Deep Learning Formulation of ECGI for Data-driven Integration of Spatiotemporal Correlations and Imaging Information

Tania Bacoyannis, Julian Krebs, Nicolas Cedilnik, Hubert Cochet, Maxime Sermesant

► To cite this version:

Tania Bacoyannis, Julian Krebs, Nicolas Cedilnik, Hubert Cochet, Maxime Sermesant. Deep Learning Formulation of ECGI for Data-driven Integration of Spatiotemporal Correlations and Imaging Information. FIMH 2019 - 10th International Conference on Functional Imaging and Modeling of the Heart, Jun 2019, Bordeaux, France. pp.20-28. hal-02108958v2

HAL Id: hal-02108958

<https://hal.inria.fr/hal-02108958v2>

Submitted on 3 May 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Deep Learning Formulation of ECGI for Data-driven Integration of Spatiotemporal Correlations and Imaging Information

Tania Bacoyannis¹, Julian Krebs¹, Nicolas Cedilnik^{1,2}, Hubert Cochet², and Maxime Sermesant¹

¹ Inria, Université Côte d'Azur, France

² Liryc Institute, Bordeaux, France

Abstract. The challenge of non-invasive Electrocardiographic Imaging (ECGI) is to re-create the electrical activity of the heart using body surface potentials. Specifically, there are numerical difficulties due to the ill-posed nature of the problem. We propose a novel method based on Conditional Variational Autoencoders using Deep generative Neural Networks to overcome this challenge. By conditioning the electrical activity on heart shape and electrical potentials, our model is able to generate activation maps with good accuracy on simulated data (mean square error, $MSE = 0.095$). This method differs from other formulations because it naturally takes into account spatio-temporal correlations as well as the imaging substrate through convolutions and conditioning. We believe these features can help improving ECGI results.

Keywords: ECGI · Deep learning · Simulation · Generative Model

1 Introduction

Electrocardiographic Imaging (ECGI) has been an active research area for decades. Important progress was achieved but there are still challenges in robustness and accuracy due to the ill-posedness of the classical formulation.

In the last few years, deep learning (DL) based methods have been used to solve inverse problems, e.g. in medical image reconstruction [7]. Autoencoders are popular for these problems, as they are specifically designed to reconstruct high dimensional data in an unsupervised fashion. Autoencoders learn an identity function in order to reconstruct the input image after having first encoded it in a latent representation and then decoded it to the original input. Such methods have been recently introduced into the ECGI problem to regularize the temporal information while processing it [4].

In this manuscript, we propose to reformulate the whole ECGI problem as a conditional variational autoencoder based on convolutional neural networks. This has four main advantages:

2 Authors Suppressed Due to Excessive Length

- Spatio-temporal correlations: the convolutional model learns interactions in space and time between signals, while most ECGI methods solve each time step independently.
- Imaging substrate: the correlation between the substrate from imaging and the signals is also learned, therefore we can seamlessly integrate any 3D image information in ECGI, while this is still difficult in the classical formulation.
- Data-driven regularisation: using a generative model from a low dimensional space should ensure smooth variations between similar cases; this could alleviate the ill-posedness problem.
- Fast computations: once trained, DL methods are very fast to evaluate.

In order to achieve this, we use Cartesian coordinates for all the data (space and time) to leverage the power of convolutional neural networks.

2 Context

2.1 Electrocardiographic Imaging

ECGI is a non invasive modality which aims to better understand the electrical activity of the heart, both quantitatively and qualitatively. ECGI allows the visualisation of the electrical potential distribution of the electrical wave on the heart surface from body surface potentials (BSP). It requires medical imaging to obtain geometrical information and methods to solve the inverse problem.

Forward Problem This refers to the estimation of the ECG data from cardiac data. The two classical numerical approaches for this are the Boundary Element Method (BEM), based on surfaces, and the Finite Element Method (FEM), where the 3D torso model is approximated by small volume elements. Both propagate the epicardial action potentials to the body surface with chosen boundary conditions, *e.g.*, null current across the body surface. There are also methods using a dipole formulation, assuming that the torso domain is homogenous and infinite [3]. In [11], the authors demonstrated that noninvasive ECGI reconstruction does not require first order approximations for torso heterogeneities.

Inverse Problem This allows to reconstruct cardiac electrical activity using BSP. The classical approach is based on a transfer matrix between epicardial potentials and the torso potentials. To compute this transfer matrix different approaches can be used such as BEM [2] or FEM [13]. However, inverting it is ill-posed [13]. Therefore, different formulations and regularization methods were proposed, see for instance the publications of the ECGI consortium³.

³ <http://www.ecg-imaging.org>

2.2 Deep Learning

Variational Autoencoders Variational Autoencoders (VAEs) are powerful probabilistic generative models [9]. VAEs consist of two connected networks: an encoder and a decoder. The former takes an input x and compresses it into a low-dimensional latent representation, which variables are denoted z , with a distribution $P(z)$, often assumed to be a centred isotropic multivariate Gaussian $P(z) = N(z : 0, I)$. The latter takes z as input and reconstructs the data from the generative distribution $P_{\Theta}(x|z)$. This likelihood distribution $P_{\Theta}(x|z)$ is learned in the decoder neural network with parameters Θ . The resulting generative process induces the distribution $P_{\Theta}(x) = E_{P(z)} P_{\Theta}(x|z)$. The term Variational in Variational Autoencoder refers to variational inference or variational Bayes. Due to the intractability of the true posterior distribution $P(z|x)$, the posterior is approximated by learning the probability $Q_{\Phi}(z|x)$ in the encoder neural network with parameters Φ .

Fig. 1. VAE model [9]

A VAE is trained in order to minimize the Kullback-Leibler (KL) divergence between the variational probability $Q_{\Phi}(z|x)$ and the true posterior distribution $P(z|x)$. As this is intractable, it can be reformulated as the evidence lower bound (ELBO) of the log marginalized data likelihood $P_{\Theta}(x)$ [9]:

$$\log P_{\Theta}(x|z) - D_{KL}(Q_{\Phi}(z|x) \parallel P(z)) \leq \log P_{\Theta}(x)$$

The VAE loss function is then defined by a reconstruction term and the KL divergence between variational and prior probability:

$$L_{VAE}(\Theta, \Phi) = -\log P_{\Theta}(x|z) + D_{KL}(Q_{\Phi}(z|x) \parallel P(z))$$

By minimizing the loss function, the lower bound of the probability of generating real data samples is maximised. To compute the gradient of the variational lower bound, we use the reparameterization trick with respect to VAE [9]. Commonly, $Q_{\Phi}(z|x)$ is Gaussian with a diagonal covariance matrix: $z \sim Q_{\Phi}(z|x) = N(z : \mu, \sigma)$ where $z = \mu + \sigma \odot \epsilon$ and $\epsilon \sim N(0, I)$.

Conditional Variational Autoencoders Conditional VAEs (CVAEs) [10] are an extension of VAEs [9] where the latent variables and the data are both conditioned on additional random variables c . The encoder of the CVAE is not only conditioned on the data x but also on the conditioning data c which results in the variational probability: $Q_{\Phi}(z|x, c)$. Respectively, the decoder is also conditioned on the conditioning data c . Thus, the generative distribution becomes: $P_{\Theta}(x|z, c)$. The CVAE is trained to maximize the conditional log-likelihood where the ELBO is:

$$L_{CVAE}(\Theta, \Phi) = -\log P_{\Theta}(x|z, c) + D_{KL}(Q_{\Phi}(z|x, c) \parallel P(z)) \leq \log P_{\Theta}(x|c)$$

4 Authors Suppressed Due to Excessive Length

β -Variational Autoencoders β -VAEs are an evolution of VAEs which intends to discover disentangled latent factors. In [8], it is shown that this model achieves similar disentanglement performance compared with VAEs, both quantitatively and qualitatively. β -VAEs and VAEs have the same goals: maximize the probability of generating real data and minimize the distance between the real and estimated posterior distributions (smaller than a constraint ε). In this approach, the prior is an isotropic Gaussian $P(z) = N(0, I)$.

The β -VAE lower bound is defined as:

$$L_{\beta\text{-VAE}}(\Theta, \Phi, \beta) = -\log P_{\Theta}(x|z) + \beta D_{KL}(Q_{\Phi}(z|x) \parallel P(z))$$

with β a Lagrangian multiplier hyperparameter under the Karush-Kuhn-Tucker conditions.

If $\beta = 1$, β -VAEs correspond to the original VAEs [9]. If $\beta > 1$, β -VAEs apply a stronger constraint on the latent bottleneck and so limit the capacity of z . It allows to learn the most efficient representation of the data. However, if β is too big, β -VAE learns an entangled latent representation because of its excessive capacity in the latent z bottleneck. Same remarks if β is too small, it will have too little capacity. To sum up, $\beta > 1$ is capital to achieve good disentanglement.

3 Methods

3.1 ECGI Forward Problem: Data Simulation

We simulated cardiac activation maps and BSP data using the Eikonal Model directly on a Cartesian grid from image segmentation [1]. The Eikonal model is a fast generic model of wave front propagation. Its inputs are the myocardial wall mask, a local conduction velocity v for each voxel x of the given wall mask and the pacing zone.

$$v(x) \|\nabla T(x)\| = 1$$

with $T(x)$ the local activation time in x . It is solved using The Fast Marching Method (FMM) [12].

BSP were generated using the dipole formulation [6] associating each activation time from the Eikonal model with an action potential signal from the Mitchell Schaeffer model. 100 torso electrodes were positioned on a 10×10 Cartesian grid in front of the heart.

In total, 120 activation maps and the corresponding BSP were simulated using a patient image segmentation from CT images. In order to accelerate computation and ease memory requirements, we currently present 2D results, where we separated the 3D simulated data in 2D slices.

3.2 Models Structure and Training

The different models described below were trained on 80% of the data, tested on the remaining 20% and have been implemented using Keras.⁴

⁴ <https://keras.io/>

VAE for BSP It follows the architecture presented in Fig. 2, with convolutional layers in order to extract spatiotemporal correlations.

Fig. 2. Variational Autoencoder architecture

The encoder consists of two convolutional layers followed by one dense layer. The bottleneck layers (μ, σ, z) are fully-connected. The decoder consists of a fully-connected and three transposed convolutional layers. For all convolutional layers except the output one: strides are set to 2, 16 filters are applied, the kernel size is 3 and the activation functions are *ReLU*. In the last layer, one filter is applied with a kernel size of 3 and a sigmoid activation function. The latent code size is set to 6. Mean squared error (MSE) is used to calculate the reconstruction loss. The optimization is performed with the NADAM solver with a batch size of 32.

VAE for Activation Maps The encoder architecture is the same as the previous one. However, network parameters and activations differ: the convolution kernel size is 3 (even in the last layer) and all activations are *tanh*. The latent code size is set to 6. MSE is used to calculate the reconstruction loss. Optimization is performed with the RMSPROP solver with a batch size of 25.

β -CVAE for Activation Maps from BSP and Images Our conditional autoencoder aims to generate heart activation maps using two conditions: the simulated ECG signals and the shape of the patient's heart. This is a novel formulation of ECGI using DL in order to learn the influence of cardiac shape/structure.

The architecture of our conditioned generative model (encoder) and our conditioned variational approximation (decoder) is described in Fig. 3. Empirically, we found that the best value for Lagrangian multiplier hyperparameter β in the loss function was 5. The ADAM optimiser was used with 10^{-4} as learning rate and a batch size of 1. Latent space is set to 16.

6 Authors Suppressed Due to Excessive Length

Fig. 3. CVAE architecture for ECGI with imaging data (C_1 : The first convolution layer in the encoder was conditioned by concatenating the input data with the BSP mapping signals and the myocardial mask. C_2, C_3 : The deconvolution layers in the decoder were conditioned by concatenating each layer's output with sub-sampled versions of the BSP mapping signals and of the mask).

4 Results

4.1 Evaluation of Body Surface Potentials VAE

Fig. 4 shows two examples of input BSP mapping signals and their corresponding generated signals by our model.

Fig. 4. Evaluation of input (red) and output (blue) signals from VAE for 2 different electrodes.

The performance of our model was evaluated measuring 5 metrics between the input BSP signals and their corresponding generated signals. We found a correlation of 0.95, a mean difference of 3.345%, an abscisse area difference of -2.152% and a maximum amplitude difference of -3.468%; 90.854% of the time, the sign of first peak was the same in the input and in the output.

4.2 Evaluation of Activation Maps VAE

Using our VAE model developed to generate the electrical activation of the heart (Fig. 5), we obtained a MSE of 0.018 for reconstruction accuracy.

Fig. 5. Predicted (left) and simulated (right) activation maps with the Activation Maps VAE model.

4.3 Evaluation of β -CVAE

β -CVAE is a probabilistic generative method, so we can generate several probable solutions for a given input. We generated ten activation maps per prediction in order to evaluate the accuracy and stability of our method. The MSE metric was 0.095 over all the tests. Fig. 6 shows an example of (a) an input activation map, (b) its corresponding generated mean activation map by our β -CVAE proposed method, (c) standard deviation map for 10 predictions, and (d) the error map.

Fig. 6. (a) Simulated and (b) predicted mean activation maps for proposed deep learning based ECGI, (c) Standard deviation map calculated over 10 predictions, (d) error map, difference between predicted and simulated activation maps.

In Fig. 6 (c) and (d), small values imply that the reconstruction performs well (small error and small standard deviation), while large values mean that the reconstruction is suffering. We can observe that the areas with the highest

8 Authors Suppressed Due to Excessive Length

standard deviation are close to areas with the highest error, therefore the probabilistic aspect of our method can help in quantifying the uncertainty in the predictions.

5 Discussion & Conclusion

As a direct application of the Forward Problem of ECGI, electrical activation and potentials of the heart were simulated using a patient CT-scan image and the Eikonal Model. The first step of our work was to understand and choose the best hyperparameters of the Variational Autoencoders to generate potential or activation maps. Finally, we proposed a novel method based on Conditional β Variational Autoencoder able to solve ECGI inverse problem in 2D. This generative probabilistic model learns geometrical and spatio-temporal information and enables to generate the corresponding activation map of the specific heart. We showed that these generated electrical activities were very similar to the simulated ones. The presented method will now be generalised to 3D and evaluated on clinical data. Theoretically there is no impediment to extend our 2D model to 3D, as all the convolution operators have a 3D version. However, it will require to simulate more data to train the model, as we will have more hyperparameters to optimise. We will also explore transfer learning to apply it on clinical data [5].

Acknowledgements

The research leading to these results has received European funding from the ERC starting grant ECSTATIC (715093) and French funding from the National Research Agency grant IHU LIRYC (ANR-10-IAHU-04).

References

1. Cedilnik, N., Duchateau, J., Dubois, R., Sacher, F., Jaïs, P., Cochet, H., Sermesant, M.: Fast Personalized Electrophysiological Models from CT Images for Ventricular Tachycardia Ablation Planning. *EP-Europace* **20** (Nov 2018)
2. Chamorro-Servent, J., Dubois, R., Potse, M., Coudière, Y.: Improving the spatial solution of electrocardiographic imaging: A new regularization parameter choice technique for the tikhonov method. In: *FIMH* (2017)
3. Chavez, C., Zenzemi, N., Coudière, Y., Alonso-Atienza, F., Alvarez, D.: Inverse Problem of Electrocardiography: estimating the location of cardiac ischemia in a 3D geometry. In: *Functional Imaging and modelling of the heart (FIMH2015)*. vol. 9126. Springer (2017)
4. Ghimire, S., Dhamala, J., Gyawali, P.K., Sapp, J.L., Horacek, M., Wang, L.: Generative modeling and inverse imaging of cardiac transmembrane potential. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*. pp. 508–516. Springer, Cham (2018)
5. Giffard-Roisin, S., Delingette, H., Jackson, T., Webb, J., Fovargue, L., Lee, J., Rinaldi, C.A., Razavi, R., Ayache, N., Sermesant, M.: Transfer Learning from Simulations on a Reference Anatomy for ECGI in Personalised Cardiac Resynchronization Therapy. *IEEE Transactions on Biomedical Engineering* **20** (2018)

6. Giffard-Roisin, S., Jackson, T., Fovargue, L., Lee, J., Delingette, H., Razavi, R., Ayache, N., Sermesant, M.: Non-Invasive Personalisation of a Cardiac Electrophysiology Model from Body Surface Potential Mapping. *IEEE Transactions on Biomedical Engineering* **64**(9), 2206 – 2218 (Sep 2017)
7. Hammernik, K., Klatzer, T., Kobler, E., Recht, M.P., Sodickson, D.K., Pock, T., Knoll, F.: Learning a variational network for reconstruction of accelerated mri data. *Magnetic Resonance in Medicine* **79**(6), 3055–3071 (2018)
8. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: β -vae: Learning basic visual concepts with a constrained variational framework. In: *International Conference on Learning Representations* (2017)
9. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2014)
10. Kingma, D.P., Mohamed, S., Jimenez Rezende, D., Welling, M.: Semi-supervised learning with deep generative models. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems 27*, pp. 3581–3589. Curran Associates, Inc. (2014)
11. Ramanathan, C., Rudy, Y.: Electrocardiographic imaging: Effect of torso inhomogeneities on noninvasive reconstruction of epicardial potentials, electrograms, and isochrones. *Journal of cardiovascular electrophysiology* **12**, 241–52 (03 2001)
12. Sermesant, M., Coudière, Y., Moreau-Villéger, V., Rhode, K., Hill, D., Razavi, R.: A fast-marching approach to cardiac electrophysiology simulation for XMR interventional imaging. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005* pp. 607–615 (2005)
13. Zemzemi, N., Dobrzynski, C., Bear, L., Potse, M., Dallet, C., Coudière, Y., Dubois, R., Duchateau, J.: Effect of the torso conductivity heterogeneities on the ECGI inverse problem solution. In: *Computing in cardiology. Nice, France* (Sep 2015)