Cardiac Activation Mapping using Physics-Informed Neural Networks Incorporating Wave Propagation Dynamics

Mohamed A. M. Gad[1], Mazen A. A. Atlam[1], Anas M. A. Elsheikh[1], Sherif Elgendy[1], Ahmed Abdelghafar[1]

Under the supervision of **Dr. Ahmed S. A. Mohamed**[2]

Systems and Biomedical Engineering Department
Faculty of Engineering
Cairo University

[1]Department of Systems and Biomedical Engineering, Faculty of Engineering, Cairo University
[2]Department of Engineering Mathematics and Physics, Faculty of Engineering, Cairo University

Abstract

Atrial fibrillation (AF) is the most common heart arrhythmia, affecting millions worldwide. Diagnosis and treatment of AF often involves creating electro-anatomic activation maps, which represent the timing of tissue activation across the heart's atria. Current mapping methods use interpolation techniques like linear or Gaussian process regression based on sparse electrode data collected within the atria. However, these techniques suffer from noise from electrode positioning and lack of prior physical knowledge of cardiac wave propagation, leading to suboptimal diagnostic accuracy. To address these challenges, we propose a physics-informed neural network (PINN) for cardiac activation mapping that incorporates the underlying wave propagation dynamics of cardiac electrical activity. Benchmarking against traditional interpolation and Gaussian process regression, the PINN model demonstrated improved diagnostic accuracy, paving the way for improved procedural efficiency and patient outcomes in atrial fibrillation diagnostics.

1 Introduction

Atrial fibrillation is the most common arrhythmia in the heart, affecting between 2.7 and 6.1 million people in the United States alone[1]. A standard procedure to diagnose and treat atrial fibrillation is the acquisition of electrical activation maps, where a catheter is inserted into the cardiac chamber and the electrode at the tip records the activation time of the tissue at a given location. This process is repeated at multiple sites to cover the entire atrium.

A study[2] identified the optimal sampling densities required for accurate activation mapping across varying levels of complexity, emphasizing the need for higher measurement densities in complex scenarios. These measurements are then interpolated to create a complete electro-anatomic map of the chamber. The most common approach to interpolate the data is to use linear functions[2, 3] or radial basis functions (RBF)[4]. Unlike traditional methods relying on linear functions, A study [3] proposed an approach that uses Gaussian Markov random fields to account for observation errors and provide probabilistic local activation time (LAT) maps. The RBF algorithm accurately reconstructed wave propagation patterns in simulated tissues with homogeneous and heterogeneous conduction properties, consistently with the data access afforded by clinical practice. These preliminary results suggest the possible integration of the method with clinically used mapping systems to favor the identification of specific propagation patterns and conduction disturbances. However, the two approaches do not incorporate the underlying physics of electrical wave propagation, focusing instead on statistical interpolation and error quantification. This can result in unrealistic interpolations with artificially high conduction velocities.

A strategy[5] incorporating the underlying physics of cardiac activity enables accurate and scalable ECG simulations by decoupling the computationally intensive lead-field calculation from the subsequent simulation process. This approach is more efficient than full-torso solutions for larger datasets, demonstrating how physics-based modeling can streamline computational workflows while preserving realism in heart simulations. However, this study's strategy assumes a constant conduction velocity field, which limits its applicability. Furthermore, no recommendations are provided for acquiring new measurements to reduce procedure time and improve accuracy.

In this paper, we propose a physics-informed neural network for cardiac activation mapping that accounts for the underlying wave propagation dynamics and addresses the limitations of the previously mentioned techniques.

2 Problem Definition

Cardiac activation mapping begins with clinicians using electrodes to measure the electrical activation times at specific points on the heart's surface. However, these data often contain gaps due to spatial limitations, particularly in regions such as fibrotic tissue with slow conduction. Traditional methods that interpolate these gaps fail to account for the underlying physics, leading to inaccuracies and unrealistic conduction velocities.

The electrical activation map of the heart can be related to a traveling wave, where the wavefront represents the location of the depolarizing cells[6]. The time at which cells depolarize is referred to as the activation time and corresponds to an increase in transmembrane potential above a certain threshold and the initiation of the cell contraction. The activation times of the traveling wave must satisfy the Eikonal equation[7, 8]:

$$\|\nabla T(\mathbf{x})\|_2 = \frac{1}{V(\mathbf{x})}$$

The cardiac activation time $T(\mathbf{x})$ represents the activation time at a spatial point $\mathbf{x}$ in the domain. It is the time at which the cardiac tissue at $\mathbf{x}$ depolarizes. The conduction velocity $V(\mathbf{x})$ denotes the local conduction velocity, the speed at which the electrical wave propagates through the cardiac tissue at $\mathbf{x}$. Physically, it depends on properties such as tissue conductivity and anisotropy. The activation time gradient $\nabla T(\mathbf{x})$ represents the direction and steepness of the wavefront at $\mathbf{x}$. The magnitude of the activation time gradient $\|\nabla T(\mathbf{x})\|$ indicates the rate at which the activation time changes rapidly at $\mathbf{x}$.

3 Methodology

Physics-Informed Neural Networks (PINNs) approximate solutions to partial differential equations (PDEs) by embedding the PDE residual and boundary conditions into the network's training process[9, 10]. For a PDE written in a residual form as $F(x, t, u_x, u_t, \dots) = 0$ over a domain Ω with initial condition $u(x, t_0) = u_0(x)$ and Dirichlet boundary condition $u(x, t) = u_\Gamma(t)$, the solution $u(x, t)$ is approximated by a feedforward neural network $\hat{u}_\theta(x, t)$. The network has layers of neurons, nonlinear activations, and trainable weights and biases (θ). The input consists of spatio-temporal coordinates (x, t), while the output approximates $u(x, t)$. Automatic differentiation enables machine-precision computation of derivatives within the PDE, offering a flexible framework for higher-dimensional systems or complex boundary conditions.

Training a PINN involves minimizing a composite loss function[9, 10], $L(\theta)$, which balances contributions from data, boundary conditions, and the PDE residual itself. The loss terms L_{data}, L_b, and L_F are calculated from data points, boundary conditions, and collocation points, respectively, using mean squared error. Weighting factors (ω_{data}, ω_b, ω_F) adjust the relative importance of these terms. Optimization is achieved through gradient descent, commonly using the Adam optimizer, with automatic differentiation facilitating efficient and precise computation of gradients with respect to θ. This process iteratively refines the network to approximate the solution by minimizing $L(\theta)$, leveraging libraries like TensorFlow or PyTorch for implementation.

To enforce the Eikonal equation[7, 8] within the framework of the Physics-Informed Neural Network (PINN), a residual form is defined:

$$R(\mathbf{x}) := V(\mathbf{x})\|\nabla T(\mathbf{x})\| - 1 = 0$$

This residual measures the deviation of the predicted values from the physical constraints described by the Eikonal equation. Minimizing this residual ensures the solution adheres to the underlying wave propagation dynamics. After defining the Eikonal equation, we approximate both the activation time $T(\mathbf{x})$ and conduction velocity $V(\mathbf{x})$ as:

$$T(\mathbf{x}) \approx \mathrm{NN}_T(\mathbf{x}, \theta_T),$$
$$V(\mathbf{x}) \approx \mathrm{NN}_V(\mathbf{x}, \theta_V),$$

where NN_T and NN_V are neural networks with parameters θ_T and θ_V, respectively, that need to be trained to obtain accurate approximations. Since the conduction velocity is strictly positive and bounded within a physiological range, we pass the output of the last layer through a sigmoid function σ, ensuring the conduction velocity neural network reads:

$$V(\mathbf{x}) = V_{\max} \cdot \sigma(NN_V(x))$$

where $V_{\max}$ represents the maximum conduction velocity specified by the user. Finally, we define the loss function used to train the model:

$$L(\theta_T, \theta_V) = \frac{1}{N_T}\sum_{i=1}^{N_T}\left(T(x_i) - \hat{T}_i\right)^2 + \frac{1}{N_R}\sum_{i=1}^{N_R} R(x_i)^2 + \alpha_{TV}\frac{1}{N_R}\sum_{i=1}^{N_R}\|\nabla V(x_i)\| + \alpha_{L2}\sum_{i=1}^{N_\theta}\theta_{T,i}^2,$$

Figure 1 illustrates the first two terms of the loss function: the first term ensures that the neural network's output aligns with the N_T activation time measurements, denoted as $\hat{T}_i$. The second term enforces that the network outputs satisfy the Eikonal equation at N_R collocation points. The third term, evaluated at the N_R collocation points, serves as a total variation regularization for the conduction velocity, allowing for discrete jumps in the solution. This is particularly beneficial for modeling slow conduction regions, such as fibrotic patches. Finally, the fourth term applies L_2 regularization to the weights of the activation time neural network, enhancing generalization and preventing overfitting. We solve the following minimization problem to train the neural networks and determine the optimal parameters:

$$\arg\min_{\theta_T, \theta_V} L(\theta_T, \theta_V).$$

All models are implemented in TensorFlow[11] and optimized using the TensorFlow ADAM optimizer[12] with default parameters and a learning rate of 0.001. To improve computational efficiency, we use a mini-batch implementation, selecting subsets of data points to compute the loss function and its gradient. For lower-dimensional problems, data points are randomly sampled using a Latin hypercube design[13], while for higher-dimensional problems, data is divided into batches, with each batch processed iteratively. This approach ensures efficient handling of large datasets during optimization.

4 Results

To evaluate the effectiveness of the proposed physics-informed neural network (PINN) framework, we designed a synthetic benchmark problem where the activation times and conduction velocity analytically satisfy the Eikonal equation. This benchmark allowed us to assess the method's ability to capture complex wave propagation dynamics, such as discontinuities in conduction velocity and collisions of wavefronts.

The benchmark problem was defined over a two-dimensional domain, with activation times given by:

$$T(x, y) = \min\left(\sqrt{x^2 + y^2}, 0.7\sqrt{(x-1)^2 + (y-1)^2}\right),$$

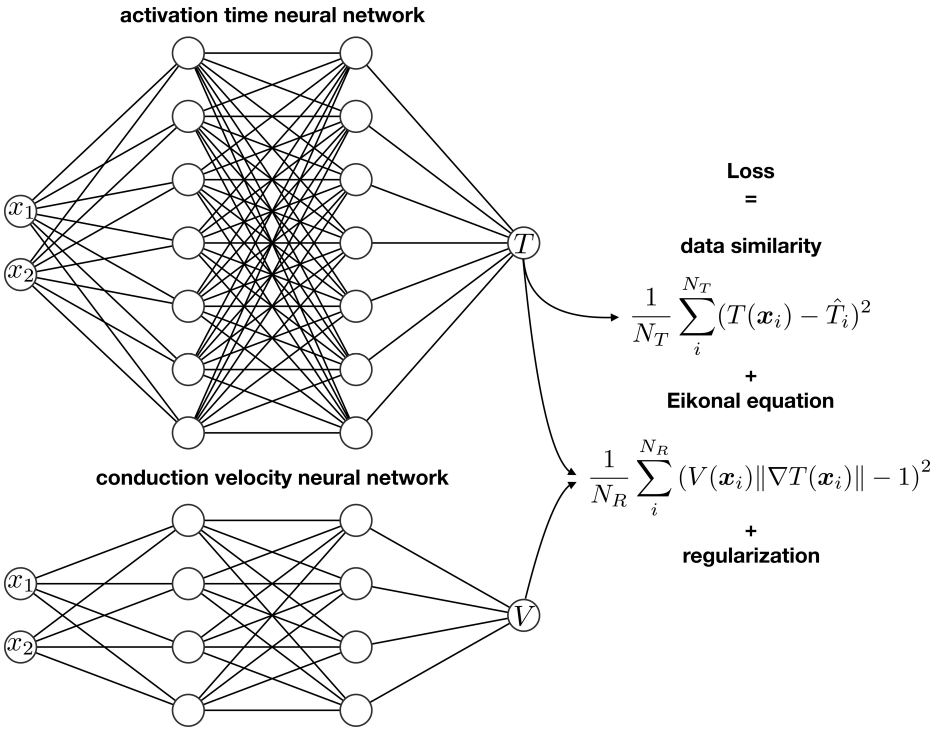

Figure 1: Physics-informed neural networks for activation mapping. We use two neural networks to approximate the activation time T and the conduction velocity V.

and conduction velocity $V(x, y)$ defined as:

$$V(x,y) = \begin{cases} 1.0 \text{ if } \sqrt{x^2+y^2} < 0.7\sqrt{(x-1)^2+(y-1)^2}, \\ \frac{1.0}{0.7} \text{ otherwise.} \end{cases}$$

These functions represent two distinct regions of conduction velocity with a collision of wavefronts in the domain $x, y \in [0, 1]$. **Figure 2**, top left, illustrates the exact mapping of activation times, and **Figure 2**, bottom left, shows the exact mapping of conduction velocity profile. We generate $N = 50$ samples with a Latin hypercube design and train our model. We only have data on the activation times, and we predict both the activation times and the conduction velocity. The results were compared against three other methods; A neural network without physics constraints, Gaussian process regression[14, 15], and Linear interpolation. In the linear interpolation case, we use the `scatteredInterpolant` function from MATLAB with linear extrapolation[2].

The PINN outperformed all baseline methods in both activation time and conduction velocity predictions. **Figure 2** displays the results, highlighting the PINN's ability to accurately capture the wavefront collision and regions of distinct conduction velocities. In contrast, the neural network without physics and Gaussian process regression failed to represent the discontinuity in conduction velocity, resulting in unrealistic gradients. Similarly, linear interpolation produced artifacts highly dependent on the location of the training points.

Quantitatively, **Table 1** summarizes the root mean squared error (RMSE) for activation times and the mean absolute error (MAE) for conduction velocity. We use the root mean squared error (RMSE) for the activation times and the mean absolute error (MAE) for the conduction velocity. We make this distinction to avoid the artificially high errors that will be reported in the root-mean squared error near the discontinuity of conduction velocity. The PINN achieved significantly lower errors compared to the baseline methods, particularly in regions near the wavefront collision where the physics-based regularization proved critical.

we evaluate the performance of all four methods to noise. We introduce Gaussian noise with a standard deviation of 1, 5, and 10% of the maximum value of activation time and run all methods 30 times with the same datasets. **Table 1**, second, third, and fourth row, summarizes the results. We can see that the physics-informed neural network outperforms other methods, except for the 5% noise case in the activation times the Gaussian process regression outperforms PINNs. For the conduction velocity, the physics-informed neural network performs better in all cases by a large margin. Gaussian process regression is as robust to noise as our approach, with similar levels of error for the activation times. Remarkably, the adding physics to the neural network re-

Noise	Quantity	PINN	Neural network	Gaussian process	Linear
0%	T (RMSE)	0.49	0.48	1.56	2.05
	V (MAE)	3.56	9.07	57.78	17.98
1%	T (RMSE)	1.91 (1.09–3.09)	4.19 (2.74–7.93)	1.92 (1.48–2.57)	2.23 (1.97–2.56)
	V (MAE)	10.24 (6.09–17.86)	63.93 (42.74–91.12)	54.08 (47.82–67.72)	27.37 (21.49–40.13)
5%	T (RMSE)	3.42 (2.22–5.34)	11.00 (6.90–18.42)	3.34 (2.63–4.13)	4.44 (3.58–5.76)
	V (MAE)	15.84 (10.40–23.20)	87.33 (68.19–128.83)	78.87 (53.77–102.81)	66.30 (42.97–171.79)
10%	T (RMSE)	6.70 (3.75–10.60)	18.25 (12.84–33.39)	6.73 (3.54–11.65)	8.55 (6.02–11.55)
	V (MAE)	23.16 (10.81–40.78)	90.09 (78.46–119.42)	96.78 (50.49–177.69)	81.06 (59.86–241.80)

Table 1: Performance of physics-informed neural network (PINN), neural network without physics, Gaussian process regression, and linear interpolation in the presence of noise. The root mean squared error (RMSE) for the activation times is normalized by 1 ms and the mean absolute error (MAE) for the conduction velocity is normalized by 1 m/s. Errors are presented as mean and range.

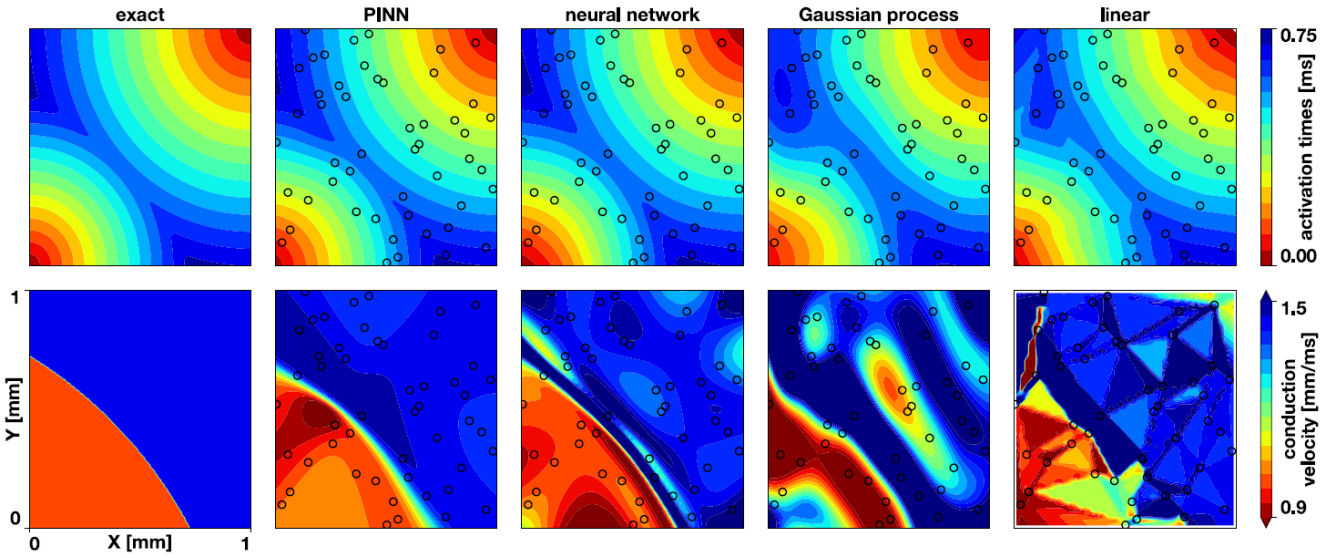

Figure 2: Benchmark problem activation times and conduction velocities. The top row shows the activation times, the bottom row the conduction velocity. The columns display the exact solution and the results of the physics-informed neural network (PINN), a neural network without physics, the Gaussian process regression, and the linear interpolation. The black circles indicate the sampling locations.

duces the error in both activation time and conduction velocity.

These results demonstrate that the incorporation of wave propagation physics into neural network training enables more accurate and physically consistent predictions of cardiac activation mapping compared to traditional methods.

5 Conclusion and Future Work

In this work, we introduced a novel framework for cardiac activation mapping using Physics-Informed Neural Networks (PINNs) to solve the Eikonal equation governing wave propagation in cardiac tissue. By embedding the principles of cardiac electrophysiology directly into the neural network architecture, our method predicts activation times and conduction velocities with high precision, offering a direct, reliable alternative to conventional estimation techniques.

A notable strength of our methodology is its ability to accurately model the dynamics of wavefront propagation, including complex scenarios such as wavefront collisions. This capability marks a significant improvement over traditional approaches, which often fail to represent such interactions accurately.

Our findings suggest that the proposed methodology can have far-reaching implications for cardiac electrophysiology. By directly predicting conduction velocities and activation times, it eliminates intermediate approximations, leading to more accurate and physiologically meaningful maps. Moreover, the computational efficiency of our model makes it suitable for integration into existing clinical workflows, even when using moderately powerful hardware.

Despite its promise, our methodology has some limitations. For instance, it does not currently account for anisotropy in cardiac tissue [6, 16] in activation time measurements. Future work will aim to address these limitations by incorporating fiber orientation data and incorporating uncertainty quantification methods. Additionally, while our results demonstrate strong performance on synthetic data, further validation on real clinical data, particularly for complex arrhythmias, remains an essential next step.

In summary, our study demonstrates the potential of PINNs for cardiac electrophysiology, offering a robust and efficient tool for activation mapping. We believe this approach will significantly benefit the diagnosis and treatment of patients with cardiac arrhythmia, paving the way for faster, more accurate clinical interventions.

References

[1] "Correction to: Heart Disease and Stroke Statistics—2018 Update: A Report From the American Heart Association". In: *Circulation* 137.12 (2018), e493–e493. URL: https://www.ahajournals.org/doi/abs/10.1161/CIR.0000000000000573.

[2] Steven E Williams et al. "Local activation time sampling density for atrial tachycardia contact mapping: how much is enough?" In: *EP Europace* 20.2 (Apr. 2017), e11–e20. ISSN: 1099-5129. DOI: 10.1093/europace/eux037.

[3] Sam Coveney et al. "Probabilistic Interpolation of Uncertain Local Activation Times on Human Atrial Manifolds". In: *IEEE Transactions on Biomedical Engineering* 67.1 (2020), pp. 99–109. DOI: 10.1109/TBME.2019.2908486.

[4] M. Masè and F. Ravelli. "Automatic reconstruction of activation and velocity maps from electroanatomic data by radial basis functions". In: *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*. 2010, pp. 2608–2611. DOI: 10.1109/IEMBS.2010.5626616.

[5] Vincent Jacquemet. "An eikonal-diffusion solver and its application to the interpolation and the simulation of reentrant cardiac activations". In: *Computer Methods and Programs in Biomedicine* 108.2 (2012), pp. 548–558. ISSN: 0169-2607. DOI: https://doi.org/10.1016/j.cmpb.2011.05.003. URL: https://www.sciencedirect.com/science/article/pii/S0169260711001180.

[6] Simone Pezzuto et al. "Evaluation of a Rapid Anisotropic Model for ECG Simulation". English. In: *Frontiers in physiology* 8.MAY (May 2017). ISSN: 1664-042X. DOI: 10.3389/fphys.2017.00265.

[7] P. Colli Franzone, L. Guerri, and S. Rovida. "Wavefront propagation in an activation model of the anisotropic cardiac tissue: asymptotic analysis and numerical simulations". In: *Journal of Mathematical Biology* 28.2 (Feb. 1990), pp. 121–176. ISSN: 1432-1416. DOI: 10.1007/BF00163143. URL: https://doi.org/10.1007/BF00163143.

[8] Piero Colli Franzone and Luciano Guerri. "Spreading of excitation in 3-d models of the anisotropic cardiac tissue. I. validation of the eikonal model". In: *Mathematical Biosciences* 113.2 (1993), pp. 145–209. ISSN: 0025-5564. DOI: https://doi.org/10.1016/0025-5564(93)90001-Q. URL: https://www.sciencedirect.com/science/article/pii/002555649390001Q.

[9] Hubert Baty and Léo Baty. "Solving differential equations using physics informed deep learning: a hand-on tutorial with benchmark tests". working paper or preprint. Apr. 2023. URL: https://hal.science/hal-04002928.

[10] M. Raissi, P. Perdikaris, and G.E. Karniadakis. "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations". In: *Journal of Computational Physics* 378 (2019), pp. 686–707. ISSN: 0021-9991. DOI: https://doi.org/10.1016/j.jcp.2018.10.045. URL: https://www.sciencedirect.com/science/article/pii/S0021999118307125.

[11] Martín Abadi et al. *TensorFlow: A system for large-scale machine learning*. 2016. arXiv: 1605.08695 [cs.DC]. URL: https://arxiv.org/abs/1605.08695.

[12] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG]. URL: https://arxiv.org/abs/1412.6980.

[13] Michael Stein. "Large Sample Properties of Simulations Using Latin Hypercube Sampling". In: *Technometrics* 29.2 (1987), pp. 143–151. DOI: 10.1080/00401706.1987.10488205. eprint: https://www.tandfonline.com/doi/pdf/10.1080/00401706.1987.10488205. URL: https://www.tandfonline.com/doi/abs/10.1080/00401706.1987.10488205.

[14] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, Nov. 2005. ISBN: 9780262256834. DOI: 10.7551/mitpress/3206.001.0001. URL: https://doi.org/10.7551/mitpress/3206.001.0001.

[15] P. Perdikaris. *Gaussian Processes: A Hands-on Tutorial*. Available online at: https://github.com/paraklas/GPTutorial. 2017.

[16] Francisco Sahli Costabal, Daniel E. Hurtado, and Ellen Kuhl. "Generating Purkinje networks in the human heart". In: *Journal of Biomechanics* 49.12 (2016). Cardiovascular Biomechanics in Health and Disease, pp. 2455–2465. ISSN: 0021-9290. DOI: https://doi.org/10.1016/j.jbiomech.2015.12.025. URL: https://www.sciencedirect.com/science/article/pii/S0021929015007332.