{INH}}(\boldsymbol{x}_i; \boldsymbol{f}) = \{\boldsymbol{x}_i + \boldsymbol{f}(\boldsymbol{x}_i)\}/\|\boldsymbol{x}_i\|^3\end{aligned} \tag{7}$$

The prior over the function $\boldsymbol{f}$ will be defined as, following Bishop [5] section 6.4:

$$\boldsymbol{f} \sim \mathcal{GP}(0, k_\theta(\boldsymbol{x}, \boldsymbol{x}')) \tag{8}$$

Variable	Distribution	Interpretation
l_k	$LogNormal(\mu_* = l_0(\boldsymbol{X}), {\sigma_*}^2 = 10)$	Expect values to be in the interval $[\frac{l_0(\boldsymbol{X})}{10}, 10 l_0(\boldsymbol{X})]$
ν_k	$LogNormal(\mu_* = \hat{\sigma}^2_{\text{kernel}}, {\sigma_*}^2 = 1.01)$	Expect values to in the interval $[\frac{\hat{\sigma}^2_{\text{kernel}}}{1.01}, 1.01\hat{\sigma}^2_{\text{kernel}}]$
$\boldsymbol{p}_0$	$\mathcal{N}(\mathbf{0}, I)$	Same as in Section 2
$\boldsymbol{p}_t$	$\mathcal{N}(\mathbf{0}, \sigma_p^2 \hat{\sigma}^2_{\text{kernel}} I)$	Same as in Section 2 but we found numerically that it was best to scale it by variance as the Gaussian process kernels

Table 1. Description of the priors used in the numerical experiments.

The prior and posterior are thus determined by the choice of the kernel function k. We use a component-wise independent scaled radial basis function (RBF) kernel defined as

$$\begin{aligned}k_{\nu,l}(x, x') &= \nu \exp(-(x - x')^2/2l^2) \\ k_\theta(\boldsymbol{x}, \boldsymbol{x}') &= \text{diag}(k_{\nu_i, l_i}(\boldsymbol{x}[i], \boldsymbol{x}'[i]), i = 1, 2, 3).\end{aligned} \tag{9}$$

The parameter ν determines the variance scale of the covariance matrix and in our case controls the magnitude of f, i.e., how far we are from the uniform dipole model. The parameter l indicates the rate at which the covariance between two points decreases. A small l would result in more variation being allowed for f between two neighboring points x, x'. In what follows, $\theta = (\nu_1, l_1, \nu_2, l_2, \nu_3, l_3)$.

Since the model errors are Gaussian, we end up with an exact posterior Gaussian over a functional space [5, section 6]. The likelihood of model (7) is:

$$\mathbb{P}(\boldsymbol{Y}|\theta, \boldsymbol{P}) = \mathbb{E}_{\boldsymbol{f}(\boldsymbol{X})|\theta}\left[\mathbb{P}(\boldsymbol{Y}|\boldsymbol{f}, \theta, \boldsymbol{P})\right] \tag{10}$$

From (8), $\boldsymbol{f}(\boldsymbol{X})|\theta$ is distributed according to a multivariate Gaussian distribution and the same is true for $\boldsymbol{Y}|\boldsymbol{f}, \theta, \boldsymbol{P}$. Using formula (2.115) from Bishop [5] the likelihood has the distribution

$$\boldsymbol{Y}|\theta, \boldsymbol{P} \sim \mathcal{N}(0, \Sigma(\theta, \boldsymbol{P}, \boldsymbol{X})) \tag{11}$$

where we refer to appendix B for the definition of $\Sigma(\theta, \boldsymbol{P}, \boldsymbol{X})$. We proceed in the same manner as in Section 2 by defining a prior over the parameters in order to compute the posterior distribution. In Table 3 we use the multiplicative mean and multiplicative standard deviation parametrization of the log normal distribution. The typical geometry scale $l_0(\boldsymbol{X})$, $\hat{\sigma}^2_{\text{kernel}}$ as well as the procedure to estimate σ_p are defined in the appendix.

In this setting, our goal is to estimate the parameter $\hat{\theta}$ by maximizing the posterior log-likelihood:

$$\hat{\theta} = \arg\max_\theta \log \mathbb{P}(\boldsymbol{Y}|\theta, \boldsymbol{P}) + \log \mathbb{P}(\theta) + \log \mathbb{P}(\boldsymbol{P}) \tag{12}$$

4. Numerical Evaluation

To solve (12), we used the Adam gradient descent algorithm [6]. For computation efficiency, a batch of elec-

trodes was sampled for each optimization step. Other parameters used in the optimization are given in appendix C. We initialized the dipole parameter $\boldsymbol{P}$ using $\mu_{\boldsymbol{P}}$ defined in (6). The metric used to evaluate goodness of fit was the coefficient of determination, known as R^2, defined as $R^2(\boldsymbol{y}, \boldsymbol{z}) = 1 - \frac{\sum_{t=1}^{T}(y_t - z_t)^2}{\sum_{t=1}^{T}(y_t - \bar{y})^2}$ where $\bar{y} = \frac{1}{T}\sum_{t=1}^{T} y_t$.

We used a detailed anatomical model of the heart and torso created from computed tomography data. Simulations were performded on a bi-ventricular mesh with a uniform edge length of 200 µm. A human ventricular ionic model [7] was used with a monodomain reaction-diffusion equation. The potentials at 252 electrode locations on the torso were computed using lead fields [8]. Three anatomical configurations were created: one with a full healthy tissue, and two with low-conductivity zones with radii of 12 and 18 mm. In these zones the conductivity was 10 times smaller than in the healthy tissue. For each configuration, 3 paced beats from different origins (A, B, C) and one normal ventricular beat were simulated. The pacing locations and the low-conductivity zones are shown in Fig. 3. In total, 12 sets of signals were simulated. For computational efficiency, the signals were sampled at 500 Hz.

We placed the dipole at the barycenter of the three stimulation points A, B, and C. We use the following notation for the figures: INH-DIP corresponds to the solution of(12), and HOM-DIP corresponds to the solution of (6).

We compared the median per patient R^2 score through the different beat origins and different sizes of the low-conduction zones in figure 1, where we see that the INH-DIP model performed considerably better than the HOM-DIP model. The corresponding reconstructed electrograms are shown in figure 2.

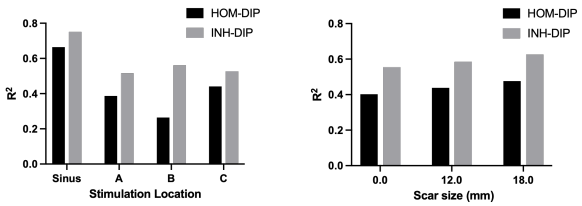

Figure 1. Comparison between the inter-patient mean of the per simulation median of the R^2 scores for each lead. We can see an enhancement of the R^2 scores throughout.

5. Discussion

The model developed in section 2 is simplistic and several approaches have been developed to reduce the model error, e.g. by Frank [3]. However, to the best of our knowledge, none of them is based on probabilistic modeling of the function $\boldsymbol{c}$ in (1). We have presented a probabilistic framework for solving the EDM problem that allows us to derive a natural extension to a patient-specific model that outperforms the equivalent unbounded dipole model in all situations tested. There are several ways to further develop and validate this model. One of them is the development of the kernel function, for example, to handle the correlation between the different components of the function f. Validation with clinical data is also a major challenge, as is the ability of the model to handle even more uncertain data (no knowledge of heart placement, varying geometries).

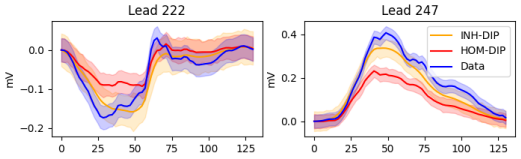

Figure 2. The means of the Likelihoods $Y|\boldsymbol{P}, \theta, X$ for each model as well as the confidence intervals (shaded) corresponding to $\mu \pm 3\sigma$ in two sample ECG leads.

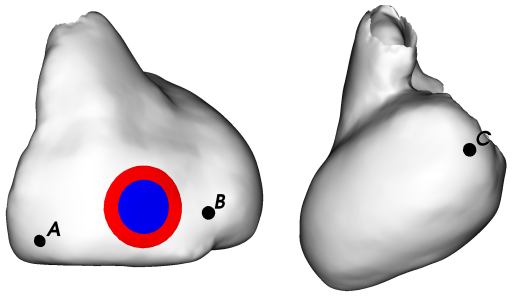

Figure 3. Mesh of the modeled heart used for simulations, with the location of the pacing sites (A, B, C) and the low-conduction zone with a radius of 12 mm in blue and a radius of 18 mm in red.

References

[1] Einthoven W, Fahr G, de Waart A. On the direction and manifest size of the variations of potential in the human heart and on the influence of the position of the heart on the form of the electrocardiogram. American Heart Journal 1950;40(2):163–211. ISSN 0002-8703. URL https://www.sciencedirect.com/science/article/pii/0002870350901657.

[2] Comprehensive electrocardiology, 2010. URL http://dx.doi.org/10.1007/978-1-84882-046-3.

[3] Frank E. An accurate, clinically practical system for spatial vectorcardiography, May 1956. URL http://dx.doi.org/10.1161/01.CIR.13.5.737.

[4] Armoundas A, Feldman A, Mukkamala R, Cohen R. A single equivalent moving dipole model: An efficient approach for localizing sites of origin of ventricular electrical activation. Annals of biomedical engineering 06 2003;31:564–76.

[5] M. C. Pattern Recognition and Machine Learning. Springer, 2016. ISBN 9780387310732.

[6] Kingma DP, Ba J. Adam: A method for stochastic optimization, 2014. URL https://arxiv.org/abs/1412.6980.

[7] ten Tusscher KHWJ, Noble D, Noble PJ, Panfilov AV. A model for human ventricular tissue. American Journal of Physiology Heart and Circulatory Physiology 2004; 286(4):H1573–H1589. URL https://doi.org/10.1152/ajpheart.00794.2003. PMID: 14656705.

[8] Potse M. Scalable and accurate ECG simulation for reaction-diffusion models of the human heart. Front Physiol 2018; 9:370.

Appendix

A. Matrix formulation equivalent uniform single dipole

$$\begin{gathered}\boldsymbol{Y}=\mathcal{N}(\mathcal{M}(\boldsymbol{X})\boldsymbol{P},\Sigma_m),\quad M(\boldsymbol{X})_{i,j}=x_i^j/\|\boldsymbol{x}_i\|^3\\ \mathcal{M}=\text{blockdiag}(\underbrace{M,\cdots,M}_{T\text{times}})\quad \Sigma_m^t=\text{diag}(\sigma_{1,t}^2,...,\sigma_{K,t}^2)\\ \Sigma_m=\text{blockdiag}(\Sigma_m^1,\cdots,\Sigma_m^T)\quad \Sigma_p=\begin{bmatrix}\sigma_0^2 & 0 & 0 & 0 & \dots & 0\\ 0 & \sigma_0^2 & 0 & 0 & \dots & 0\\ 0 & 0 & \sigma_0^2 & 0 & \dots & 0\\ 0 & 0 & 0 & \sigma^2 & \dots & 0\\ \dots & \dots & \dots & \dots & \dots & 0\\ 0 & 0 & 0 & 0 & \dots & \sigma^2\end{bmatrix}\\ R_{i,j}=\begin{cases}1, & \text{if } i=j\\ -1, & \text{if } i=(j-3); i\geq 3\\ 0, & \text{otherwise}\end{cases},\quad R_{reg}=R^T\Sigma_p^{-1}R\end{gathered} \tag{13}$$

B. Matrix formulation Gaussian process

$$\begin{gathered}C_k(\boldsymbol{P})=\begin{pmatrix}p_0^k\mathbb{I}_{K,K}\\ \cdots\\ p_T^k\mathbb{I}_{K,K}\end{pmatrix},\text{diag}(\|\boldsymbol{x}_1\|^3,...,\|\boldsymbol{x}_K\|^3)^{-1}\\ \Sigma(\theta,\boldsymbol{P},X)=\Sigma_m+\sum_{k=1}^{3}C_k(\boldsymbol{P})K_{\nu_k,l_k}(X,X)C_k^T(\boldsymbol{P})\end{gathered} \tag{14}$$

C. Optimization GP

For the ensuring $l > 0$ and $\nu^2 > 0$ of (12) for each kernel we optimize the log of those quantities. The typical scale length for a given geometry $\boldsymbol{X}$ is $l_0 = \sum_{i=1}^{K}\min\{\|\boldsymbol{x}_i-\boldsymbol{x}_j\||j\neq i\}/K$ The other parameters are specified below:

Variable	Value	Variable	Value
Batch size	32	Learning rate dipole	0.5
Learning rate $logl$	0.01	Learning rate $log\nu$	0.01
Number of iterations	10000	β (from Adam)	$(0.9, 0.999)$
ϵ (from Adam)	10^{-8}	σ_{kernel}	$0.1\frac{Tr(\Sigma_m)}{Tr(S)\boldsymbol{P}}$
σ_p	$\frac{4\|\boldsymbol{Y}_{\max_t}\|}{T\|M(\boldsymbol{X})(1,1,1)\|}$		

where $\boldsymbol{Y}_{\max_t}$ is the K-dimensional vector having each component being the maximum over time of y_t^i.

Acknowledgements

This work was supported by the Région Nouvelle-Aquitaine, grant nr. 2017 – 1R50109 – 00013434; the European Research Council (H2020 grant agreement number 715093, ECSTATIC); and the French National Research Agency, grant references ANR-10-IAHU04-LIRYC and ANR-11-EQPX-0030. This work was granted access to the HPC resources of CINES under the allocation 2021-A0110307379 made by GENCI.