



Generative Modeling and Inverse Imaging of Cardiac Transmembrane Potential

Sandesh Ghimire¹(✉), Jwala Dhamala¹, Prashnna Kumar Gyawali¹,
John L. Sapp², Milan Horacek², and Linwei Wang¹

¹ Rochester Institute of Technology, Rochester, NY 14623, USA
sg9872@rit.edu

² Dalhousie University, Halifax, NS, Canada
<http://www.sandeshgh.com>

Abstract. Noninvasive reconstruction of cardiac transmembrane potential (TMP) from surface electrocardiograms (ECG) involves an ill-posed inverse problem. Model-constrained regularization is powerful for incorporating rich physiological knowledge about spatiotemporal TMP dynamics. These models are controlled by high-dimensional physical parameters which, if fixed, can introduce model errors and reduce the accuracy of TMP reconstruction. Simultaneous adaptation of these parameters during TMP reconstruction, however, is difficult due to their high dimensionality. We introduce a novel model-constrained inference framework that replaces conventional physiological models with a deep generative model trained to generate TMP sequences from low-dimensional generative factors. Using a variational auto-encoder (VAE) with long short-term memory (LSTM) networks, we train the VAE decoder to learn the conditional likelihood of TMP, while the encoder learns the prior distribution of generative factors. These two components allow us to develop an efficient algorithm to simultaneously infer the generative factors and TMP signals from ECG data. Synthetic and real-data experiments demonstrate that the presented method significantly improve the accuracy of TMP reconstruction compared with methods constrained by conventional physiological models or without physiological constraints.

Keywords: Inverse problem · ECG imaging
Sequential variational auto-encoder · Bayesian inference

1 Introduction

Noninvasive electrophysiological (EP) imaging involves the reconstruction of cardiac electrical activity from high-density body-surface electrocardiograms (ECGs) [6]. It solves an ill-posed inverse problem that deteriorates as the imaging depth increases from the epicardium to the endocardium [9]. One type of increasingly utilized regularization considers knowledge about the well-defined physiological process of cardiac electrical propagation. This is often realized in

a model-constrained approach, where the optimization or statistical inference of cardiac electrical activity is constrained by a pre-defined model describing local activation/repolarization and its spatial propagation [4, 11, 12]. Earlier models include step jump functions [10], logistic functions [11], and 3D curve models [4] empirically parameterized to mimic the physiological process. Recently, more expressive cardiac EP simulation models have also been used [7, 12].

These model-constrained approaches are afflicted with a common challenge: they are controlled by high-dimensional parameters often associated with local tissue properties and the origin of electrical activation that are unknown *a priori*. The more expressive the model is, the more parameters it has. To fix these model parameters in optimization/inference, as is common in existing approaches [12], model errors may be introduced decreasing the accuracy of the estimated electrical activity [12]. To adapt these model parameters to the observed data, as is desired for accurate inference, is however difficult due to their high-dimensionality and nonlinear relationship with the observed ECG data [3].

In this paper, we introduce a novel model-constrained inference framework that replaces the conventional physiological models with a deep generative model that is trained to generate the spatiotemporal dynamics of transmembrane potential (TMP) from a low-dimensional set of *generative factors*. These generative factors can be viewed as a low-dimensional abstraction of the high-dimensional physical parameters, which allows us to efficiently adapt the prior physiological knowledge to the observed ECG data (through inference of the generative factors) for an improved reconstruction of TMP dynamics.

In specific, the presented method consists of two novel contributions. First, to obtain a generative model that is sufficiently expressive to reproduce the temporal sequence of 3D spatial TMP distributions, we adopt a novel sequence-to-sequence variational auto-encoder (VAE) [2] with cascaded long short-term memory (LSTM) networks. This VAE is trained on a large database of simulated TMP dynamics originating from various myocardial locations and with a wide range of local tissue properties. Second, once trained, the VAE decoder describes the likelihood of the TMP conditioned on a low-dimensional set of generative factors, while the encoder learns the posterior distributions of the generative factors conditioned on the training data. We utilize these two components within the Bayesian inference, and present a variation of the expectation-maximization (EM) algorithm to jointly estimate the generative factors and transmural TMP signals from observed ECG data. In a set of synthetic and real-data experiments, we demonstrate that the presented method is able to improve the accuracy of transmural EP imaging in comparison to statistical inference either constrained by a conventional physiological model [12] or without physiological constraints.

2 Generative Modeling of TMP via Sequential VAE

To learn to generate the spatiotemporal TMP sequences, we use a sequential variation of VAE [8] based on the use of LSTM networks [2].

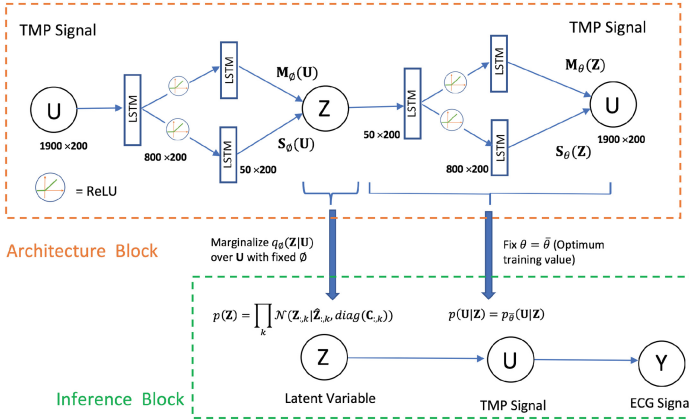


Fig. 1. Red block: VAE architecture. Green block: graphical model in inference.

VAE Architecture: The architecture of the sequential VAE is summarized in the red block in Fig. 1. Both the encoder and the decoder consists of two layers of LSTM, where the second layer includes separate mean and variance networks. The spatial dimension decreases from the original TMP signal \mathbf{U} to the latent representation \mathbf{Z} , while the temporal relationship is modeled by the LSTMs. Note that while the random variables in a standard VAE are vectors, a sequential VAE deals with matrices. By defining the conditional distribution of a matrix as the product of distributions over its columns, we obtained the likelihood distribution $p_{\theta}(\mathbf{U}|\mathbf{Z})$ and the variational posterior distribution $q_{\phi}(\mathbf{Z}|\mathbf{U})$ as:

$$p_{\theta}(\mathbf{U}|\mathbf{Z}) = \prod_k \mathcal{N}(\mathbf{U}_{:,k} | \mathbf{M}_{\theta}(\mathbf{Z})_{:,k}, \text{diag}(\mathbf{S}_{\theta}(\mathbf{Z})_{:,k})) \quad (1)$$

$$q_{\phi}(\mathbf{Z}|\mathbf{U}) = \prod_k \mathcal{N}(\mathbf{Z}_{:,k} | \mathbf{M}_{\phi}(\mathbf{U})_{:,k}, \text{diag}(\mathbf{S}_{\phi}(\mathbf{U})_{:,k})) \quad (2)$$

where $\mathbf{M}_{\phi}(\mathbf{U})$ and $\mathbf{S}_{\phi}(\mathbf{U})$ are output from the mean and variance networks of the encoder parameterized by ϕ , and $\mathbf{M}_{\theta}(\mathbf{Z})$ and $\mathbf{S}_{\theta}(\mathbf{Z})$ are output from the mean and variance networks of the decoder parameterized by θ .

VAE Training: Training of the VAE is performed by maximizing the variational lower bound on the likelihood of the training data given as:

$$\mathcal{L}_{ELB}(\theta, \phi; \mathbf{U}^{(i)}) = -KL(q_{\phi}(\mathbf{Z}|\mathbf{U}^{(i)}) || p_{\theta}(\mathbf{Z})) + E_{q_{\phi}(\mathbf{Z}|\mathbf{U}^{(i)})}(\log p_{\theta}(\mathbf{U}^{(i)}|\mathbf{Z})) \quad (3)$$

where $p_{\theta}(\mathbf{Z})$ is an isotropic Gaussian prior. The calculation of the KL divergence and cross entropy loss for the presented sequential architecture is carried out in a manner similar to that described in [8]. The training data is generated by the Aliev-Panfilov (AP) model [1], simulating spatiotemporal TMP sequences originated from different ventricular locations with different tissue properties.

3 Transmural EP Imaging

The biophysical relationship between cardiac TMP, \mathbf{U} and body-surface ECG, \mathbf{Y} can be described by a linear measurement model: $\mathbf{Y} = \mathbf{H}\mathbf{U}$, where \mathbf{H} is specific to the heart-torso model of an individual. To estimate \mathbf{U} from \mathbf{Y} is severely ill-posed and requires the regularization from additional knowledge about \mathbf{U} .

Probabilistic Modeling of the Inverse Problem: We formulate the inverse problem in the form of statistical inference. We define the likelihood distribution of \mathbf{Y} given \mathbf{U} by assuming zero-mean measurement errors with variance β^{-1} :

$$p(\mathbf{Y}|\mathbf{U}, \beta) = \prod_k \mathcal{N}(\mathbf{Y}_{:,k} | \mathbf{H}\mathbf{U}_{:,k}, \beta^{-1}\mathbf{I}) \quad (4)$$

To incorporate physiological knowledge about \mathbf{U} , we model its prior distribution conditioned on \mathbf{Z} using the VAE decoder with trained parameter $\bar{\theta}$:

$$p_{\bar{\theta}}(\mathbf{U}|\mathbf{Z}) = \prod_k \mathcal{N}(\mathbf{U}_{:,k} | \mathbf{M}_{\bar{\theta}}(\mathbf{Z})_{:,k}, \text{diag}(\mathbf{S}_{\bar{\theta}}(\mathbf{Z})_{:,k})) \quad (5)$$

To further utilize the knowledge about the generative factor \mathbf{Z} learned by the VAE from a large training dataset, we also utilize the VAE-encoded marginal posterior distribution of \mathbf{Z} as its prior distribution in Bayesian inference. In specific, we approximate samples from this marginalized distribution to be Gaussian:

$$p(\mathbf{Z}) = \prod_k \mathcal{N}(\mathbf{Z}_{:,k} | \bar{\mathbf{Z}}_{:,k}, \text{diag}(\mathbf{C}_{:,k})) \quad (6)$$

With this, we complete the statistical formulation of our problem. Our goal is to estimate the joint posterior distributions $p(\mathbf{U}, \mathbf{Z}|\mathbf{Y}) \propto p(\mathbf{Y}|\mathbf{U})p(\mathbf{U}|\mathbf{Z})p(\mathbf{Z})$.

Inference: Due to the presence of a deep neural network, the posterior $p(\mathbf{U}, \mathbf{Z}|\mathbf{Y})$ is analytically intractable. To address this issue, we note that conditioned on \mathbf{Z} , the distribution of \mathbf{U} is Gaussian in each column; thus, $p(\mathbf{U}|\mathbf{Y}, \mathbf{Z})$ is analytically available. We leverage this fact and employ a variant of the expectation maximization (EM) algorithm to obtain the *maximum a posteriori* (MAP) estimate of \mathbf{Z} along with the posterior distribution of \mathbf{U} given the MAP estimate of \mathbf{Z} .

E-step: Conditioned on an estimated value of \mathbf{Z} (say $\hat{\mathbf{Z}}$), we calculate $\hat{p}(\mathbf{U}|\mathbf{Y}, \hat{\mathbf{Z}}) = \prod_k \mathcal{N}(\mathbf{U}_{:,k} | \hat{\mathbf{U}}_{:,k}, \hat{\Sigma}_{:,:,k})$, with the covariance and mean of the k^{th} column of \mathbf{U} as:

$$\hat{\Sigma}_{:,:,k} = (\beta\mathbf{H}^T\mathbf{H} + \mathbf{D}_k^{-1})^{-1}, \quad \hat{\mathbf{U}}_{:,k} = \hat{\Sigma}_{:,:,k}(\beta\mathbf{H}^T\mathbf{Y}_{:,k} + \mathbf{D}_k^{-1}\mathbf{m}_k) \quad (7)$$

where $\mathbf{D}_k = \text{diag}(\mathbf{S}_{\bar{\theta}}(\hat{\mathbf{Z}})_{:,k})$, and $\mathbf{m}_k = \mathbf{M}_{\bar{\theta}}(\hat{\mathbf{Z}})_{:,k}$ and $\mathbf{S}_{\bar{\theta}}(\hat{\mathbf{Z}})_{:,k}$ are the k^{th} column output of the VAE decoder network when $\hat{\mathbf{Z}}$ is input to it.

M-step: Given $\hat{p}(\mathbf{U}|\mathbf{Y}, \hat{\mathbf{Z}})$, we update \mathbf{Z} by maximizing $E_{\hat{p}(\mathbf{U}|\mathbf{Y}, \hat{\mathbf{Z}})} \log(p(\mathbf{Y}, \mathbf{U}, \mathbf{Z}))$

$$\mathcal{L} = E_{\prod_k \mathcal{N}(\mathbf{U}_{:,k}|\hat{\mathbf{U}}_{:,k}, \hat{\Sigma}_{:,k})} [\log(p_{\hat{\theta}}(\mathbf{U}|\mathbf{Z}))] + \log(p(\mathbf{Z})) + \text{constant} \quad (8)$$

Realizing that a complete optimization of \mathcal{L} with respect to \mathbf{Z} would be expensive, we instead take a few gradient descent steps towards the optimum. The gradient of the second term is analytically available. The gradient of the first term is calculated by backpropagation through the decoder network.

The EM steps iterate until convergence, at which we obtain both the MAP value of \mathbf{Z} and the posterior distribution of \mathbf{U} conditioned on \mathbf{Z} and \mathbf{Y} .




Fig. 2. Examples of TMP signals generated by samples from two different distributions: Left- marginalized posterior density encoded by the VAE; Right- isotropic Gaussian.

4 Results

Synthetic Experiments: Synthetic experiments are carried out on two image-derived human heart-torso models. On each heart, the VAE is trained using around 850 simulated TMP signals considering approximately 50 different origins of ventricular activation in combination with 17 different tissue property configurations. As an initial study, here we focus on tissue properties representing local regions of myocardial scars with varying sizes and locations.

The presented method incorporating the trained VAE model is then tested on simulated 120-lead ECG data from three different settings, each with 20 experiments. The three settings include (1) presence of myocardial scar not included in training data, (2) origin of ventricular activation different from those used in training, and (3) both myocardial scar and activation origin not seen in training. In all experiments, the performance of the presented method is compared to 0-order Tikhonov regularization with temporal constraint (Greensite method) [5] and conventional EP model constrained inference with fixed parameters [12].

The reconstruction accuracy is measured with three metrics: (1) normalized RMSE given by the ratio of Frobenius norm of the error matrix to that of the truth TMP matrix, (2) Euclidean distance between the reconstructed and true origins of ventricular activation, and (3) Dice coefficient of the reconstructed S_1 and true regions of scar S_2 as $= 2|S_1 \cap S_2| / (|S_1| + |S_2|)$. In the two physiologically constrained methods, region of scar is defined based on absence or delay of

activation and shortening of action potential duration; in Greensite method, since the reconstructed signal no longer preserves the temporal shape of TMP, the region of scar is defined based on the peak amplitude of the signal.

Computational Cost: Training of the VAE takes approximately 40 h on a 4 GB Nvidia Quadro P1000 GPU. Generation of training data for each heart takes about 7 h and inference around 30 min on Quadcore CPU.

TMP Generation: Fig. 2 shows examples of local TMP signals generated by the trained VAE decoder against TMP signals simulated by the AP model [1]. Note that, when generating from a isotropic Gaussian (Fig. 2 right), noisy rather than meaningful TMP signals may also be generated. In comparison, when sampling from the approximated posterior distribution of \mathbf{Z} as described in Eq. (6), the generated signals closely resemble the simulated TMP signals.

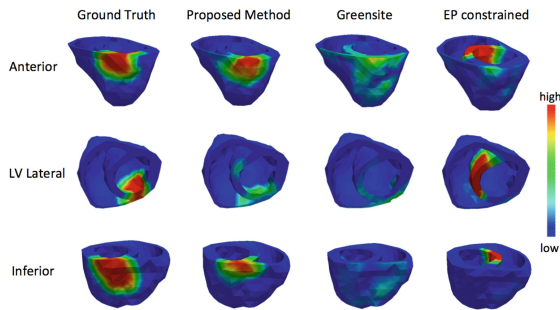


Fig. 3. Snapshots of early TMP pattern reconstructed by the three methods in comparison to the ground truth. The origin of activation is noted on the left in each row.

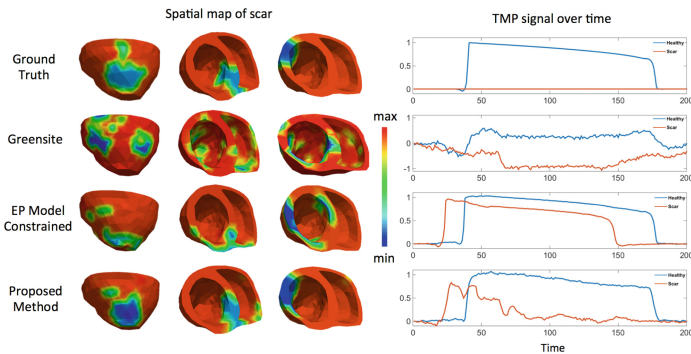


Fig. 4. Spatial distributions of scar tissues and temporal TMP signals obtained by the three methods in comparison to the ground truth.

Imaging TMP from Various Origins: Fig. 3 shows a snapshot from the early stage of ventricular activation reconstructed by the three methods in comparison to the ground truth. Since the EP model constrained approach assumes general sinus-rhythm activation, it introduces model error that incorrectly dominates the results. The simple Greensite method, free from erroneous model assumption, actually does a better job in comparison. By adapting model generative factors to the data, the presented method demonstrates a significantly improved ability to reconstruct TMP sequence resulting from unknown origins.

Imaging TMP at the Presence of Myocardial Scar: Fig. 4 shows the spatial distribution of scar tissue obtained by the three different methods, along with temporal TMP signals reconstructed in healthy and scar regions, in comparison to the ground truth. Without prior physiological knowledge, the Greensite method is not able to preserve the temporal TMP shape, resulting in high RMSE error as shown in Table 1. By thresholding the maximum amplitude of the reconstructed signals, the identified region of scar has high false positives and resembles poorly with the ground truth. The EP model constrained approach does a better job in retaining the temporal TMP shape. However, without prior knowledge about the scar, the model error again affects the accuracy of TMP reconstruction, especially in the early stage of activation when a smaller amount of ECG data is available for correcting the model error. The presented method, in comparison, is able to recognize the presence of scar tissue, adapting the physiological constraint for improved TMP reconstructions and scar identifications.

Summary: Table 1 summarizes the quantitative comparison of the three methods tested in the three settings as described earlier. Although the test cases were not seen by the VAE during training, the proposed method shows a significant improvement in inverse reconstruction (paired t-test, $p < 0.001$) when compared

Table 1. Quantitative accuracy of the three methods in three settings. Test data is simulated with **(1) Top:** scar not in VAE training, **(2) Middle:** activation origin not in training, **(3) Bottom:** both myocardial scar and activation origin not in training.

	Greensite	EP constrained	Proposed Method
Normalized RMSE	1.005 ± 0.006	0.3 ± 0.04	0.23 ± 0.05
Dice coefficient	0.19 ± 0.04	0.25 ± 0.09	0.52 ± 0.2

	Greensite	EP constrained	Proposed Method
Normalized RMSE	1.001 ± 0.003	0.28 ± 0.05	0.11 ± 0.08
Euclidean Distance	18.5 ± 10.96	39.47 ± 6.3	14.37 ± 14.0

	Greensite	EP constrained	Proposed Method
Normalized RMSE	1.005 ± 0.003	0.39 ± 0.03	0.29 ± 0.09
Dice coefficient	0.20 ± 0.07	0.21 ± 0.05	0.48 ± 0.24
Euclidean Distance	18.7 ± 9.3	65.5 ± 11.02	17.89 ± 10.6

with the other two methods in all settings and metrics except with Euclidean distance using Greensite method, where improvement is only marginal. It shows the importance of physiological knowledge and its adaptation to observed data during model-constrained inference.




Fig. 5. Real-data experiments: regions of scar tissues identified by the presented method and conventional EP model constrained method, in comparison to bipolar voltage data (red: scar core; green: scar border; purple: healthy tissue).

Real Data Experiments: Two case studies are performed on real-data from patients who underwent catheter ablation due to scar-related ventricular arrhythmia. Spatiotemporal TMP is reconstructed from 120-lead ECG data using the presented method and the EP model constrained method. In Fig. 5, scar regions (red regions with low voltage) identified from the reconstructed TMP are compared with scar regions (red regions) in the *in-vivo* bipolar voltage data. In both cases, while the scar tissue identified by two methods are generally in similar locations, the presented method shows less false positives and higher qualitative consistency with bipolar voltage maps.

5 Discussion and Conclusions:

To our knowledge, this is the first work that integrates a generative network learned from numerous examples into a statistical inference framework to allow the adaptation of prior physiological knowledge via a small number of generative factors. The results show the ability of this concept to improve model-constrained inference. Since the present formulation is in a personalized setting, we intend to extend this architecture to learn a geometry-invariant generative model that can be trained on multiple heart models and applied on a new subject.

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

References

1. Aliev, R.R., Panfilov, A.V.: A simple two-variable model of cardiac excitation. *Chaos Solitons Fractals* **7**(3), 293–301 (1996)
2. Bowman, S.R., Vilnis, L., Vinyals, O., Dai, A.M., Jozefowicz, R., Bengio, S.: Generating sentences from a continuous space. arXiv preprint [arXiv:1511.06349](https://arxiv.org/abs/1511.06349) (2015)
3. Ghimire, Sandesh, Sapp, John L., Horacek, Milan, Wang, Linwei: A variational approach to sparse model error estimation in cardiac electrophysiological imaging. In: Descoteaux, Maxime, Maier-Hein, Lena, Franz, Alfred, Jannin, Pierre, Collins, D.Louis, Duchesne, Simon (eds.) MICCAI 2017. LNCS, vol. 10434, pp. 745–753. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_84
4. Ghodrati, A., Brooks, D.H., Tadmor, G., MacLeod, R.S.: Wavefront-based models for inverse electrocardiography. *IEEE TBME* **53**(9), 1821–1831 (2006)
5. Greensite, F., Huiskamp, G.: An improved method for estimating epicardial potentials from the body surface. *IEEE TBME* **45**(1), 98–104 (1998)
6. Gulrajani, R.M.: The forward and inverse problems of electrocardiography. *IEEE Eng. Med. Biol. Mag.* **17**(5), 84–101 (1998)
7. He, B., Li, G., Zhang, X.: Noninvasive imaging of cardiac transmembrane potentials within three-dimensional myocardium by means of a realistic geometry anisotropic heart model. *IEEE TBME* **50**(10), 1190–1202 (2003)
8. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013)
9. Plonsey, R., Barr, R.C.: *Bioelectricity: A Quantitative Approach*. Springer, New York (2007). <https://doi.org/10.1007/978-0-387-48865-3>
10. Pullan, A., Cheng, L., Nash, M., Bradley, C., Paterson, D.: Noninvasive electrical imaging of the heart: theory and model development. *Ann. Biomed. Eng.* **29**(10), 817–836 (2001)
11. Van Dam, P.M., Oostendorp, T.F., Linnenbank, A.C., Van Oosterom, A.: Non-invasive imaging of cardiac activation and recovery. *Ann. Biomed. Eng.* **37**(9), 1739–1756 (2009)
12. Wang, L., Zhang, H., Wong, K.C., Liu, H., Shi, P.: Physiological-model-constrained noninvasive reconstruction of volumetric myocardial transmembrane potentials. *IEEE Trans. Biomed. Eng.* **57**(2), 296–315 (2010)