Assessing Ionic Current Blockades and Electromechanical Biomarkers' Interrelations Through a Novel Multi-Channel Causal Variational Autoencoder

Safaa Al-Ali[1], Maria Teresa Mora[2], Maxime Sermesant[1], Beatriz Trenor[2], Irene Balelli[1]

[1] Centre INRIA d'Université Côte d'Azur, EPIONE team, Sophia Antipolis, France
[2] Centro de Investigación e Innovación en Bioingeniería (Ci2B), Universitat Politècnica de València, Valencia, Spain

Abstract

Drug-induced Torsade de pointes (TdP) is a critical arrhythmia that can lead to sudden cardiac death. Besides ionic current blockades, in-silico electrophysiological and mechanical biomarkers can provide mechanistic proarrhythmic information for TdP-risk assessment, and specific torsadogenic indices have been developed for that purpose, yet, determining the causal relationships between these variables is challenging. This study aims to employ a novel Multi-Channel Causal Variational Autoencoder (MC^{2}VAE), to identify causal relationships between ionic current blockades, electrophysiological biomarkers, and torsadogenic indices, considered as three distinct sources of information (channels) for drug-induced TdP risk. Our approach interestingly suggests the existence of latent causal relationships between the considered channels, allowing for a better reconstruction of all observed features. Further, MC^{2}VAE can quantify the strengths of the identified causal relationships, opening up a viable avenue for actionable interventions on the established causal graph. Finally, we consider the downstream task of drug classification for TdP risk on the latent channels and show a clear improvement in classification performances when combining the three considered channels. Overall, our results provide a strong rationale for causally combining multi-channels biomarkers in TdP-risk characterization.

1. Introduction

Some molecules can interfere with cardiac electrophysiology and induce the life-threatening arrhythmia known as Torsade de Pointes (TdP): therefore, pro-arrhythmia assessment is needed in the preclinical stages of drug development. The limited accuracy of strategies focusing on hERG block and QT prolongation led to the Comprehensive In Vitro Proarrhythmia Assay (CiPA) initiative, which promotes in silico simulations to improve predictions by providing a mechanistic classification [1]. Electrophysiological models considering the blockage of seven ionic currents, instead of only the main repolarizing current I_{Kr}, are used to compute biomarkers able to better identify torsadogenic drugs. Moreover, detailed and complex biophysical models have been developed to better represent myocyte activity, including excitation-contraction coupling, the process by which changes in membrane potential trigger calcium release to activate myofilaments. Therefore, drug effects on ionic currents can be translated from excitation to contraction. So far, multiple indices have been proposed for TdP-risk assessment [2], most of them based on drug properties and action potential, and the latest increasing tendency is to combine several electrophysiological properties, including Ca-derived features [3, 4]. Besides, instead of striving to find a single optimal predictor, the strategy of combining multiple features and analyzing them with machine learning tools seems promising [5]. In that sense, the main dilemma is whether or not we are over-exploiting the model by using as input many interrelated variables.

Therefore, the challenge lies in effectively combining the different sources of information (*channels*) to achieve a more informative TdP-risk assessment. We hypothesize that discovering causal relationships between the different channels could help our understanding of the underlying mechanisms leading to TdP, and improve its characterization. However, the process of analyzing data from different channels is challenging, due to their heterogeneity and the potential presence of redundant shared information. We consider here Variational Autoencoders (VAEs), Bayesian generative models that encode data into a meaningful lower-dimensional latent space, from which the original data can be effectively generated [6]. VAEs are flexible enough to adapt to a variety of data types and structures and have already shown promising in the context of multi-channel data analysis [7].

In this work, we propose to use a novel Multi-Channel Causal Variational Autoencoder [8] to identify latent causal relationships between ionic current blockades, elec-

 ISSN: 2325-887X DOI: 10.22489/CinC.2024.408

trophysiological biomarkers and torsadogenic indices considered as three distinct channels. In Sec. 2 we describe the data and give details about the mathematical formulation of MC^{2}VAE. Sec.3 presents our results, while Sec.4 provides some conclusions and perspectives.

2. Materials and method

2.1. Dataset collection and preprocessing

Electrophysiological cellular simulations with drugs were performed using a modified version of the human endocardial ventricular O'Hara et al. model [9, 10]. A total of 109 drugs from CredibleMeds (accessed March 9, 2020) with known torsadogenic risk (37 with known, 14 with possible, 13 with conditional, and 45 with no proven risk) were used in this study. Each drug consists of data coming from 3 channels, I, E, and T, described below.

I: The drug-induced blockade of the seven most important ionic currents according to the CiPA initiative, computed using the simple pore block model, at the effective free therapeutic plasma concentration (EFTPC) [11].

E: Ten in-silico electrophysiological biomarkers that consist of direct features from the action potential (APD90, APD50, Tri9050, Tri9030, qNet) and from the calcium transient signals (Casyst, Cadiast, CaTD90, CaTD50), including a surrogate of the electromechanical window (Emw = CaTD90-APD90) [5].

T: Three in-silico derived features that have been proposed as torsadogenic indices: T_x, the ratio between the concentration of a drug that provokes a 10% prolongation of the APD90 in control conditions and the EFTPC; T_{qNet}, the ratio between the net charge carried by the net current when exposed to 10 times the EFTPC with respect to the net charge in control conditions; and T_{triang}, the ratio between Tri9030 for a drug concentration of 10 times EFTPC and triangulation in control conditions [9].

A total of 20 features are observed per drug. To ensure their compatibility with the model's assumptions (Sec. 2.2), we applied a power transformation to I and a logarithmic transformation to T and E, followed by standardization.

2.2. MC^{2}VAE

The structure of MC^{2}VAE is composed of three main components: 1) encoding, 2) causal layer, and 3) decoding (Fig. 1). Each channel is projected into a one-dimensional latent space through its channel-specific encoder. The obtained latent variables are fed to the causal layer, where a causal graph is learned to discover latent causal relationships between the channels, hence transformed according to the learned causal graph. Finally, causal latent variables pass through their respective channel-specific decoder which reconstructs the input data [8]. In the following, we briefly outline the formal definition of MC^{2}VAE, and its optimization. The main code is made publicly available on GitLab.

2.2.1. Latent structural causal model

We denote by $\mathbf{X} = \{I, E, T\}$ the dataset consisting of $N = 109$ drugs, where I denotes the ionic current blockade fractions, E the electrophysiological biomarkers and T the torsadogenic indices. $\mathbf{X}_i := (I_i, E_i, T_i)$ represents the dataset of the i^{th} drug.

We assume independence between the drugs, and hypothesize linear causal relationships across the channel latent variables, $\mathbf{z}^c := (z^{I,c}, z^{E,c}, z^{T,c})$, and a normal prior for their respective noise terms, $\mathbf{z} := (z^I, z^E, z^T)$. These assumptions define a latent structural causal model (SCM):

$$\mathbf{z}^c = A_{\boldsymbol{\gamma}}^T \mathbf{z}^c + \mathbf{z} = (\mathbb{I} - A_{\boldsymbol{\gamma}}^{\mathrm{T}})^{-1}\mathbf{z}, \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbb{I}), \quad (1)$$

where $\mathbb{I}$ denotes the identity matrix, and $A_{\boldsymbol{\gamma}}$ is a weighted adjacency matrix, to be learned: $(A_{\boldsymbol{\gamma}})_{ij}$ provides the strength of the causal linear relationships of the parent variable $z^{i,c}$ to the children variable $z^{j,c}$, i, j in $\{I, E, T\}$.

2.2.2. MC^{2}VAE loss function

The encoder and decoder channel-specific operators are neural networks parametrized by $\boldsymbol{\Theta} := (\boldsymbol{\theta}_m)_{m=I,E,T}$ and $\boldsymbol{\Phi} := (\boldsymbol{\phi}_m)_{m=I,E,T}$ respectively. To optimize the parameters of MC^{2}VAE, we should maximize the marginal log-likelihood of $\mathbf{X}$, $\mathcal{L}(\mathbf{X}; \boldsymbol{\Phi}, \boldsymbol{\Theta}, A_{\boldsymbol{\gamma}})$ and derive the true posterior $p(\mathbf{z}|\mathbf{X})$ over the latent space. Due to analytical intractability, we apply variational Bayes and introduce a tractable posterior $q(\mathbf{z}|\mathbf{X}; \boldsymbol{\Theta})$ which approximates $p(\mathbf{z}|\mathbf{X})$ [12]. We get:

$$\begin{aligned}\mathcal{L}(\mathbf{X}; \boldsymbol{\Phi}, \boldsymbol{\Theta}, A_{\boldsymbol{\gamma}}) &= \sum_{i=1}^{N} \log\left[p(\mathbf{X}_i; \Phi, A_{\boldsymbol{\gamma}})\right] \\ &\geq \mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{X};\boldsymbol{\Theta})}\left\{\log\left[p(\mathbf{X}|\mathbf{z}^{\mathbf{c}}; \boldsymbol{\Phi})\right]\right\} \\ &+\mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{X};\Theta)}\left\{\log\left[p(\mathbf{z}_{\mathbf{i}}^{\mathbf{c}}|\mathbf{z}_{\mathbf{i}}; A_{\boldsymbol{\gamma}})\right]\right\} \\ &-\mathcal{D}_{\mathrm{KL}}\left(q(\mathbf{z}|\mathbf{X}; \boldsymbol{\Theta})||p(\mathbf{z})\right) := \mathcal{E},\end{aligned}$$

where $p(\mathbf{X}|\mathbf{z}^c; \boldsymbol{\Phi})$ is the probability distribution of the decoder, $p(\mathbf{z}_i^c|\mathbf{z}_i; A_{\boldsymbol{\gamma}})$ is the Markov factorization of the joint distribution of the latent variables according to the assumed SCM, $p(\mathbf{z})$ is the normal prior of the latent causal variables (*c.f.* Eq. (1)), and $\mathcal{D}_{\mathrm{KL}}$ denotes the Kullback-Leibler divergence. Since we seek for directional cause-and-effect relationships among causal latent variables, we enforce the acyclicity of the causal graph, represented by $A_{\boldsymbol{\gamma}}$, through a penalization term inspired by [13].

3. Results

For the experiments presented here, all encoders and decoders consist of a single linear layer. MC^{2}VAE learning is

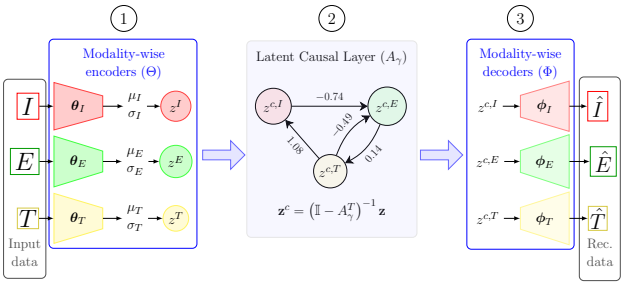

Figure 1: MC^{2}VAE pipeline for $\mathbf{X} = \{I, E, T\}$. In step 2) we represent the causal graph learned by our model.

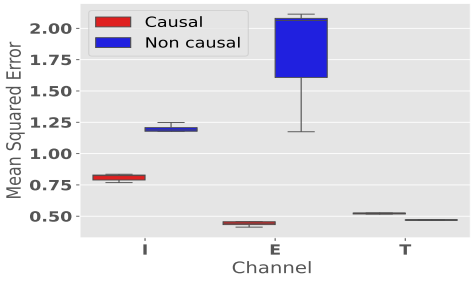

Figure 2: Mean squared errors per channel.

efficiently carried out through stochastic gradient descent using the Adam optimizer [14]. We consider 1000 epochs with an initial learning rate of $1e{-}2$ which allows us to reach convergence. All results are obtained by performing 5-fold cross-validation. We apply a prior constraint to our graph, imposing no causal relationship from mechanical biomarkers towards ionic current blockades.

The resulting causal graph reveals latent causal relationships between the three channels: in Fig. 1 (2) we report the average values obtained for the A_γ weights, which quantify the strengths of each causal link. MC^{2}VAE correctly identifies a causal link between drug-induced blockades (I) and electrophysiological biomarkers (E), which was expected. The causal relationship between T, the torsadogenic indices, to both I and E can be justified in the light of the definition of T's parameters themselves, which explicitly take into account drug concentration, affecting ionic current blockades and action potential-related features. Finally, it should be noted that our method exhibits a bidirectional causal relationship between E and T, meaning that it can not totally exclude the existence of a causal link from E to T. However, the absolute weight of the latter causal arrow is very low, hence can be considered as less relevant.

To quantify the impact of the discovered causal graph in defining a meaningful latent representation of the channels, we evaluate MC^{2}VAE's ability to reconstruct the input data, measured through mean squared error (MSE): the causal learning step enables efficient reconstruction of the input data, and a net improvement with respect to the case in which we perform separate training of 3 independent VAEs, one per channel (Fig. 2), blue boxplots).

In Fig. 3 (left) we show the three-dimensional latent representation of each drug obtained by sampling 50 times from their learned distribution, each axis being the latent coordinate for a specific channel. Despite MC^{2}VAE is a fully unsupervised method, we can clearly see the separation of the drugs in the latent 3D space with respect to their known TdP risk (considered here as binary, unsafe, in orange, for known or possible TdP risk drugs, and safe, in purple, for the remaining categories). For the sake of clarity, we further project these representations in each 2D plane (Fig. 3, right panels), and show two unsafe and two safe drugs: ibutilide (orange circle), disopyramide (orange square), loratadine (purple circle), and diltiazem (purple square), as an example.

Finally, we challenge our latent representation for the downstream binary classification task and study the impact of an increasing number of channels to characterize drug-induced TdP risk. Fig. 4 shows that the inclusion of each of the considered channels brings additional valuable information for TdP risk assessment, generating evidence and a strong rationale for including them in the analysis.

4. Conclusion

In this paper, we have used a novel Multi-Channel Causal Variational Autoencoder (MC^{2}VAE) to identify causal relationships between ionic current blockades, electrophysiological biomarkers, and torsadogenic indices. Our approach reveals and quantifies the hidden causal relationships among the different sets of biomarkers, and helps to justify the integration of these three channels, as their joint causal analysis shows that they produce a much more effective and actionable characterization of TdP risk than considering one channel at a time. This work opens up several directions: for instance, we plan to include additional information in the model, either as a new channel, or for a semi-supervised approach, or finally as a causal node to be related to the latent channels. We are also interested in studying more complex causal relationships while encoding-decoding structures can be diversified to better adapt to the intrinsic nature of each channel.

Acknowledgments

This work was supported by the SimCardioTest project which has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No. 101016496, and by the French government, through the 3IA Côte d'Azur Investments in the

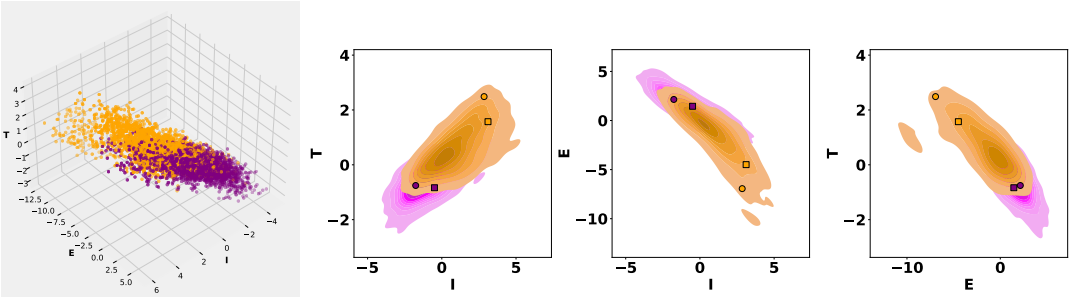

Figure 3: Left: scatter plot of all drugs in the 3D latent space. Right: Kernel density distribution of drugs for each pairwise combination of latent channels. Unsafe drugs are in orange; safe drugs are in purple.

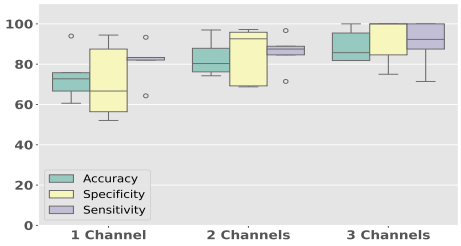

Figure 4: Drugs TdP-risk classification performance using a different number of modalities based on the voting classifier with Random-Forest, K-Nearest-Neighbors, Gaussian Naive Bayes and Linear Discriminant Analysis.

Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002.

References

[1] Colatsky T, Fermini B, Gintant G, Pierson JB, Sager P, Sekino Y, Strauss DG, Stockbridge N. The comprehensive in vitro proarrhythmia assay (CiPA) initiative—update on progress. Journal of Pharmacological and Toxicological Methods 2016;81:15–20.

[2] Parikh J, Di Achille P, Kozloski J, Gurev V. Global sensitivity analysis of ventricular myocyte model-derived metrics for proarrhythmic risk assessment. Frontiers in Pharmacology 2019;10.

[3] Lancaster MC, Sobie E. Improved prediction of drug-induced torsades de pointes through simulations of dynamics and machine learning algorithms. Clinical Pharmacology Therapeutics 2016;100(4):371–379.

[4] Passini E, Trovato C, Morissette P, Sannajust F, Bueno-Orovio A, Rodriguez B. Drug-induced shortening of the electromechanical window is an effective biomarker for in silico prediction of clinical risk of arrhythmias. British Journal of Pharmacology 2019;176(19):3819–3833.

[5] Llopis-Lorente J, Trenor B, Saiz J. Considering population variability of electrophysiological models improves the in silico assessment of drug-induced torsadogenic risk. Computer Methods and Programs in Biomedicine 2022; 221:106934.

[6] Kingma DP, Welling M. Stochastic gradient VB and the variational auto-encoder. In Proceedings of the 2nd International Conference on Learning Representations (ICLR). 2014; Volume 19.

[7] Abi Nader C, Ayache N, Frisoni GB, Robert P, Lorenzi M, Alzheimer's Disease Neuroimaging Initiative. Simulating the outcome of amyloid treatments in alzheimer's disease from imaging and clinical data. Brain Communications 2021;3(2).

[8] Al-Ali S, Balelli I. Multi-Channel Causal Variational Autoencoder, August 2024. URL https://hal.science/hal-04666466. Working paper or preprint.

[9] Llopis-Lorente J, Gomis-Tena J, Cano J, Romero L, Saiz J, Trenor B. In silico classifiers for the assessment of drug proarrhythmicity. Journal of Chemical Information and Modeling 2020;60(10):5172–5187.

[10] O'Hara T, Virág L, Varró A, Rudy Y. Simulation of the undiseased human cardiac ventricular action potential: model formulation and experimental validation. PLoS Computational Biology 2011;7(5):e1002061.

[11] Al-Ali S, Llopis-Lorente J, Mora MT, Sermesant M, Trenor B, Balelli I. A causal discovery approach to streamline ionic currents selection to improve drug-induced tdp risk assessment. In 2023 Computing in Cardiology (CinC), volume 50. 2023; 1–4.

[12] Blei DM, Kucukelbir A, McAuliffe JD. Variational inference: A review for statisticians. Journal of The American Statistical Association 2017;112(518):859–877.

[13] Zheng X, Aragam B, Ravikumar PK, Xing EP. Dags with no tears: Continuous optimization for structure learning. Advances in Neural Information Processing Systems 2018; 31.

[14] Kingma DP, Ba J. Adam: A method for stochastic optimization, 2017. URL https://arxiv.org/abs/1412.6980.

Address for correspondence:

Safaa Al-Ali
Centre INRIA d'Université Côte d'Azur, EPIONE team,
06902 Sophia Antipolis, France.
safaa.al-ali@inria.fr