Deep Learning Formulation of ECGI Integrating Image & Signal Information with Data-driven Regularisation

Tania Bacoyannis¹, Buntheng Ly¹, Nicolas Cedilnik^{1,2}, Hubert Cochet²,
and Maxime Sermesant^{1,2}

¹*Inria, Université Côte d’Azur, Epione team, Sophia Antipolis, France*

²*IHU Liryc, Univ. Bordeaux, Bordeaux, France*

Abstract

Aims Electrocardiographic Imaging (ECGI) is a promising tool to map the electrical activity of the heart non-invasively using body surface potentials (BSP). However, it is still challenging due to the mathematically ill-posed nature of the inverse problem to solve. Novel approaches leveraging progress in artificial intelligence could alleviate these difficulties.

Methods We propose a Deep Learning (DL) formulation of ECGI in order to learn the statistical relation between BSP and cardiac activation. The presented method is based on Conditional Variational Autoencoders (CVAE) using deep generative neural networks. To quantify the accuracy of this method, we simulated activation maps and BSP data on six cardiac anatomies.

Results We evaluated our model by training it on five different cardiac anatomies (5 000 activation maps) and by testing it on a new patient anatomy over 200 activation maps. Due to the probabilistic property of our method, we predicted 10 distinct activation maps for each BSP data. The proposed method is able to generate volumetric activation maps with a good accuracy on the simulated data: the mean absolute error is 9.40 ms with 2.16 ms standard deviation on this testing set.

Conclusion The proposed formulation of ECGI enables to naturally include imaging information in the estimation of cardiac electrical activity from body surface potential. It naturally takes into account all the spatio-temporal correlations present in the data. We believe these features can help improve ECGI results.

Keywords: Electrocardiographic Imaging; Inverse Problem; Deep learning; Computational Modelling; Generative Model

What's New

- The proposed deep learning formulation of ECGI leverages the available cardiac imaging to better reconstruct the spatio-temporal electrical activity of the heart.
- This conditional generative model is able to learn spatio-temporal correlations as well as the influence of cardiac shape/structure through convolutions.
- As a deep learning model, its objective is to generalise the learned aptitude in order to predict accurate new activation for unseen data.
- Using a generative model from a low dimensional space should ensure smooth variations between similar cases; this could alleviate the ill-posedness issue and enables uncertainty quantification.
- The new ECGI formulation inherits the benefits of deep learning methods: computing power and memory capacity. Once trained, our model is fast to evaluate.

Introduction

Cardiac arrhythmias often require invasive catheter measurements in order to precisely diagnose the pathology and plan the therapy. This information on the local electrical activity in the myocardium can theoretically be obtained non-invasively through electrocardiographic imaging (ECGI). ECGI allows a panoramic three-dimensional visual reconstruction of the electrical activity of the heart from body surface potentials (BSP), see Fig. 1. This technique requires anatomical information of the heart and torso from 3D imaging, BSP measured with electrodes and their localisation on the torso surface, and mathematical methods to solve the inverse problem. ECGI has been an active research area for decades. The first sketch has been proposed in 1977 [1] and its history

and clinical applications were recently reviewed [2]. Important progress was achieved but there are still scientific challenges due to the ill-posedness of the classical mathematical formulation. Therefore, this emerging modality is still at its infancy in clinical practice due to its currently limited accuracy and resolution.

Figure 1: ECGI process. Using 1.a) cardiac imaging and 2.a) signal processing, we obtain information on 1.b) the cardiac geometry and the position of the electrodes used to record the ECG signals, and 2.b) the surface ECG signals. Both are then used by 3.a) numerical methods to reconstruct the 3.b) cardiac electrical activity.

The classical approach is based on a transfer matrix between epicardial potentials and the torso potentials. The forward problem of ECGI refers to the computation of the electrical potential on the torso surface from the cardiac electrical activity. To compute this transfer matrix different approaches were proposed such as: the Boundary Element Method (BEM) [3] based on surfaces, or the Finite Element Method (FEM) [4] whereas the 3D torso anatomy is approximated by small volume elements. Both methods propagate the epicardial action potentials to the body surface with chosen boundary conditions, *e.g.* null current across the body surface. In [5], the authors demonstrated that ECGI reconstruction does not require the detailed inclusion of torso heterogeneities.

However, the inverse problem is ill-posed [4] which means that its solution is not unique, and that its resolution will be very sensitive to noise. Small errors in input will cause large errors in the output. Therefore, a variety of formulations and regularisation methods were proposed to stabilise this problem, see for instance the publications of the

ECGI consortium¹. Recently, 15 different algorithms to solve the inverse problem were compared and the conclusion was that each of them had benefits which varied according to bi-/ , left-/ and right-ventricular pacing [6]. Moreover, to solve the ECGI inverse problem, it is possible to include a multitude of physiological, physical and mathematical priors for the forward and inverse problem but also for describing the cardiac sources [7] as well as to perform the ECG signal processing [8]. An overview of most of these methods was published recently [9].

However, ECGI could achieve more accurate and robust results by leveraging all the information present in the different imaging modalities acquired, as well as by introducing prior knowledge from previous cases. This is what we propose in this manuscript, through the use of deep learning and generative models.

Deep learning (DL) has generated a huge interest in healthcare applications and has become the standard method for solving a large variety of computer vision problems. Moreover, these last years DL-based methods have proven their ability to solve inverse problems, e.g. in medical image reconstruction [10]. Recently, [11] drew up a state-of-the-art of DL approaches for solving inverse problems, and has showed the benefits of such data-driven methods. [12] introduced autoencoders into ECGI for the regularisation of temporal information.

Our new numerical method is based on DL and autoencoders in order to solve the whole inverse problem and reconstruct accurate 3D activation maps. A preliminary version of this method in 2D and on a single geometry was presented in [13].

In this manuscript, we propose to reformulate the whole ECGI problem as a β -conditional variational autoencoder based on convolutional neural networks.

Using a DL model as autoencoder trained on a large database, i.e activations with different pacing locations, allows to get more robust methods and reduces the prediction time. The developed method was used to predict activation maps for different cardiac geometries with their corresponding simulated BSP. The model was trained and tested on a simulated database using a personalised cardiac simulation pipeline [14]. Compared with

¹<http://www.ecg-imaging.org>

the classical mathematical formulation of ECGI, this approach has five main advantages:

- Spatio-temporal correlations: the convolutional model learns interactions in space and time between signals, while most ECGI methods solve each time step independently or use a temporal prior without spatial correlations.
- Imaging substrate: the correlation between the substrate from imaging (i.e. myocardial scar) 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 issue.
- Volumetric predictions: this 3D convolutional approach generates volumetric activation maps within the whole myocardium, not limited to the epicardium and the endocardium.
- Fast computations: once trained, DL methods are very fast to evaluate.

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

In the following section, we will first present our DL model, then the method used to simulate activation maps from CT images and we will describe the computation details of our model.

Methods

Deep Learning Approach

We base our reformulation of ECGI on a probabilistic generative model, namely a Conditional Variational AutoEncoder (CVAE) [15, 16]. A CVAE is an extension of a Variational AutoEncoder (VAE) [17] which is itself a powerful probabilistic generative version of autoencoders [17].

An autoencoder is an unsupervised learning technique which compresses the relevant information of the input data into a low-dimensional latent code in order to reconstruct output data with the least possible amount of difference with the input. This dimensionality reduction allows to capture the more important features of the input data.

In a Variational AutoEncoder (VAE), the term Variational refers to variational inference or variational Bayes used to approximate complex distributions. As the autoencoder, the VAE consists 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 prior 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 $P_{\Theta}(x) = E_{P(z)}P_{\Theta}(x|z)$. Due to the intractability of the true posterior distribution $P(z|x)$, the posterior is approximated by learning a simpler probability $Q_{\Phi}(z|x)$ in the encoder neural network with parameters Φ . A VAE is trained in order to maximise the probability of generating real data samples, by maximising the Evidence Lower Bound (ELBO). The ELBO is composed by two terms: the Kullback-Leibler (KL) divergence which is a measure between the variational probability $Q_{\Phi}(z|x)$ and the true posterior distribution $P(z|x)$; and the log likelihood of $P_{\Theta}(x)$ under the probability distribution $P_{\Theta}(x|z)$. By minimising the loss function, the ELBO is maximised.

In a Conditional Variational AutoEncoder (CVAE), the input data is conditioned on an additional random variable c that modulates the prior on the latent vector. The encoder of the CVAE has therefore the data x but also the conditioning data c as input, which results in the variational probability: $Q_{\Phi}(z|x, c)$. Respectively, the decoder also has as input the conditioning data c . Thus, the generative distribution becomes: $P_{\Theta}(z|x, c)$. Based on VAE, the CVAE is also trained to minimise the loss function:

$$L_{CVAE}(\Theta, \Phi) = -\log P_{\Theta}(x|z, c) + D_{KL}(Q_{\Phi}(z|x, c) \parallel P(z)) \quad (1)$$

Commonly, $Q_{\Phi}(z|x, c)$ is Gaussian with a diagonal covariance matrix: $z \sim Q_{\Phi}(z|x, c) =$

$N(z : \mu, \sigma)$ where: $z = \mu + \sigma \odot \epsilon$, $\odot$ denotes the Hadamard product and $\epsilon \sim N(0, I)$.

A β -VAE intends to discover disentangled latent factors by adding a penalty coefficient on the KL divergence. In [18], it is shown that with $\beta > 1$ the model achieves better disentanglement performance than classical VAEs.

We use a β -CVAE which loss function is therefore defined by a reconstruction term $P_{\Theta}(x|z, c)$ and a weighted KL divergence between variational and prior probabilities term:

$$L_{\beta\text{-CVAE}}(\Theta, \Phi, \beta) = -\log P_{\Theta}(x|z, c) + \beta D_{KL}(Q_{\Phi}(z|x, c) \parallel P(z)) \quad (2)$$

with β the penalty coefficient.

Artificial Neural Network Architecture

The proposed method aims to generate cardiac activation maps using two constraints: the 3D imaging data of the patient heart and the BSP signals. Therefore we use a multimodal input β -CVAE, which consists of three elements: the encoder, the latent vector and the decoder. The architecture of our deep learning model is described in Fig. 2.

First, the encoder takes as input data the normalised activation maps x and as conditioning data c the BSP signals and the cardiac image. As these data come from different modalities, BSP and imaging data do not have the same dimension. To obtain the same shape for the signals, we applied two successive 2D convolutional layers. Then these three input are concatenated to be used as a single input to the encoder model. This enables the network to directly learn spatiotemporal correlations within these input. The encoder consists of two successive convolutional layers (see subsection *Computation Details* for the implementation details) following the grey arrows, which compress the input.

The 3D output of these convolutional layers is flattened to 1D before being fed to a dense layer to further compress the output vector into 16 dimensions. We then generate from it the multivariate Gaussian means and standard deviations μ and σ . In order to sample the latent vector z we used the Gaussian reparameterisation method.

The decoder network consists of a main branch which takes the vector z as input, and the conditional branches, where the conditional input c is integrated. The main

Figure 2: β -CVAE model architecture. Training: input data is BSP mapping signals concatenated with myocardial mask and encoded in z . Activation maps are decoded from z with conditioning by sub-sampled versions of the BSP mapping signals and of the cardiac mask. Prediction: decoder is used to generate myocardial electrical activation maps from a new cardiac segmentation and its corresponding BSP mapping signals.

branch applies a dense layer to the input z , before reshaping the 1D vector to a 3D matrix. Then, we use the combination of transpose and normal convolutional layers at each level to increase the output dimension back to its original size. At each level, we use two convolutional layers (strided followed by non-strided) on the condition inputs c to generate the condition input in the same shape as the output from the main branch. Then, we concatenate the condition input with the main branch output. At the end of the decoder, we multiply the output with the myocardial image binary mask M to only provide predictions within the myocardial wall.

Supervised Learning Loss Function

During the training of the CVAE, we use a modified version of the Root Mean Squared Error (RMSE), Eq. (3), to calculate the reconstruction loss.

In order to put more emphasis on the small values, i.e the earliest activation times of

the normalised propagation map A , we reverse the order of the values by using $2M - A$ as input. We also only calculate the reconstruction error in the myocardial binary mask M , named RMSE_M . Therefore we use a weighted RMSE formula RMSE_w defined as follow:

$$\text{RMSE}_w(A) = \text{RMSE}(A, 2M - A) = \sqrt{\text{MSE}(A) \times (2M - A)} \quad (3)$$

There are two terms used in the reconstruction loss. First, we calculate a reconstruction loss between the encoded data from the z -latent vector and the ground truth activation maps L_z , Eq. (4), to get a more consistent loss function:

$$L_z = \text{RMSE}_w(x, P(x|z, c)), \text{ where } z = Q(z|x, c) \quad (4)$$

Second, we add a reconstruction loss from the multivariate Gaussian $L_{\bar{z}}$, Eq. (5), to improve the generative aspect:

$$L_{\bar{z}} = \text{RMSE}_w(x, P(x|\bar{z}, c)), \text{ where } \bar{z} \sim \mathcal{N}(0, 1) \quad (5)$$

$P(x|\bar{z}, c)$ can be generated independently from the input activation map, x , while it integrates the information from the conditions c . This can be used to generate activation maps from ECG signals and cardiac masks only.

The total loss function in our model is a sum of the two different RMSE_w and the KL divergence loss $L_{\beta\text{-CVAE}}$. Therefore, the total loss function to minimise when training our model is:

$$L_{ECGI} = L_{\beta\text{-CVAE}} + L_z + L_{\bar{z}} \quad (6)$$

Synthetic Data Generation: ECGI Forward Problem

In order to evaluate the proposed methodology, we built a synthetic database from six cardiac anatomies based on CT-scans provided by the CHU Bordeaux, France, see Fig. 3. All six patients showed post-infarction scar resulting in severe thinning of the left ventricular wall. We generated synthetic electrical data with the forward problem of electrocardio-

graphy. In our case, we use a dipole formulation, assuming that the torso domain is homogeneous and infinite [19].

Figure 3: Five cardiac anatomies from CT imaging used for training and one (*) for testing. The thickness maps show the infarct scars.

Figure 4: Synthetic data simulation process: myocardial wall thickness from patient image is used to parameterise the Eikonal model that generates activation maps (LAT: local activation time). The Mitchell-Schaeffer model with the dipole formulation are used to generate BSP mapping signals.

We proposed a fast method to simulate the activation maps of the heart and their

corresponding body surface potentials, Fig. 4c. It is based on the Eikonal Model on a Cartesian grid from image segmentation [20]. The Eikonal model is a fast generic model of wave front propagation; its results in terms of local activation times are comparable to more sophisticated models [21] and it is efficiently computed with the Fast Marching method (FMM) [22]. The Eikonal model takes as input the myocardial wall mask, a local conduction speed s for each voxel v of the given wall mask and the pacing zone. The model equation writes:

$$s(v) \|\nabla T(v)\| = 1 \quad (7)$$

with $T(v)$ the local activation time in voxel v .

For these simulations, we vary the local conduction velocity according to the local myocardial wall thickness, as chronic infarcts generate thinner walls with slower conduction velocity; this approach has been shown to efficiently reproduce intra-cardiac activation maps recordings [20]. For each simulation, to initialise the wave front propagation, we chose randomly a single pacing onset where the heart had a wall thickness superior to 6 mm. We defined a stopping value for the propagation at 250 ms.

To numerically compute the Forward problem, i.e estimate the BSP from the cardiac data, we first model the 100 torso electrodes by a 10×10 Cartesian grid on the front-side of the torso. We consider that the torso-domain is homogeneous and infinite. To generate the BSP we use the dipole formulation from the volume conductor theory where every element of the cardiac volume is a time varying current dipole [23]. Moreover, each activation time obtained with the Eikonal model is associated with an action potential from the cardiac cell Mitchell-Schaeffer model [24]. The local current density j_{eq} is written as:

$$j_{eq} = -\sigma_E \cdot \nabla V = -\sigma_E \cdot \frac{\partial V}{\partial t} \frac{\partial t}{\partial x} = -\sigma_E \cdot \frac{\partial V}{\partial t} \nabla T \quad (8)$$

with σ_E the local electrical conductivity, ∇V the spatial gradient of the potential V and ∇T the gradient of the activation map given by the Eikonal model [14]. We can therefore directly compute the local dipole from the gradient of the activation map and the time derivative of a transmembrane potential model (Mitchell-Schaeffer).

In total, we simulated 6 000 activation maps and the corresponding BSP using 6 patient image segmentations from CT images. We resampled the activation map and the myocardial wall to have the same dimensions $100 \times 100 \times 100$. The signal input is kept at $10 \times 10 \times 250$.

Computational Details

The model was trained on 5 patient geometries, see Fig. 4a, including a total of 5 000 activation maps and their corresponding BSP, which were divided into 8:2 ratio for training and validation datasets. We then tested the trained model on 200 activation maps and their corresponding BSP from a new geometry unseen during the training Fig. 4b. The neural network design, training and testing were implemented using Keras API of Tensorflow 2.0 ².

To adapt the signal shape, we used two 3×3 2D convolutional layers with the number of strides of 2 and 5 and the number of output features of 200 and 100, respectively. The encoder of our network consists of two successive stride 2 convolutional layers (Fig. 2), one flatten layer followed by one dense layer. Each layer has a $3 \times 3 \times 3$ kernel and all the activation functions are *tanh*. The bottleneck layers (μ, σ, z) are fully-connected. From different tests, the latent code size was set to 25. The decoder has a fully-connected layer which is then reshaped into a 3D matrix before applying a succession of transpose and normal convolutional layers. The model was optimised using the total loss described in Eq. 6, without any additional weight applied to the loss terms. The best value for β , the penalisation term in the loss function, was 10. We used Adam [25] optimiser with a learning rate of 10^{-4} , with a batch size of one for 240 epochs. The computation time for training was 64 hours on a NVidia Tesla V100 (32GB of RAM).

²https://www.tensorflow.org/api_docs/python/tf/keras

Results

When testing the method, our probabilistic generative model allows us to generate multiple solutions from a single input, by sampling the latent space. To compare our predictions with the ground truth activation, we generated 10 different activation maps for each case and computed the mean and standard deviation of these predictions. Fig. 5 shows a 3D visualisation and a slice of the estimated mean activation map over the 10 predictions (b) in comparison with the ground-truth simulated cardiac electrical activity (a). (c) represents the standard deviation map for these 10 predictions, and (d) the error map between the generated and the ground truth maps.

Figure 5: (a): Ground Truth activation map, (b): Activation map generated by the presented method, (c): standard deviation map calculated over 10 predictions, (d): error-map, difference between predicted and simulated activation maps. All results are in 3D, showing the middle x-axis slice in the bottom of the figure.

We evaluated our model by testing it on a new patient anatomy over 200 activation maps. Due to the probabilistic property of our method, we predicted 10 distinct activation maps for each BSP data.

The proposed method is able to generate volumetric activation maps with a good accuracy on this simulated data: the mean absolute error is 11.33 ms with 4.10 ms standard deviation on this testing set.

Small errors and small standard deviation in Fig. 5c and Fig. 5d indicate that the reconstruction performs well and in a consistent way, while large values mean that the reconstruction is suffering. The zones with the highest error correspond partly with the zones with the highest standard deviation in the maps. This underlines the potential use of our probabilistic generative model to help in quantifying the uncertainty in the predictions.

From the representative results shown in Fig. 5 we can observe that the weight applied to the loss function influences the performance of the reconstruction in the later regions of the wave propagation. In opposition, zones where the propagation starts are better reconstructed.

Discussion

The proposed approach demonstrated great accuracy on this synthetic database with a good robustness to anatomical variations. However the goal of this work is to apply it clinically for diagnosis and therapy planning. So it is necessary to detail the limitations we have at this present stage in order to consider the required extensions.

To generate our database, we only considered single pacing sites. A first extension to our work will be to simulate data with multiple onsets. This is straightforward to implement, but the accuracy of the results has to be evaluated.

Second, we use the Eikonal model, which is a simplified model in terms of transmembrane potential capabilities in comparison to more complex models as the Bidomain one. However, we believe that this model can be used in many clinical cases. Indeed, the spatial dynamics of the Eikonal model and the reaction-diffusion equations are extremely similar [26]. The local conduction velocity modifications can represent the impact of infarcts on activation maps [20]. Thus, we believe that this fast method to simulate activation maps is relevant in a context of myocardial infarction. The limit of this model is reached for arrhythmias or pathologies modifying local transmembrane potential profile. From a more global and mathematical point of view, the Eikonal model provides a very strong prior on spatiotemporal correlations for solving the ECGI inverse problem.

Moreover, our method makes the hypothesis about isotropic propagation in the heart. Although heart anisotropy may not play an important role for ECGI because of the spatial scale at which it is working, we know that it has an important impact on local propagation front and thus in the forward problem. In the future, we are going to include the fibres in our forward problem formulation, which will be an easy extension.

In order to evaluate our model, we generated 10 activation maps for each case. It would be interesting for each one of these cases to vary the number of activation maps we generated and to estimate the effects on the accuracy of the model and the uncertainty quantification.

Moreover, this method has to be evaluated on real data. The differences between the training data and the testing data can severely impact performance. Robustness to signal noise and missing leads, a classical problem in ECGI, has still to be quantified. We believe that adding noise to our simulated database should help. We also intend to train the model on a wider variety of simulated data and also on some real data with partial ground-truth from catheter mapping. To help moving from synthetic to real data, we may use transfer learning [27].

Conclusion

Our proposed β -CVAE is able to generate a 3D volumetric map of the cardiac electrical activity for a new geometry from non-invasive BSP and cardiac imaging data. During the training, our model is designed to learn spatiotemporal as well as geometrical information. This novel formulation enables to directly introduce all the imaging information on cardiac substrate within the ECGI methodology and therefore fully exploits the available patient information.

This is the first time an ECGI method is able to reconstruct the activation pattern throughout the myocardium in a volumetric manner at such resolution, while integrating 3D substrate information from imaging and within a probabilistic framework. This promising formulation now has to be evaluated on real clinical data.

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] R. C. Barr, M. Ramsey, and M. S. Spach. Relating epicardial to body surface potential distributions by means of transfer coefficients based on geometry measurements. *IEEE Transactions on Biomedical Engineering*, BME-24(1):1–11, 1977.
- [2] Helder Pereira, Steven Niederer, and Christopher A Rinaldi. Electrocardiographic imaging for cardiac arrhythmias and resynchronization therapy. *EP Europace*, 08 2020. euaa165.
- [3] Judit Chamorro-Servent, Rémi Dubois, Mark Potse, and Yves Coudière. Improving the spatial solution of electrocardiographic imaging: A new regularization parameter choice technique for the tikhonov method. In Mihaela Pop and Graham A Wright, editors, *Functional Imaging and Modelling of the Heart*, pages 289–300, Cham, 2017. Springer International Publishing.
- [4] Nejib Zemzemi, Cecile Dobrzynski, Laura Bear, Mark Potse, Corentin Dallet, Yves Coudière, Remi Dubois, and Josselin Duchateau. Effect of the torso conductivity heterogeneities on the ECGI inverse problem solution. In *Computing in cardiology*, Nice, France, September 2015.
- [5] Charu Ramanathan and Yoram Rudy. Electrocardiographic imaging: Effect of torso inhomogeneities on noninvasive reconstruction of epicardial potentials, electrograms, and isochrones. *Journal of cardiovascular electrophysiology*, 12:241–52, 03 2001.

- [6] Amel Karoui, Laura Bear, Pauline Migerditichan, and Nejib Zemzemi. Evaluation of fifteen algorithms for the resolution of the electrocardiography imaging inverse problem using ex-vivo and in-silico data. *Frontiers in Physiology*, 9:1708, 2018.
- [7] A Van Oosterom and V Jacquemet. A parameterized description of transmembrane potentials used in forward and inverse procedures. *Folia Cardiologica*, 12(suppl D):111, 2005.
- [8] Laura R Bear, Y Serinagaoglu Dogrusoz, Jana Svehlíková, Jaume Coll-Font, Wilson Good, Eelco van Dam, R MacLeod, Emma Abell, R Walton, Ruben Coronel, et al. Effects of ecg signal processing on the inverse problem of electrocardiography. In *2018 Computing in Cardiology Conference (CinC)*, volume 45, pages 1–4. IEEE, 2018.
- [9] Matthijs Cluitmans, Dana H. Brooks, Rob MacLeod, Olaf Dössel, María S. Guillem, Peter M. van Dam, Jana Svehlikova, Bin He, John Sapp, Linwei Wang, and Laura Bear. Validation and opportunities of electrocardiographic imaging: From technical achievements to clinical applications. *Frontiers in Physiology*, 9:1305, 2018.
- [10] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated mri data. *Magnetic Resonance in Medicine*, 79(6):3055–3071, 2018.
- [11] Simon Arridge, Peter Maass, Ozan Öktem, and Carola-Bibiane Schönlieb. Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174, 2019.
- [12] Sandesh Ghimire, Jwala Dhamala, Prashnna Kumar Gyawali, John L. Sapp, Milan Horacek, and Linwei Wang. Generative modeling and inverse imaging of cardiac transmembrane potential. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pages 508–516, Cham, 2018. Springer.
- [13] Tania Bacoyannis, Julian Krebs, Nicolas Cedilnik, Hubert Cochet, and Maxime Sermesant. Deep learning formulation of ecgi for data-driven integration of spatiotemporal correlations and imaging information. In Yves Coudière, Valéry Ozenne, Ed-

- ward Vigmond, and Nejib Zemzemi, editors, *Functional Imaging and Modeling of the Heart*, pages 20–28, Cham, 2019. Springer International Publishing.
- [14] Nicolas Cedilnik, Josselin Duchateau, Frédéric Sacher, Pierre Jaïs, Hubert Cochet, and Maxime Sermesant. Fully automated electrophysiological model personalisation framework from ct imaging. In Yves Coudière, Valéry Ozenne, Edward Vigmond, and Nejib Zemzemi, editors, *Functional Imaging and Modeling of the Heart*, pages 325–333, Cham, 2019. Springer International Publishing.
- [15] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, pages 3581–3589. Curran Associates, Inc., 2014.
- [16] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28, pages 3483–3491. Curran Associates, Inc., 2015.
- [17] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *In: Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.
- [18] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. β -vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.
- [19] Carlos Eduardo Chávez, Nejib Zemzemi, Yves Coudière, Felipe Alonso-Atienza, and Diego Álvarez. Inverse problem of electrocardiography: Estimating the location of cardiac ischemia in a 3d realistic geometry. In Hans van Assen, Peter Bovendeerd, and Tammo Delhaas, editors, *Functional Imaging and Modeling of the Heart*, pages 393–401, Cham, 2015. Springer International Publishing.

- [20] Nicolas Cedilnik, Josselin Duchateau, Rémi Dubois, Frédéric Sacher, Pierre Jaïs, Hubert Cochet, and Maxime Sermesant. Fast Personalized Electrophysiological Models from CT Images for Ventricular Tachycardia Ablation Planning. *EP-Europace*, 20, November 2018.
- [21] M. Potse, B. Dube, J. Richer, A. Vinet, and R. M. Gulrajani. A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart. *IEEE Transactions on Biomedical Engineering*, 53(12):2425–2435, 2006.
- [22] Maxime Sermesant, Yves Coudière, Valérie Moreau-Villéger, Kawal Rhode, Derek Hill, and Reza Razavi. A fast-marching approach to cardiac electrophysiology simulation for XMR interventional imaging. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005*, pages 607–615, 2005.
- [23] Sophie Giffard-Roisin, Thomas Jackson, Lauren Fovargue, Jack Lee, Hervé Delingette, Reza Razavi, Nicholas Ayache, and Maxime Sermesant. Non-Invasive Personalisation of a Cardiac Electrophysiology Model from Body Surface Potential Mapping. *IEEE Transactions on Biomedical Engineering*, 64(9):2206 – 2218, September 2017.
- [24] Colleen Mitchell and David Schaeffer. A two-current model for the dynamics of cardiac membrane. *Bulletin of mathematical biology*, 65:767–93, 10 2003.
- [25] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- [26] M. Wallman, N. P. Smith, and B. Rodriguez. A comparative study of graph-based, eikonal, and monodomain simulations for the estimation of cardiac activation times. *IEEE Transactions on Biomedical Engineering*, 59(6):1739–1748, 2012.
- [27] Sophie Giffard-Roisin, Hervé Delingette, Thomas Jackson, Jessica Webb, Lauren Fovargue, Jack Lee, Christopher A Rinaldi, Reza Razavi, Nicholas Ayache, and Maxime

Sermesant. Transfer Learning from Simulations on a Reference Anatomy for ECGI in Personalised Cardiac Resynchronization Therapy. *IEEE Transactions on Biomedical Engineering*, 20, 2018.