



HAL
open science

Interpretable Prediction of Post-Infarct Ventricular Arrhythmia using Graph Convolutional Network

Buntheng Ly, Sonny Finsterbach, Marta Nuñez-Garcia, Pierre Jaïs, Damien Garreau, Hubert Cochet, Maxime Sermesant

► **To cite this version:**

Buntheng Ly, Sonny Finsterbach, Marta Nuñez-Garcia, Pierre Jaïs, Damien Garreau, et al.. Interpretable Prediction of Post-Infarct Ventricular Arrhythmia using Graph Convolutional Network. STACOM 2022 - 13th Workshop on Statistical Atlases and Computational Modelling of the Heart, Sep 2022, Singapore, Singapore. hal-03829609

HAL Id: hal-03829609

<https://inria.hal.science/hal-03829609>

Submitted on 25 Oct 2022

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.

Interpretable Prediction of Post-Infarct Ventricular Arrhythmia using Graph Convolutional Network

Buntheng Ly¹, Sonny Finsterbach², Marta Nuñez-Garcia³, Pierre Jais³,
Damien Garreau⁴, Hubert Cochet³, and Maxime Sermesant¹

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

² CHU Bordeaux, Université de Bordeaux, Bordeaux, France

³ IHU Liryc, Université de Bordeaux, Bordeaux, France

⁴ Université Côte d’Azur, Nice, France

Contact: `maxime.sermesant@inria.fr`

Abstract. Heterogeneity of left ventricular (LV) myocardium infarction scar plays an important role as anatomical substrate in ventricular arrhythmia (VA) mechanism. LV myocardium thinning, as observed on cardiac computed tomography (CT), has been shown to correlate with LV myocardial scar and with abnormal electrical activity. In this project, we propose an automatic pipeline for VA prediction, based on CT images, using a Graph Convolutional Network (GCN). The pipeline includes the segmentation of LV masks from the input CT image, the short-axis orientation reformatting, LV myocardium thickness computation and mid-wall surface mesh generation. An average LV mesh was computed and fitted to every patient in order to use the same number of vertices with point-to-point correspondence. The GCN model was trained using the thickness value as the node feature and the atlas edges as the adjacency matrix. This allows the model to process the data on the 3D patient anatomy and bypass the “grid” structure limitation of the traditional convolutional neural network. The model was trained and evaluated on a dataset of 600 patients (27% VA), using 451 (3/4) and 149 (1/4) patients as training and testing data, respectively. The evaluation results showed that the graph model (81% accuracy) outperformed the clinical baseline (67%), the left ventricular ejection fraction, and the scar size (73%). We further studied the interpretability of the trained model using LIME and integrated gradients and found promising results on the personalised discovering of the specific regions within the infarct area related to the arrhythmogenesis.

Keywords: Graph Neural Network, Ventricular Arrhythmia, Interpretable AI, Cardiac CT

1 Introduction

VA is the an abnormal heart rhythm most observed leading to the sudden cardiac death (SCD), ranking among the highest causes of mortality in the developed countries [6]. The current gold standard predictor of SCD is still the left

ventricular ejection fraction (LVEF), despite a majority of SCDs occurred in patients with preserved LVEF ($> 35\%$). The post-infarction scar on the LV myocardium (LVMYO) is recognised as the anatomical substrate of the VA mechanism. Nonetheless, identifying the precise arrhythmogenesis characteristics of the scar regions is still challenging, thus limiting its application in the clinical practice. Recent study has shown that the LVMYO scar can be located through the analysis of the LV wall thinning in CT imaging [5], as the alternative to the late gadolinium enhancement (LGE) MRI. CT imaging is generally more accessible in terms of machine availability and patient compatibility (i.e. patients with metallic device) compared to MRI. For image processing, CT imaging also have higher image resolution and quantitative image intensity consistency between manufacturers, thus reducing the potential error in the automatic processing pipeline.

In recent years, Deep Learning (DL), notably the convolution neural network (CNN), has made a remarkable impact in medical image processing, ranging from semantic segmentation to diagnosis and outcome prediction. Although, the grid-like property of the CNN can limit its efficacy on the organs with specific 3D anatomy, such as the LVMYO. Our previous work proposed transforming the CT input image into the 2D bullseye representation of the LV thickness to remove the blank voxels of the endocardium [4]. However, the flattening of the 3D led to the distortion of the 3D heterogeneity of the thickness map, impacting the subsequent fitting of the DL model. On the other hand, graph neural networks have gained more applications in medical image processing and bioinformatics [16], thanks to its adaptability to the specific data geometry.

In this study, we investigated the graph-level classification task using a model built with graph convolutional network (GCN) layers [3], and the LV thickness map as input. Moreover, we studied the interpretability of the graph model using LIME [8] and integrated gradients [11] to locate the specific regions contributed to the VA prediction and to gain better insight of the explicit arrhythmogenic regions from the model perspective.

2 Method

Our image processing pipeline formulates the 3D CT scan into a graph input. The processing steps were image segmentation, short-axis (SAX) reorientation, LVMYO thickness calculation, LV mid-wall meshing and average mesh fitting. From there, a graph classification network was built using the atlas mesh and the thickness value as input. The image processing pipeline and the model architecture are described in section 2.1 and section 2.2.

2.1 Image Processing

The 3D ventricular masks were segmented from the input CT image using a Dual-UNet segmentation model, composing of a “coarse” UNet [9] for ROI segmentation and a “refine” UNet for the cropped input segmentation. The required

masks included the epicardium and endocardium of the LV and the epicardium of the RV. Using the LV and RV mask as landmarks, automatic SAX view re-orientation was done using the method described in [7]. The SAX ventricular masks were then resampled isotropically at 0.5mm. The LV wall thickness and the mid-wall points were calculated using the Eulerian partial differential method as proposed by [15]. The thickness value was normalised by clipping at 10mm then divided by 10. Then, the surface mid-wall mesh was generated from the mid-wall points using marching cube algorithm and uniformly remeshed using approximated centroid voronoi diagrams (ACVD) method[13]. Finally, an average mid-wall mesh was generated and fitted to every patient based on the large deformation diffeomorphic metric mapping (LDDMM) framework. The atlas fitting was done using the deterministic atlas function provided by the Deformetrica software [1]. Rigid registration of the meshes was done prior to atlas fitting. As the mesh rotation and size were already registered in the SAX reorientation step, we translated the centre of mass of all the meshes toward a reference from a randomly selected mesh. At inference, the atlas mesh could be directly registered to the new mesh using affine registration. The image process steps are shown in Fig. 1.I.

2.2 Graph Convolutional Network Model

We built the model using the GCN layer [3], which used the following convolutional operation:

$$X' = \hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{\frac{1}{2}} XW + b, \quad (1)$$

where $X' \in \mathbb{R}^{N \times F}$ is the output of N nodes and F filters, $X \in \mathbb{R}^{N \times C}$ is the input of C features, $W \in \mathbb{R}^{C \times F}$ is the weights, b is the bias, $\hat{D} \in \mathbb{R}^{N \times N}$ the degree matrix, and $\hat{A} \in \mathbb{R}^{N \times N}$ is the self-loop adjacency matrix ($\hat{A} = I_N + A$).

We considered the LV mid-wall mesh as an indirect graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with N nodes $v_i \in \mathcal{V}$ and edges $(v_i, v_j) \in \mathcal{E}$. We used the thickness value of each node v_i as the input feature $X \in \mathbb{R}^{N \times 1}$. Since the input were uniformly remeshed, no edge feature was use. The adjacency matrix A was extracted from the atlas mesh edges and used for every inputs. The LV mesh could then be represented as a graph $\mathcal{G} = (X, A)$.

Index Pooling. We employed the pooling method as proposed by [10], where a coarse graph was generated to pre-define the corresponding pooling index and the new adjacency matrix during the model construction. In our case, the coarse graph was generated using the ACVD uniform meshing, with reduced number of points, and the corresponding pooling indexes were defined using k-dimensional tree nearest neighbour search (KD-tree NN). In this project, we built the index pooling layer using max pooling method. The graph coarsening and pooling indexes searching are computationally expensive, which is unfit for online processing. This pooling method is more suitable with a point-to-point correspondent graph input, as the pooling indexes of the atlas mesh can be applied to all the input graphs. The index pooling steps are shown in Fig. 1.III.

Model Architecture. The model architecture is shown in Fig. 1.II. Starting with an input graph of N nodes, the model was built using two consecutive GCN layers, followed by an index pooling layer, reducing the aggregated graph to N_{down} nodes. The pooling output was then passed through three more GCN layers, before feeding to the fully-connected network (FCN) classification block. We set the filter size to 64 for all the GCNs, except the last GCN, where the filter was set to 1. The classification block was built with 4 fully-connected layers with [128, 64, 32, 2] units, respectively. We applied linear rectifier activation the first 3 layers and the softmax activation for the output layer. The model was trained using the binary cross-entropy loss. The model was built using tools provided by Spektral ¹ (GCN layer) and Tensorflow ².

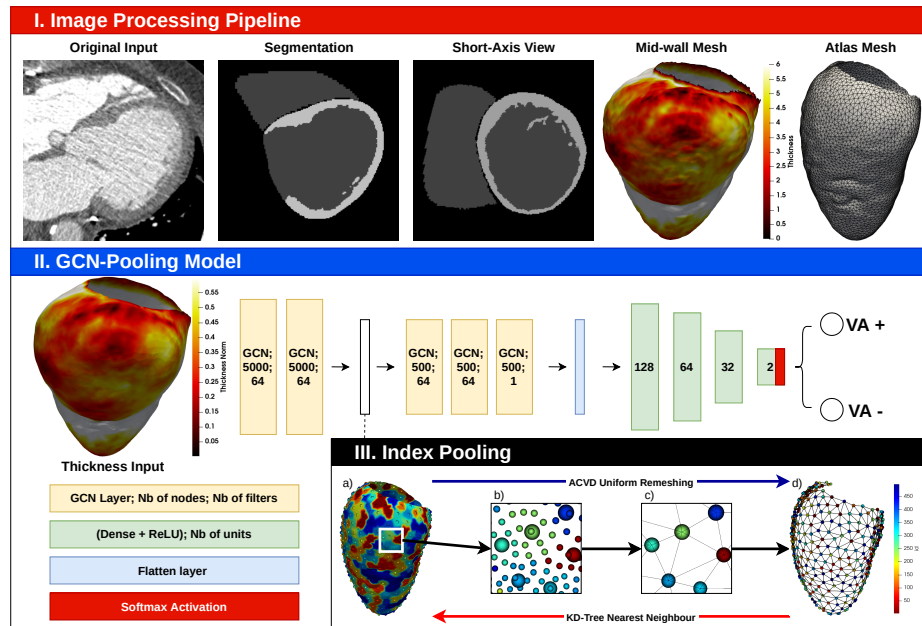


Fig. 1. Image processing pipeline, GCN-Pooling architecture and Index Pooling layer. The thickness input mesh was the result after fitting the average mesh to the mid-wall mesh, ensuring the point-to-point correspondence of every input. *Index Pooling*: a). pooling patches (from KD-tree NN); b). original (smaller spheres) and pooling nodes (bigger spheres); c & d). coarse graph.

¹ <https://github.com/danielegrattarola/spektral>

² <https://www.tensorflow.org/>

2.3 Interpretability Study

We studied the interpretability of the graph neural network by adapting to our context two existing methods: Local Interpretable Model-agnostic Explanation (LIME, [8]), and integrated gradients [11].

LIME. To summarize briefly, in the case of image data LIME computes a local surrogate linear model based on the absence or presence of superpixels. Our use-case is similar if we consider the input graph $\mathcal{G} = (X, A)$ as an image with pixel values defined as the thickness. Following [2], our method can be described as:

1. segment the mesh into d patches;
2. randomly *turn on/off* the patches to create new inputs $\mathcal{G}_1 = (X_1, A), \dots, \mathcal{G}_n = (X_n, A)$;
3. run the model prediction on all the new inputs $y_i = f(\mathcal{G}_i)$;
4. compute $\hat{\beta}_n$ by fitting the y_i s to a local weighted surrogate model.

We used two segmentation methods in step 1. The first is *quick shift* [14], as is default for the image version of LIME, which bases the segmentation on both the nodes 3D coordinates and the thickness value³. We scaled the magnitude of the node coordinates and thickness value using scikit-learn standard feature scaler⁴. The quickshift algorithm was then applied with the distance threshold of 0.2. The second is *grid segmentation* with KD-tree NN, which is based only on the nodes 3D coordinates. The mesh was first downsampled to 30 mm radius using voxel downsampling, before running the KD-Tree search, which generated a quasi uniform grid segmentation of the input mesh. Turning the patches *i on* meant replacing the original value of patch with the replacement value, and vice versa. To study the case of VA+ classification, the replacement value was set to 1 to represent the normalised healthy LVMYO, under the hypothesis that the completely healthy LV would be classified as VA-.

Integrated gradients. Integrated gradients is computed as the integral of back-propagated gradients of the straightline path from the baseline input x' to the original input x . To adapt the method to the graph network, the integrated gradients along the dimension i^{th} for a model $f(\mathcal{G})$ with the input graph $\mathcal{G} = (X, A)$ and baseline graph $\mathcal{G}' = (X', A)$ was calculated as:

$$\text{IG}_i(\mathcal{G} = (X, A)) := (X_i - X'_i) \times \int_{\alpha=0}^1 \frac{\partial f(X' + \alpha(X - X'), A)}{\partial X_i} d\alpha, \quad (2)$$

where $\alpha \in [0, 1]$ denotes the step coefficients between the baseline input ($\alpha = 0$) to the input ($\alpha = 1$). We set the number of step to 50. We set the baseline feature X' the same way as the replacement input in LIME by changing the thickness feature value to 1.

³ <https://github.com/Nick-01/MedoidShift-and-QuickShift>

⁴ <https://scikit-learn.org/>

3 Experimental Setup

We studied our method using a retrospective dataset of 600 patients, collected between 2010 and 2020. The dataset included the CT images captured at diastolic of patients with history myocardial infarction (MI) at least 1 month before the scan date and without any history of LV surgery. The dataset included 27.5% of patients labelled as VA+. The VA inclusion criteria were any episode of ventricular tachycardia (VT) or fibrillation (VF), and aborted cardiac arrest.

The total population was randomly divided at 1 : 4 ratio, resulting in 451 and 149 patients for training and testing, respectively. Two-sample t-test was used on all the available clinical characteristics to ensure that there was no significant difference between the training and testing population. The available characteristics in the dataset were age, gender, hypertension, diabetes, dyslipidaemia, smoking, LVEF, scar size and scar age. The scar size was calculated as the area on the LV surface with $< 5\text{cm}$ thickness, and the scar age was calculated as the delay between the MI and the scan date.

3.1 Baseline Models

DL Model. We studied the performance of the GCN-Pooling model against the baseline DL models using 10-fold cross-validation on the training dataset.

Using the 2D thickness bullseye prediction pipeline proposed by [4], we tested two models: the conditional variational autoencoder classification (CVAE-Class) model and the direct CNN classification model (using the same layers as the CVAE-Class without the decoder). To test the performance of the traditional CNNs, we also run the cross-validation test with the 3D LV wall mask as input, using the 3D variant of the VAE-Class (without the condition) and direct CNN model. We used the SAX oriented LV wall masks to ensured the rotation and the voxel spacing consistency between inputs.

For the graph input, we generated two sets of LV mid-wall meshes at 500 and 5000 nodes using the pipeline described in section 2.1. With the atlas-fitted inputs, we run cross-validation on the GCN-Pooling, GCN (without index pooling) and FCN model. The baseline GCN model was built with the same layers as the GCN-Pooling, and the FCN model was built the same layers as the classifier block (4-consecutive fully-connected layers). As ablation study, we also run cross-validation on the GCN and FCN model using the rigid registered remesh inputs without the atlas fitting.

Clinical Baseline. Finally, we compared the best DL model against the clinical characteristics. With training dataset, we trained the best DL model and run univariate analysis to find the significant variables associated with VA. Then, the optimal cut-off values were calculated with the receiver operator curve (ROC) analysis [12] using the following equation:

$$\text{Cutoff}_{\text{Optimal}} = \min(|\text{AUC} - \text{Spe}_c| + |\text{AUC} - \text{Sen}_c|). \quad (3)$$

We then compared the generalisability of the DL model and the cutoff values on the testing dataset.

4 Results

The Tab. 1 shows the cross-validation results of the DL models. The GCN-Pooling model outperformed the baselines models with the cross-validate accuracy of 0.818 ± 0.06 . It outperformed the 2D and 3D input models (both CVAE-Class and CNN-Class), thus proving that the LV wall thickness could be formulated as graph to accurately predict the presence/absence of VA. The 3D models also underperformed compared to the 2D model, which further demonstrated the limit of the 3D CNNs in the classification task of the 3D anatomy such as the LV wall.

For the FCN and GCN models, the point-to-point correspondence was crucial for the models optimisation, as proven by the plummet in the accuracy of the models trained with non atlas fitting inputs. Without the index pooling, the 500-node models (FCN and GCN) outperformed the 5000-node models, however their accuracy could only reach 0.78. The prediction accuracy was increased for the 500-node models with additional node coordinates as the input feature. By adding the index pooling layer, the GCN-Pooling was able to properly learn the more complex heterogeneity of the 5000-node inputs and achieved higher prediction accuracy.

The univariate analysis on the training population clinical characteristics resulted in 4 significant variables ($p < 0.01$): LVEF, gender, scar size and scar age. The optimal cutoff values calculated from the dataset were 43%, 65cm^2 and 132months for the LVEF, scar size and scar age, respectively.

The evaluation results on the testing dataset are shown in Tab. 2. The 5000-node GCN-Pooling model achieved the highest prediction accuracy at 0.812, followed by the scar-based markers at 0.725 (scar size) and 0.689 (scar age). The LVEF had the lowest performance at 0.671 accuracy.

4.1 Model Interpretability

The Fig. 2 shows the output coefficients of LIME and integrated gradients on a true positive prediction from the testing population. We could note the consistency of the high coefficient (blue) regions between the two methods. The high coefficient regions were also faithful among LIME outputs when the different segmentation methods were used.

We observed that the regions with prominent coefficient correlated strongly with the thinning regions of the LV MYO. This validates the hypothesis that the post-infarction LV thinning is associated with the presence of VA. On top of that, a closer inspection of the top coefficient patches showed that the GCN-Pooling model did not base its VA+ prediction on the entire thinning regions. Rather, higher coefficients were focused on smaller distinctive regions within or adjacent to the scar regions, as indicates by the yellow ellipsoids in the Fig. 2.

Table 1. 10-fold cross-validation of the DL baseline models. The results were displayed as *mean* (\pm *std*). (*) Non average fitted input, (**) Non average fitted input with position and thickness features. **Bold row:** best accuracy model.

Model	Input Shape	Accuracy	Sensitivity	Specificity
CNN-class	256×256	0.775 (± 0.05)	0.616 (± 0.10)	0.834 (± 0.06)
	$240 \times 240 \times 288$	0.759 (± 0.06)	0.508 (± 0.19)	0.853 (± 0.05)
(C)VAE-Class	256×256	0.788 (± 0.06)	0.691 (± 0.18)	0.825 (± 0.05)
	$240 \times 240 \times 288$	0.775 (± 0.06)	0.457 (± 0.18)	0.875 (± 0.03)
FCN	$N = 500^*$	0.604 (± 0.04)	0.525 (± 0.15)	0.643 (± 0.09)
	$N = 500^{**}$	0.700 (± 0.06)	0.266 (± 0.17)	0.862 (± 0.05)
	$N = 500$	0.770 (± 0.06)	0.716 (± 0.18)	0.790 (± 0.07)
	$N = 5000^*$	0.706 (± 0.07)	0.250 (± 0.11)	0.878 (± 0.09)
	$N = 5000^{**}$	0.743 (± 0.03)	0.120 (± 0.09)	0.975 (± 0.03)
	$N = 5000$	0.761 (± 0.12)	0.775 (± 0.17)	0.756 (± 0.17)
GCN	$N = 500^*$	0.701 (± 0.05)	0.626 (± 0.09)	0.742 (± 0.11)
	$N = 500^{**}$	0.718 (± 0.07)	0.766 (± 0.17)	0.700 (± 0.09)
	$N = 500$	0.784 (± 0.08)	0.716 (± 0.17)	0.809 (± 0.09)
	$N = 5000^*$	0.679 (± 0.12)	0.633 (± 0.17)	0.696 (± 0.21)
	$N = 5000^{**}$	0.677 (± 0.09)	0.566 (± 0.20)	0.718 (± 0.16)
	$N = 5000$	0.779 (± 0.08)	0.716 (± 0.19)	0.803 (± 0.10)
GCN-Pooling	$N = 500$	0.784 (± 0.05)	0.808 (± 0.07)	0.775 (± 0.06)
	$N = 5000$	0.818 (± 0.06)	0.766 (± 0.10)	0.837 (± 0.07)

Table 2. Evaluation results on the testing population of the GCN-Pooling model and the clinical baselines.

	Accuracy	Sensitivity	Specificity
GCN-Pooling ($N = 5000$)	0.812	0.780	0.824
Scar Size > $65.49cm^2$	0.725	0.683	0.741
Scar Age > 132months	0.698	0.839	0.640
LVEF < 43%	0.671	0.707	0.657

5 Discussion

Although, the testing dataset were not used during the designing and tuning of the model, the cross-validation and evaluation results were still based on mono-centre dataset. The retrospective nature of dataset could also introduce additional biases relating to the center imaging guideline, which skewed toward a selected population. Therefore, the current study population might not be the realistic representation of the general public.

The interpretability analysis only provided a perspective into the relation between the model prediction and the input. The high coefficient regions could be read as highly arrhythmogenic, simply because they were the regions allowing the model to classify the input as VA+. While the choice of the segmentation method in LIME is not trivial, we focused the analysis only on two segmen-

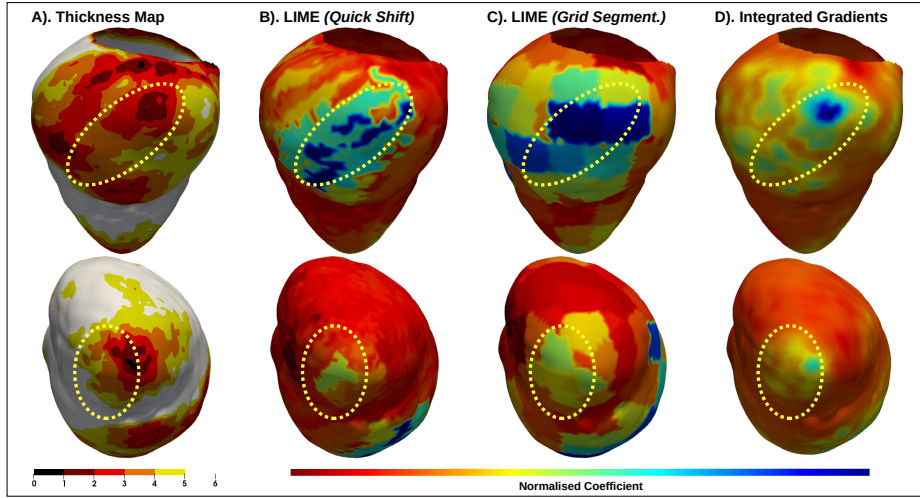


Fig. 2. Interpretability studies of the GCN-Pooling model on a true positive prediction case from testing population. The yellow ellipsoids highlight the regions with consistently high coefficient among different methods. First row: inferior LV view; second row: apex view.

tation methods, to limit the scope of this study. The quick shift method was chosen as it incorporates both physical coordinates and the thickness values of the node and was parameterised to generate small segments (by setting smaller threshold). On contrast, the grid segmentation method does not consider the thickness value and was parameterised to generate larger segments (using larger neighbour searching radius). Therefore, the consistency of the two LIME methods, as well as the integrated gradients, further solidify the tolerance of our interpretability approaches toward superpixel segmentation bias. Nevertheless, limiting by the retrospective aspect of the available data, the current analysis of the interpretability outputs were restricted to hypothesising.

A prospective and multi-centre dataset is necessary to confirm the generalisability of the prediction pipeline and the validity of the detected arrhythmogenic site.

6 Conclusion

We proposed a novel automatic pipeline for VA prediction using the cardiac CT images. Formulating the LV thickness map into a graph input allowed the model to perform an accurate prediction of VA based on the 3D anatomy of the LV and achieved a better score than the model trained on the 2D input. The evaluation on the testing dataset showed that the GCN-Pooling model outperformed the clinical markers, especially comparing to the current gold standard predictor (LVEF). The interpretability studies proved that the model derived its prediction

from the very specific regions within the greater scar regions, an encouraging evidence for the future work on the personalised identification of arrhythmogenic site of the LV.

References

1. Bône, A., Louis, M., Martin, B., Durrleman, S.: Deformetrica 4: An Open-Source Software for Statistical Shape Analysis. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (2018)
2. Garreau, D., Mardaoui, D.: What does LIME really see in images? In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, PMLR (2021)
3. Kipf, T.N., Welling, M.: Semi-Supervised Classification with Graph Convolutional Networks. 5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings (2017)
4. Ly, B., Finsterbach, S., Nuñez-Garcia, M., Cochet, H., Sermesant, M.: Scar-Related Ventricular Arrhythmia Prediction from Imaging Using Explainable Deep Learning. In: Ennis, D.B., Perotti, L.E., Wang, V.Y. (eds.) Functional Imaging and Modeling of the Heart. Springer International Publishing, Cham (2021)
5. Mahida, S., Sacher, F., Dubois, R., Sermesant, M., Bogun, F., Haïssaguerre, M., Jais, P., Cochet, H.: Cardiac Imaging in Patients with Ventricular Tachycardia. *Circulation* (2017)
6. Nielsen, J.C., Lin, Y.J., de Oliveira Figueiredo, M.J., Shamloo, A.S., Alfie, A., Boveda, S., Dagres, N., Di Toro, D., Eckhardt, L.L., Ellenbogen, K., Hardy, C., Ikeda, T., Jaswal, A., Kaufman, E., Krahn, A., Kusano, K., Kutiyifa, V., Lim, H.S., Lip, G.Y., Nava-Townsend, S., Pak, H.N., Diez, G.R., Sauer, W., Saxena, A., Svendsen, J.H., Vanegas, D., Vaseghi, M., Wilde, A., Jared Bunch, T., Buxton, A.E., Calvimontes, G., Chao, T.F., Eckardt, L., Estner, H., Gillis, A.M., Isa, R., Kautzner, J., Maury, P., Moss, J.D., Nam, G.B., Olshansky, B., Pava Molano, L.F., Pimentel, M., Prabhu, M., Tzou, W.S., Sommer, P., Swampillai, J., Vidal, A., Deneke, T., Hindricks, G., Leclercq, C.: European Heart Rhythm Association (EHRA)/Heart Rhythm Society (HRS)/Asia Pacific Heart Rhythm Society (APHRS)/Latin American Heart Rhythm Society (LAHRS) expert consensus on risk assessment in cardiac arrhythmias: Use the right tool for the right outcome, in the right population. *Europace* (2020)
7. Nuñez-Garcia, M., Cedilnik, N., Jia, S., Sermesant, M., Cochet, H.: Automatic Multiplanar CT Reformatting from Trans-Axial into Left Ventricle Short-Axis View. In: STACOM 2020 - 11th International Workshop on Statistical Atlases and Computational Models of the Heart. Lima, Peru (Oct 2020)
8. Ribeiro, M.T., Singh, S., Guestrin, C.: “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '16, Association for Computing Machinery, New York, NY, USA (2016)
9. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (2015)
10. Simonovsky, M., Komodakis, N.: Dynamic Edge-Conditioned Filters in Convolutional Neural Networks on Graphs. Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017 (2017)

11. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic Attribution for Deep Networks. In: Precup, D., Teh, Y.W. (eds.) Proceedings of the 34th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 70, pp. 3319–3328. PMLR (06–11 Aug 2017)
12. Unal, I.: Defining an Optimal Cut-Point Value in ROC Analysis: An Alternative Approach. Computational and Mathematical Methods in Medicine (2017)
13. Valette, S., Chassery, J.M.: Approximated Centroidal Voronoi Diagrams for Uniform Polygonal Mesh Coarsening. Computer Graphics Forum (2004)
14. Vedaldi, A., Soatto, S.: Quick Shift and Kernel Methods for Mode Seeking. In: In European Conference on Computer Vision, volume IV (2008)
15. Yezzi, A.J., Prince, J.L.: An Eulerian PDE Approach for Computing Tissue Thickness. IEEE Transactions on Medical Imaging (2003)
16. Zhang, X.M., Liang, L., Liu, L., Tang, M.J.: Graph Neural Networks and Their Current Applications in Bioinformatics. Frontiers in Genetics (2021)