

Atrial Fibrosis Distribution Generation based on the Diffusion Models

Alexander M. Zolotarev¹, Caroline H. Roney¹

¹ Queen Mary University of London, UK

Abstract

Deep learning (DL) models have the potential to accurately predict atrial fibrillation (AF) ablation outcomes based on patient-specific anatomical and physiological features, including fibrotic remodelling. However, it is challenging to collect enough data for training deep learning models using clinical data alone. With this motivation, we developed a method to artificially generate additional datasets. In this study, we aim to generate artificial atrial fibrosis distributions via diffusion models to increase our training dataset size by imitating independent personalized AF episodes. Specifically, a Denoising Diffusion Probabilistic Model was trained on real fibrosis distributions to generate 100 artificial fibrosis distributions using a 2D representation of the atria (Universal Atrial Coordinates). AF episodes before and after pulmonary vein isolation (PVI) ablation were simulated on left atrial anatomy incorporating these fibrosis maps.

We tested whether these generated fibrosis distributions were suitable for extending a training dataset for a DL binary classifier, which predicts whether AF is sustained post-PVI ablation. Fibrosis and dominant frequency maps extracted from pre-ablation AF simulations were used as inputs to predict AF sustainability after PVI. We compared training a binary classifier on generated fibrosis data to training it on 100 real fibrosis distributions on the same LA anatomy. For the baseline classifier trained and tested on real fibrosis data, the ROC-AUC score was 0.96. In contrast, training the classifier using generated cases, and testing on the real fibrosis data, resulted in an ROC-AUC score of 0.92. The results indicate that the artificial fibrosis distributions correspond well with the real ones and can be used for dataset expansion. Code will be available at https://github.com/pcmlab/cinc23_qmul.

1. Introduction

Atrial Fibrillation (AF) remains the most common heart rhythm disorder and is associated with an increased risk of stroke, with currently estimated prevalence of AF in adults between 2% and 4% [1]. The surgical treatment of AF is ablation, with the main goal being to eliminate the patho-

logical sources of AF signals. Conventional AF ablation therapy is pulmonary vein isolation (PVI) which allows to blockage signals from of the most frequent source of AF paroxysms [2].

Biophysical simulations of AF, when combined with Deep Learning (DL) algorithms, can create patient-specific models for optimizing the strategy of ablation or predicting the success rate of ablation on clinical timescales. These biophysical simulations utilize the cardiac monodomain equation for wavefront propagation coupled to a human atrial cell model, solved on atrial anatomies (meshes constructed from imaging data) with different fibre and fibrosis distributions. It has been shown that atrial fibrosis is correlated with the conduction velocity of electrical propagation and may be used as a prediction metric for ablation success rate [3], indicating that fibrosis plays an important role in AF maintenance.

There are several ways to analyze the resulting biophysical simulations – from testing the different ablation strategies or patient-specific effects of AF drugs to predicting the long-term freedom rate from AF. To run this analysis, one can feed the anatomical and physiological features of AF simulation to a DL prediction pipeline. However, data requirements for training a robust pipeline are high.

One of the possible solutions to address this issue of large data requirements is to create realistic artificial representations of atrial fibrosis that can capture patient variability in electrophysiological features for successful biophysical simulations. Generating these fibrosis distributions at scale enables large *in-silico* trials for evaluating existing and novel treatment approaches.

2. Methods

2.1. Biophysical atrial simulations

We used 100 patient-specific atrial models constructed from late gadolinium enhancement (LGE) MRI data from our previous study [4]. Briefly, personalized models were constructed for 100 atrial fibrillation (43 paroxysmal, 41 persistent, 16 long-standing persistent) patients, undergoing first ablation. The left atrial surface meshes were initially derived from MRI scans with consecutive manual

segmentation of the atrial blood pools and reconstruction of an atrial surface mesh.

A 2D representation of the 3D atrial surface, *universal atrial coordinates (UAC)*, is calculated for each mesh by solving a Laplace equation with boundary conditions; the full pipeline is described in [5]. UAC are necessary for registering scalar and vector features, such as fiber and fibrosis distributions, across different anatomies. We also separate the atrial tissue into different regions with varied electrophysiology parameters using the calculated UAC.

To create a bi-layer model, the endocardial mesh of the Left Atrium (LA) was duplicated for each case and was projected at the same distance from the original mesh for each node to construct the epicardial surface mesh, and the corresponding nodes were connected through line elements. A fiber field was added to the endocardial and epicardial surfaces of each mesh from a human diffusion tensor MRI dataset [6].

Original fibrosis distributions (n=100) were obtained from late-gadolinium enhancement (LGE) MRI, as described in [4]. Next, LGE MRI distributions, measured in image intensity ratio (IIR) relative to the blood pool intensity, were assigned to the bi-layer models through UAC registration. Based on these distributions the ionic and conductivity properties of regions with values above the threshold were changed to represent the fibrotic remodelling, following our previous studies [7].

Finally, the differential equation is numerically solved based on the conduction properties of different atrial regions and the fibrosis distributions. We utilised the Courtemanche human atrial model [8] and monodomain tissue model for AF simulation using openCARP solver. AF was initiated with conditions corresponding to four Archimedean spiral waves on the atrial surface, with more details at [9]. We also simulate PVI ablation for AF episodes by adding two non-conducting rings around the left and right pulmonary vein antra [4].

The simulated AF episodes (15-sec duration) were processed by normalization and applying the Fourier transformation to create Dominant Frequency maps. These maps show the frequency of the highest peak of the Fourier spectrum (the most common frequency in signal) and therefore may indicate the possible source of the wavefront. 3D coordinates of mesh nodes were converted to 2D coordinates based on the UAC to create a 48 by 48 pixels map which can be used as input to the Machine Learning prediction pipeline.

2.2. Generative models

Diffusion models have been successful in various computer vision and augmented reality applications. The core mechanism of these models is the generation of artificial images by restoring them from noise distributions. Gaus-

sian noise is gradually added to the training images, and the model learns to reverse this process [10].

For the generation of artificial fibrosis distributions, we used the Denoising Diffusion Probabilistic Model (DDPM) from MONAI Generative Models software [11]. It was trained on 100 original fibrosis distributions in 2D in the format of 48 by 48 pixel maps. The training was conducted on NVIDIA GeForce RTX 3080 video card for 500 epochs using Mean Squared Error loss and Adam optimizer with a learning rate of 2.5×10^{-5} . As a backbone architecture, we used the default model - the original DDPM scheduler containing 1000 timesteps in its Markov chain, and a 2D U-Net with attention mechanisms in the 2nd and 3rd levels, each with 1 attention head.

2.3. Deep Learning Predictions

The Deep Learning pipeline was used to predict AF sustainability after PVI based on features from pre-ablation AF simulations. The area under the curve for the receiver operating characteristic (ROC-AUC) was used to assess the prediction ability of the Deep Learning model across the testing dataset. The ground truth for the pipeline (labels of AF sustainability) was assessed manually by reviewing the movie of the AF simulation using meshalyzer software.

The pipeline structure is shown in Figure 1. There are two heads for anatomical (2D fibrosis map) and physiological (Dominant Frequency map) inputs. Each head consists of 3 convolutional blocks with the following ReLU activation function and dropout layer. The convolution with 3 by 3 kernel, ReLU activation function, batch normalization and max pooling construct the convolution block. The outputs of both branches were fed to the fully connected layer followed by sigmoid activation.




Figure 1. Deep learning pipeline for AF sustainability prediction (probability, p) after PVI ablation based on pre-ablation fibrosis and DF maps. PVI - Pulmonary Vein Isolation, DF - Dominant Frequency.

The Binary Cross Entropy loss function and Adam optimiser were used for training. Each dataset was separated into training and testing sets with a ratio of 90:10. The result of the pipeline is the prediction of whether the AF simulation is sustained (the probability between 0 and 1). Our motivation to choose this architecture was based on the ability of convolution layers to extract meaningful features from images. We did not use a more complicated convolution network because of the small size of the training dataset and the risk of overfitting.

We also utilize this pipeline to check if the AF simulations, run on atrial meshes with artificial fibrosis distributions, are relevant to ones based on real distributions. For this goal, we trained the model on artificial cases and tested it on real ones. Then we compared the ROC-AUC with ROC-AUCs of training and testing within one domain (AF episodes of either real or artificial fibrosis distributions).

3. Results & Discussion

We generated 100 artificial fibrosis distributions; 5 random examples of them are shown on the bottom row in Figure 2. Firstly, the artificial cases were compared with the real distributions to find possible differences. The mean intensity of real distributions is 1.09 versus 1.28 IIR for artificial distributions.




Figure 2. Randomly selected examples of fibrosis distributions in 2D - real at the top and artificial at the bottom.

Secondly, we performed AF simulations with artificial distributions. The number of sustained AF episodes (5 seconds) is 33 (out of 100) cases in the artificial dataset. For comparison, while running simulations with real distributions, 40 cases are sustained AF episodes.

Thirdly, we checked if we can successfully predict the AF sustainability based on anatomical and physiological features of AF simulations. Initially, we performed the experiments separately on both datasets.

We trained two models independently on the training sets of real and artificial datasets and tested on the testing sets without merging them. The resulting ROC-AUC metrics are shown in Table 1; they are close or equal to 1 for both datasets. It is likely that the ROC-AUC of 1 for the artificial dataset indicates the relative similarity of ar-

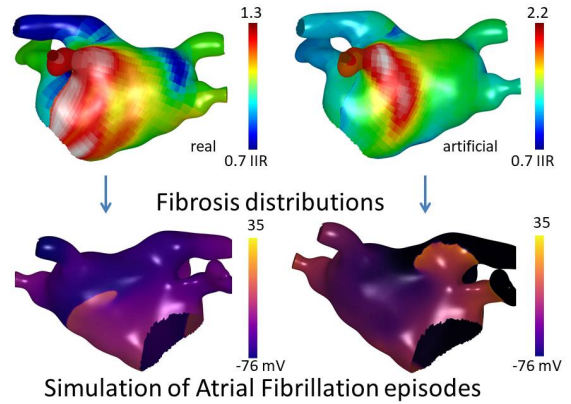


Figure 3. The top row shows fibrosis distributions in IIR (LGE-MRI image intensity ratio); the bottom row - example snapshots of AF simulations using these distributions. A real fibrosis distribution is shown on the left column and an artificially generated one on the right.

tificial cases meaning that it is easy for the Deep Learning model to learn the correct answer for this dataset.

Finally, we tested if we could use artificial distributions for the expansion of the dataset with real fibrosis distributions. The model was trained on 100 artificial distributions and tested on 100 real distributions to predict AF sustainability. The ROC-AUC was 0.92, which shows the high level of similarity between the two domains and the possibility of artificial expansion of the *in-silico* datasets.

We also checked if there is a real need for Diffusion models in obtaining artificial fibrosis distributions. To do this, 100 random Gaussian noise distributions were generated by Pytorch and used as fibrosis distributions for AF simulations as described in Section 2.1. We found that the percentage of AF sustained cases is higher (45 %) than for AF cases of real fibrosis distributions (33 %) and those of artificially generated by the Diffusion model (40 %). However, testing on real cases (n=100) after training on random ones shows that ROC-AUC achieves 0.89 in comparison with training on artificially generated cases (ROC-AUC = 0.92). This indicates the advantage of the developed methodology for generating artificial fibrosis distributions.

4. Conclusion

This study is the first to our knowledge to generate artificial fibrosis distributions for AF computer simulations. This approach may be used to generate large virtual populations for *in-silico* trials of different ablation therapies, and for improving patient-specific ablation prediction using digital twins.

	Fibrosis distributions		
	Train on real	Train on artificial	Train on random
Percentage of sustained post-PVI AF episodes	40	33	45
ROC-AUC if train and test on the same domain	0.96	1	1
ROC-AUC if test on real cases	0.96	0.92	0.89

Table 1. The comparison between AF simulations after PVI ablation based on different domains of fibrosis distributions - real, artificial and random (Gaussian noise). ROC-AUC shows the performance of prediction of AF post-PVI sustainability by DL pipeline. ROC-AUC - the area under the curve for the receiver operating characteristic.

Acknowledgments

Dr Caroline Roney acknowledges support from a UKRI Future Leaders Fellowship (MR/W004720/1). We acknowledge Archer2 simulation funding and the open-CARP software team.

References

[1] Hindricks G, Potpara T, Dagres N, Arbelo E, Bax JJ, Blomström-Lundqvist C, Boriani G, Castella M, Dan GA, Dilaveris PE, Fauchier L, Filippatos G, Kalman JM, La Meir M, Lane DA, Lebeau JP, Lettino M, Lip GYH, Pinto FJ, Thomas GN, Valgimigli M, Van Gelder IC, Van Putte BP, Watkins CL, Group ESD. 2020 ESC Guidelines for the diagnosis and management of atrial fibrillation developed in collaboration with the European Association for Cardio-Thoracic Surgery (EACTS): The Task Force for the diagnosis and management of atrial fibrillation of the European Society of Cardiology (ESC) Developed with the special contribution of the European Heart Rhythm Association (EHRA) of the ESC. *European Heart Journal* 08 2020; 42(5):373–498. ISSN 0195-668X.

[2] Haïssaguerre M, Jaïs P, Shah DC, Takahashi A, Hocini M, Quiniou G, Garrigue S, Le Mouroux A, Le Métayer P, Clémenty J. Spontaneous initiation of atrial fibrillation by ectopic beats originating in the pulmonary veins. *New England Journal of Medicine* 1998;339(10):659–666.

[3] Althoff TF, Eichenlaub M, Padilla-Cueto D, Lehrmann H, Garre P, Schoechlin S, Ferro E, Invers E, Ruile P, Hein M, Schlett C, Figueras i Ventura RM, Prat-Gonzalez S, Mueller-Edenborn B, Bohnen M, Porta-Sanchez A, Tolosana JM, Guasch E, Roca-Luque I, Arbelo E, Neu-

mann FJ, Westermann D, Sitges M, Brugada J, Arentz T, Mont L, Jadidi A. Predictive value of late gadolinium enhancement cardiovascular magnetic resonance in patients with persistent atrial fibrillation: dual-centre validation of a standardized method. *European Heart Journal Open* 12 2022;3(1):oeac085. ISSN 2752-4191.

[4] Roney CH, Sim I, Yu J, Beach M, Mehta A, Solis-Lemus JA, Kotadia I, Whitaker J, Corrado C, Razeghi O, Vigmond E, Narayan SM, O’Neill M, Williams SE, Niederer SA. Predicting atrial fibrillation recurrence by combining population data and virtual cohorts of patient-specific left atrial models. *Circulation Arrhythmia and Electrophysiology* 2022;15(2):e010253.

[5] Roney CH, Pashaei A, Meo M, Dubois R, Boyle PM, Trayanova NA, Cochet H, Niederer SA, Vigmond EJ. Universal atrial coordinates applied to visualisation, registration and construction of patient specific meshes. *Medical Image Analysis* 2019;55:65–75. ISSN 1361-8415.

[6] Roney CH, Bendikas R, Pashakhanloo F, Corrado C, Vigmond EJ, McVeigh ER, Trayanova NA, Niederer SA. Constructing a human atrial fibre atlas. *Annals of Biomedical Engineering* 2020;49:233 – 250.

[7] Beach M, Sim I, Mehta A, Kotadia I, O’Hare D, Whitaker J, Solis-Lemus JA, Razeghi O, Chiribiri A, O’Neill M, Williams S, Niederer SA, Roney CH. Using the universal atrial coordinate system for mri and electroanatomic data registration in patient-specific left atrial model construction and simulation. In Ennis DB, Perotti LE, Wang VY (eds.), *Functional Imaging and Modeling of the Heart*. Cham: Springer International Publishing, 2021; 629–638.

[8] Courtemanche M, Ramirez RJ, Nattel S. Ionic mechanisms underlying human atrial action potential properties: insights from a mathematical model. *American Journal of Physiology Heart and Circulatory Physiology* 1998; 275(1):H301–H321. PMID: 9688927.

[9] Roney CH, Beach ML, Mehta AM, Sim I, Corrado C, Bendikas R, Solis-Lemus JA, Razeghi O, Whitaker J, O’Neill L, Plank G, Vigmond E, Williams SE, O’Neill MD, Niederer SA. In silico comparison of left atrial ablation techniques that target the anatomical, structural, and electrical substrates of atrial fibrillation. *Frontiers in Physiology* 2020;11. ISSN 1664-042X.

[10] Ho J, Jain A, Abbeel P. Denoising diffusion probabilistic models. *arXiv preprint arxiv200611239* 2020;.

[11] Pinaya WHL, Graham MS, Kerfoot E, Tudosiu PD, Dafflon J, Fernandez V, Sanchez P, Wolleb J, da Costa PF, Patel A, Chung H, Zhao C, Peng W, Liu Z, Mei X, Lucena O, Ye JC, Tsaftaris SA, Dogra P, Feng A, Modat M, Nachev P, Ourselin S, Cardoso MJ. Generative ai for medical imaging: extending the monai framework, 2023.

Address for correspondence:

Alexander M. Zolotarev
Queen Mary University of London
327 Mile End Rd, Bethnal Green, London, UK E1 4NS
a.zolotarev@qmul.ac.uk