

A unified framework for geometry-independent operator learning in cardiac electrophysiology simulations

Bei Zhou¹, Cesare Corrado¹, Shuang Qian¹, Maximilian Balmus^{1,2},
Angela W. C. Lee¹, Cristobal Rodero¹, Marco J.W. Götte³,
Luuk H.G.A. Hopman³, Mengyun Qiao⁴, Steven Niederer^{1, 2}

¹Imperial College London, London, UK.

²Alan Turing Institute, London, UK.

³Amsterdam University Medical Center, Amsterdam, The Netherlands.

⁴University College London, London, UK.

Abstract

Accurate maps of atrial electrical activation are essential for personalised treatment of arrhythmias, yet biophysically detailed simulations remain computationally intensive for real-time clinical use or population-scale analyses. Here we introduce a geometry-independent operator-learning framework that predicts local activation time (LAT) fields across diverse left atrial anatomies with near-instantaneous inference. We generated a dataset of 308,700 simulations using a GPU-accelerated electrophysiology solver, systematically varying multiple pacing sites and physiologically varied conduction properties across 147 patient-specific geometries derived from two independent clinical cohorts. All anatomical and functional data are expressed in a Universal Atrium Coordinate system, providing a consistent representation that decouples electrophysiological patterns from mesh topology. Within this coordinate space, we designed a neural operator with a vision-transformer backbone to learn the mapping from structural and electrophysiological inputs to LAT fields. With a mean prediction error of 5.1 ms over a 455 ms maximum simulation time, the model outperforms established operator-learning approaches and performs inference in 0.12 ms per sample. Our framework establishes a general strategy for learning domain-invariant biophysical mappings across variable anatomical domains and enables integration of computational electrophysiology into real-time and large-scale clinical workflows.

Introduction

Atrial fibrillation (AF) is the most common sustained cardiac arrhythmia, affecting around 46.3 million individuals worldwide [1], and represents a major cause of stroke [2], heart failure [3], and sudden cardiac death [4]. Computational simulations have emerged as powerful tools in AF research and clinical trials, offering insights into arrhythmia mechanisms, enabling the prediction of ablation outcomes, and personalised therapeutic strategies [5–8]. Many of these applications rely on calibrating patient-specific biophysical models to local activation time (LAT) maps, which capture the spatiotemporal propagation of electrical wavefronts across the myocardium. This calibration typically involves solving large ensembles of high-fidelity finite element method (FEM) simulations that numerically solve the governing partial differential equations (PDEs) of cardiac electrophysiology on anatomically accurate patient-specific geometries [9, 10].

Although highly accurate, FEM workflows are computationally expensive and often take hours to days to solve, making them impractical for intraoperative use. Consequently, patient-specific models are usually prepared in advance from pre-procedure MRI or CT scans that define anatomy and tissue properties. In many clinical scenarios, however, such imaging is unavailable, and the only anatomical and electrophysiological data come from intraoperative electro-anatomical mapping (EAM). While EAM-derived LATs can be used for retrospective calibration [11], the turnaround time of FEM-based methods prevents their use for real-time decision-making during ablation.

Recent advances in machine learning have opened new opportunities to accelerate the numerical solution of cardiac electrophysiology models. Physics-informed neural networks (PINNs) have been proposed as a physics-constrained alternative to purely data-driven surrogates, embedding the governing PDE residuals directly into the loss function [12]. They are typically trained for a single PDE configuration, defined by a fixed geometry, parameter set, and boundary conditions [13]. However, they do not naturally generalise to new configurations without retraining: if the domain shape, material parameters, or boundary conditions change, the PINN must be re-trained from scratch or at least fine-tuned.

Neural operators directly address this limitation by learning a mapping between infinite-dimensional function spaces rather than a single instance of a PDE [13]. Foundational neural operators such as the Deep Operator Network (DeepONet) [14] and the Fourier Neural Operator (FNO) [15] have demonstrated strong performance on canonical problems defined over structured domains. By learning this operator mapping, they can efficiently predict solution fields for new inputs, even for unseen geometries or different discretisation resolutions, without the need for time-consuming retraining. Applying neural operators to complex real-world problems, such as predicting cardiac activation maps, presents significant challenges. First, these aforementioned neural operator frameworks are built to handle data on structured, canonical domains, making them ill-suited for diverse and irregular geometries [16, 17]. Second, training a reliable neural operator is contingent on the availability of large, high-quality datasets of diverse simulation results [18], and generating such datasets is a major challenge [19].

Several approaches have been proposed to address the challenge of anatomical variability. Diffeomorphic Mapping Operator Learning (DIMON) [20] uses diffeomorphic transformations to map individual geometries to a shared, standardised space. This allows a neural operator to learn latent solutions on the reference domain, with the original shape encoded as an auxiliary input. While effective for relatively simple geometries, such as the semi-spherical left ventricle (LV), this strategy faces substantial difficulties when extended to anatomies with more intricate topologies, such as the multi-holed structure of the left atrium, where constructing topologically consistent reference domains is non-trivial. Similarly, methods like GeoFNO require a diffeomorphic mapping to simple domains (e.g., cubes or tori) where the Fourier transform can be applied [16]; however, the topologically complex 3D geometries of the human left atrium, with multiple openings, make such mappings impossible [21]. Graph-based methods like the Graph Fourier Neural Kernel (G-FuNK) [21] learn directly on mesh graphs, avoiding domain mapping. However, their spectral bases are sensitive to mesh quality and topology, reducing robustness and generalisability in heterogeneous clinical datasets.

Here, we introduce a method that overcomes the dependency of neural operators on structured domains for electrophysiology simulations. Our approach projects inputs and outputs onto a common, anatomically standardised reference using the Universal Atrial Coordinates (UAC) system [22]. The UAC defines a mesh-independent coordinate space for the left atrium (LA), preserving physiologically relevant spatial relationships while unifying geometry and topology across patients. This representation enables neural operators such as DeepONet and FNO to consistently interpret spatial features, free from irregular mesh geometries or topological inconsistencies.

To support large-scale training, we developed a GPU-accelerated FEM solver, drastically reducing the time required to compute LAT maps and enabling the generation of a high-fidelity, comprehensive dataset. Using this dataset, we train a neural operator to approximate the underlying biophysical mapping, achieving near real-time inference without compromising physiological accuracy. Our model, an encoder-decoder architecture with a Vision Transformer (ViT) [23] backbone, predicts full-field LAT maps from patient-specific anatomical features, conductivity, and pacing sites. It achieves a mean absolute error of 5.1 ± 0.2 ms relative to FEM simulations, exhibiting high pointwise accuracy, spatial smoothness, and robust generalisation across independent clinical cohorts, overcoming the domain shift inherent in multi-centre medical data.

By integrating GPU-accelerated simulation and neural operator training into a unified framework, our approach addresses both the computational bottleneck of large-scale electrophysiology simulations and the geometric variability of patient-specific atrial anatomies. It enables accurate prediction of full-field LAT maps in just 0.12 ms per sample on a GPU, offering a pathway toward rapid, patient-specific, physiology-informed decision support for AF interventions.

Results

Large-scale EP simulations on GPUs

We generated a comprehensive set of biophysically detailed EP simulations using a custom GPU-accelerated FEM solver applied to two independent cohorts of AF patients (Cohort A: $n = 100$; Cohort B: $n = 47$). This process constructed two large-scale datasets, Dataset A and Dataset B, respectively (see Methods). To ensure robustness across diverse anatomical and physiological conditions, for each of the 147 patient-specific LA, we simulated LAT maps from seven distinct pacing sites across the LA. At each site, we sampled 300 combinations of the key conduction parameters, the longitudinal conductivity (σ_l) and the anisotropy ratio (σ_l/σ_t), using Latin hypercube sampling [24] to cover a wide range of clinically relevant tissue properties. This yields a total of 308,700 simulated LAT maps, establishing one of the largest computational datasets of its kind for training neural operators.

All simulated LAT maps were subsequently projected onto the standardised UAC domain to obtain a consistent, geometry-independent representation. Fig. 1a illustrates an LAT map on the 3D atrial surface together with its corresponding 2D UAC projection, with the coronary sinus pacing site producing full atrial activation within 204 ms.

To assess potential cohort-level differences, we performed dimensionality reduction on the full set of UAC-projected LAT maps. High-dimensional LAT fields were first denoised using principal component analysis (PCA), and the resulting principal components were subsequently embedded using Uniform Manifold Approximation and Projection (UMAP) [25]. The 2D embedding (Fig. 1b) reveals a clear separation between the two datasets. Within each cohort, seven well-defined clusters emerge, each corresponding to one of the pacing sites. Clusters associated with the same pacing location in Datasets A and B do not overlap but instead form displaced pairs in the latent space. This consistent displacement provides visual evidence of a domain shift between the cohorts, arising not from physiological differences, but from systematic variations in anatomical distributions, acquisition protocols, or preprocessing pipelines. This domain shift has a measurable impact on cross-domain generalisation performance.

We next quantified the primary determinants of total activation time across all simulations. Three factors exert the strongest influence: pacing site, LA surface area, and conduction velocities. The maximum LAT varies across pacing locations in Dataset A (Fig. 1c) and Dataset B (Fig. 1d). Across both datasets, maximum LAT increases approximately linearly with LA surface area and decreases with increasing conductivity values. These consistent trends confirm that surface area and conductivity jointly modulate global activation timing, motivating their inclusion as key anatomical features in the subsequent analyses.

Input feature sensitivity analysis

We evaluated the contribution of each anatomical input modality by removing spatial coordinates, fibre orientation, and surface area features in turn, and retraining




Fig. 1 Construction and characterisation of large-scale electrophysiology datasets generated from biophysically detailed simulations. **a**, Example local activation time (LAT) map from a 3D left atrium (LA) electrophysiology simulation with coronary sinus pacing, together with its projection onto the standardised Universal Atrial Coordinate (UAC) domain. **b**, UMAP embedding of all UAC-projected LAT maps from Datasets A and B, showing seven pacing-site clusters per cohort and systematic displacement between corresponding clusters, indicating a pronounced domain shift between the two datasets. **c**, Distribution of maximum LAT across pacing locations in Dataset A, demonstrating linear dependence on LA surface area and sensitivity to conductivity. **d**, Corresponding analysis for Dataset B, showing the same physiological trends despite cohort-level differences.

the model under each condition. Performance was assessed using two complementary metrics: Mean Absolute Error (MAE), which quantifies the numerical accuracy (error in milliseconds), and the Structural Similarity Index Measure (SSIM) [26], which measures the spatial fidelity and structural quality of the predicted activation maps. All analyses in this section were performed using 5-fold cross-validation. Unless stated otherwise, experiments were conducted on Dataset A, with cross-domain evaluations involving both Datasets A and B.

The results in Fig. 2a show that incorporating all three feature types achieves accurate and spatially smooth LAT predictions ($\text{MAE } 6.09 \pm 0.26 \text{ ms}$; $\text{SSIM } 0.963 \pm 0.001$). Removing the surface area feature results in a pronounced degradation in performance ($\text{MAE } 9.88 \pm 0.81 \text{ ms}$), indicating that global anatomical scaling is an essential contextual signal for correct wavefront organisation. Similarly, excluding spatial coordinates reduces both numerical accuracy and structural similarity ($\text{MAE } 9.04 \pm 0.44 \text{ ms}$; SSIM




Fig. 2 Comprehensive evaluation of model performance for LAT map prediction. **a**, Our model performance from feature sensitivity tests on Dataset A, demonstrating the contribution of each input feature. **b**, Evaluation of the impact of two regularisation terms, namely total variance (TV) and Laplacian, on LAT map prediction on Dataset A. **c**, Performance comparison between our proposed model and established deep learning architectures, including DeepONet, Fourier Neural Operator (FNO), Wavelet Neural Operator (WNO), U-Net, and ResNet on Dataset A. **d**, Spatial distribution of the averaged LAT prediction error (in ms) across five test cases in a representative validation fold. White circles highlight regions with the highest prediction errors, and adjacent numbers show the average error within those specific regions. **e**, Performance comparison between Single Pacing Location Models and a Unified Model across seven pacing sites. **f**, Cross-domain performance comparison on two distinct datasets. The notation $X \rightarrow Y$ indicates that the model was trained on source dataset X and evaluated on target dataset Y .

0.948 ± 0.002), demonstrating that explicit localisation in the UAC domain is critical for resolving conduction pathways.

In contrast, removing fibre orientation led to improved performance across both metrics ($\text{MAE } 5.15 \pm 0.25 \text{ ms}$; $\text{SSIM } 0.978 \pm 0.002$). This suggests that the fibre orientation information, as currently represented, may introduce noise or redundancy relative to the geometric cues already encoded in the coordinate and surface area channels. We discuss this observation further in the discussion section. Accordingly, fibre orientation features were excluded from all subsequent experiments.

Evaluating the effect of regularisation

To assess the impact of regularisations on LAT map reconstruction, we evaluated the contribution of two common spatial regularisation terms, total variation and second-order Laplacian loss, when added individually or in combination. These regularizers aim to encourage spatial smoothness and suppress high-frequency artefacts in the predicted LAT maps.

To evaluate the impact of different loss formulations on model performance, we conducted systematic experiments by varying the regularisation coefficient $\lambda \in \{5 \times 10^{-2}, 3 \times 10^{-2}, 1 \times 10^{-2}, 5 \times 10^{-3}, 1 \times 10^{-3}\}$ in the loss function.

For each fold in 5-fold cross-validation, the model achieving the best performance on the validation set was retained, and the corresponding metrics were recorded. The performance for a given regularisation weight λ was derived by averaging these metrics across the five folds. To probe the effect of each regularisation configuration, such as TV alone or the combination of TV and Laplacian, we further averaged the fold-averaged metrics across all tested λ values (Fig. 2b).

Among all evaluated configurations, the H^1 loss combined with both TV and Laplacian regularisations achieved the best overall performance, with a mean MAE of 5.16 ms and a mean SSIM of 0.978. This demonstrates that combining both regularisation terms effectively enhances structural fidelity while preserving numerical accuracy. The addition of TV or Laplacian individually yielded smaller improvements, suggesting that their combined effect is necessary to fully capture the relevant spatial structures.

Benchmarking model performance

We benchmarked a range of neural operator frameworks and deep-learning architectures capable of modelling cardiac activation patterns from multimodal inputs. For each model, we selected three primary hyperparameters governing capacity and constructed six configurations by varying these across representative ranges, ensuring that the comparison is both systematic and approximately capacity-matched across architectures (Fig. 2c). Supplementary Section 2 provides further details on architecture design and training procedures.

We first evaluated operator learning approaches. While the Fourier Neural Operator (FNO) [15] is theoretically well-suited for learning global mappings in function space, its empirical performance was moderate (MAE 10.14 ± 0.07 and SSIM 0.938 ± 0.005), indicating potential challenges in handling heterogeneous and noisy physiological data. The Wavelet Neural Operator (WNO) [27] exhibited the weakest performance among evaluated models (MAE 10.08 ± 0.28 and SSIM 0.938 ± 0.007 , likely reflecting sensitivity to measurement noise and anatomical variability inherent in real-world datasets. The DeepONet framework, using our proposed ViT-based architecture for both branch and trunk networks, achieved strong performance (MAE 7.12 ± 0.37 ms; SSIM 0.963 ± 0.005).

Among convolutional baselines, the U-Net architecture [28] outperformed the simpler ResNet model [29]. The U-Net’s multi-scale encoder-decoder structure, with its strong inductive bias for spatial localisation, yielded a balanced performance of MAE

8.48 ± 0.28 ms and SSIM 0.965 ± 0.004 . However, these convolutional models were fundamentally limited in their ability to capture global pacing effects and long-range spatial dependencies

Our model surpassed all baselines, achieving the best results with an MAE of 5.15 ± 0.25 ms and SSIM of 0.978 ± 0.002 . This confirms the model’s superior capability to synthesise heterogeneous inputs into spatially coherent, smooth, and highly accurate LAT map predictions.

While MAE and SSIM provide global measures of model performance, they do not reveal where errors concentrate spatially. We therefore visualised the prediction error distributions across the UAC domain. Fig. 2d displays the spatial distribution of prediction error on LAT maps for six test cases, each averaged over seven pacing locations and 300 combinations of conductivity velocities. Across all six test cases, a visually consistent and highly coherent source of localised prediction error is observed near the horizontal central line of the UAC map. This localised band of elevated error is the most prominent spatial feature in the error distribution plots. Beyond this central band, the prediction errors generally remain low and diffuse across the rest of the UAC domain, indicating high model fidelity in most tissue regions.

Beyond predictive accuracy, we further assessed the computational efficiency of our model to evaluate its suitability for real-time applications. Inference on batches of 1000 samples was performed on an NVIDIA RTX 5090 GPU in evaluation mode with CUDA synchronisation to ensure precise wall-clock timing. The model achieved an average inference time of 122.23 ms per batch, or approximately 0.12 ms per sample. This represents a five to six orders of magnitude acceleration compared with conventional finite-element electrophysiology solvers, which require approximately one minute per simulation with our solver in this study. Such efficiency enables near real-time prediction of LAT maps across diverse atrial anatomies and pacing configurations, highlighting the potential of our model for integration into time-critical clinical workflows and large-scale population studies.

Model generalisation across pacing sites

Understanding whether a single model can generalise across diverse pacing conditions is critical for building scalable and clinically useful predictors of atrial activation, as it allows us to determine whether it can match or exceed the performance of individually trained models. Success in this setting would indicate that the model captures shared electrophysiological structure, enabling robust predictions without prior knowledge of stimulation location.

We thus assessed the predictive performance of two modelling paradigms for LAT map reconstruction: (i) Single Pacing Location Models, trained independently for each of seven pacing sites, and (ii) a Unified Model, trained jointly across all sites to learn a shared inductive bias over anatomical and physiological structure. Across all pacing conditions, the Unified Model yields consistent and statistically significant improvements over its independently trained counterparts. Specifically, it achieves a 16.5% reduction in mean MAE from 6.3 ms to 5.2 ms (Fig. 2e, left), demonstrating enhanced predictive accuracy across heterogeneous inputs. The model also delivers a

1.67% increase in average SSIM from 0.960 to 0.976 (Fig. 2e, right), reflecting superior preservation of spatial structure and clinically relevant activation patterns.

The challenge of latent data heterogeneity

Our initial experiments confirmed the strong performance of our model on a single dataset (Dataset A), demonstrating its ability to capture anatomical features relevant to the target task within a controlled domain. However, high performance within the same dataset does not guarantee translational reliability. Here, we distinguish between internal performance, which measures the model’s accuracy when trained and tested on the same dataset, and cross-domain performance, which measures generalisation to an independent dataset. Evaluating both is critical to determine whether the model has learned biologically meaningful features or relies on dataset-specific cues.

Cross-domain generalisation was evaluated on two levels. First, we measured model performance when trained and tested on the same dataset (internal performance: $A \rightarrow A$ and $B \rightarrow B$) compared with performance when trained on one dataset and tested on the other (cross-domain performance: $A \rightarrow B$ and $B \rightarrow A$). Second, we assessed whether including data from an additional domain ($(A + B) \rightarrow A$, $(A + B) \rightarrow B$, and $(A + B) \rightarrow (A + B)$) during training improves generalisation to the original dataset, using internal performance as a baseline.

Each evaluation was performed on models trained with 18 combinations of hyperparameters to ensure statistical significance, with results highlighting consistent trends (Fig. 2f; full results in Supplementary Section 3). The strong internal performance confirms that the UAC projection effectively separates anatomical representation from patient-specific shape. In contrast, the pronounced drop in cross-domain performance provides clear evidence of latent bias. The model relies on image-specific cues that mislead predictions when applied to a new domain. This finding highlights a second-order challenge in translational modelling. Models are highly sensitive to systematic variations in data representation and annotation rather than true anatomical differences.

Integrating heterogeneous datasets during training mitigates these domain-specific biases. Training on the combined datasets $(A + B)$ encourages the model to focus on consistent anatomical features, ignoring spurious correlations tied to individual acquisition or annotation pipelines. Across both datasets, MAE decreased and SSIM increased in comparison with the internal performance, demonstrating that exposure to independent cohorts acts as a form of regularisation. For completeness, we also evaluated the model trained and tested on the combined datasets $(A + B) \rightarrow (A + B)$, obtaining (5.77 ± 0.16) for MAE and (0.9734 ± 0.003) for SSIM.

Discussion

The increasing complexity of modern biomedical simulations, particularly in fields like computational cardiology [30], has created a growing need for methods that can bridge the gap between high-fidelity mechanistic models and the demands of real-time clinical applications. The results demonstrate that our proposed framework can accurately and robustly learn the complex, nonlinear mapping from a patient’s

anatomy and electrophysiological parameters to their full-field LAT map, achieving high fidelity to biophysical simulations in milliseconds. This represents an advancement over conventional approaches and a crucial step toward enabling in-procedure, real-time, personalised cardiac electrophysiology analysis.

Geometry-independent neural operator learning

Our approach demonstrates that a well-designed coordinate system is a powerful tool for bridging the gap between irregular, real-world data and the structured-domain requirements of neural networks, allowing a single neural operator to learn a geometry-independent mapping. By leveraging the UAC system, we established a canonical representation that enabled neural operators to robustly capture the complex, nonlinear physics of cardiac conduction [31–34]. A key finding is the model’s low MAE of (5.15 ± 0.25) ms. This level of accuracy is significant in a clinical context, as it is well within the typical measurement noise and physiological variability observed in cardiac electroanatomic mapping systems [35] (typically on the order of 5–10 ms [36]). This demonstrates that our model’s predictions are not only numerically accurate but also fall within a clinically acceptable margin, making the approach viable for non-invasive patient stratification, procedure planning, and guiding ablation therapies.

A central tenet of our work is the understanding that the success of a neural operator relies on the availability of a large, diverse, and high-quality training dataset. By developing a GPU-accelerated FEM solver, we were able to generate a large simulation dataset across diverse patient-specific left atrial anatomies, providing a robust foundation for learning generalizable principles of cardiac electrophysiology. This data generation pipeline is a critical component of our framework, as it bypasses the prohibitive computational bottleneck that has historically limited the application of data-driven methods in this domain.

Beyond supporting our framework, this large, curated dataset provides a foundation for future data-driven studies in cardiac electrophysiology. Its scale, diversity, and consistency enable predictive models to generalise across anatomies and protocols, a capability previously limited by data scarcity. For comparison, G-FuNK [21], trained on 25 left atrial meshes from atrial fibrillation patients, concluded that additional geometries could improve performance. DIMON [20] used 1,006 left ventricular patient data, and each was simulated at seven pacing sites, totalling 7,042 simulations. In contrast, our simulations preserve these locations while also varying tissue conductivity, allowing neural operators to capture electrophysiological responses across diverse stimulation patterns and heterogeneous, patient-specific left atrial anatomies. Overall, our dataset spans 308,700 simulations, offering a far richer foundation for training and benchmarking state-of-the-art neural operators.

Comparative model analysis

Our model achieves the best performance, which can be attributed to two key factors. First, the self-attention mechanism enables global context modelling [37], allowing the network to capture long-range dependencies such as conduction wavefront spread from distant pacing sites. Second, by using convolutional patch embeddings and transposed

convolution decoding, ViT maintains strong spatial localisation [38], which is critical for resolving the high-frequency components of activation fronts. This balance between global awareness and local detail explains ViT’s superior pointwise accuracy and its strong preservation of spatial gradients.

We benchmarked our approach against several state-of-the-art neural operators and two different CNN architectures. Our model’s performance surpassed these baselines for reasons rooted in their fundamental design. For instance, DeepONet, which in this work employs our model for its branch or trunk networks, leverages the same global context modelling advantages as our model. This enables DeepONet to perform reasonably well, particularly in SSIM and H^1 loss, by effectively capturing smooth functional mappings and spatial dependencies. However, despite incorporating ViT components, DeepONet’s architecture focuses on decomposing input-output relationships into branch and trunk networks, which may limit its end-to-end spatial locality learning compared to the full ViT model. Consequently, its MAE performance lags behind ViT, reflecting some difficulty in the precise localisation of rapid feature changes critical for accurate cardiac wavefront propagation. In contrast, FNOs are designed to learn global mappings in function space using spectral convolutions. While FNOs can capture global trends efficiently, their reliance on low-frequency Fourier modes makes them ill-suited for capturing sharp transitions and local gradient variations common in LAT maps, particularly near pacing initiation sites. This is reflected in their relatively poor H^1 loss and MAE, despite modest SSIM performance. WNOs attempt to address this by incorporating multi-resolution representations, theoretically improving sensitivity to localised patterns. However, WNO’s performance degrades substantially in this setting, likely due to instability in wavelet decomposition on structured but irregular biomedical domains and sensitivity to small-scale alignment errors, leading to high variance in H^1 loss.

Among the convolutional baselines, U-Net delivers strong SSIM, indicating that it generates visually plausible LAT maps. Its hierarchical multi-scale architecture with skip connections allows for reconstruction of structural details. However, convolutional models like U-Net and ResNet-style CNN architectures are inherently local in receptive field, making it difficult to capture global pacing effects or subtle long-range interactions without very deep networks or large kernels [39].

Toward domain-invariant generalisation

Prior studies have evaluated model performance primarily under variations in pacing site [20]. In contrast, a central contribution of our work is demonstrating that a single model can generalise across both diverse pacing locations and simulation parameters. This capability is essential for clinical applicability, since the pacing site is often variable or unknown in practice. Our results also show that the unified model, trained jointly across all seven pacing sites, consistently and significantly outperforms models trained for individual pacing locations. The unified model achieves a 16.5% reduction in mean MAE, and a 1.6% increase in average SSIM. This demonstrates that our model effectively extracts generalizable electrophysiological principles that are independent of the specific pacing location, making it a robust and scalable solution for real-world clinical applications where prior knowledge of the stimulation site is unavailable.

Our cross-domain experiments reveal a critical insight. The apparent limitation of single-domain models is not a shortcoming of the ViT architecture, but a reflection of latent heterogeneity in clinical imaging data. By explicitly combining datasets from distinct acquisition and annotation protocols, we demonstrate that the model naturally discards protocol-specific shortcuts and focuses on the fundamental anatomical structure. This approach not only rescues performance on external cohorts but also slightly improves internal predictions.

This finding establishes a generalizable strategy for achieving domain-invariant medical image analysis. Our UAC projection provides the necessary geometry-invariant foundation, while training on heterogeneous data acts as a regularization method in feature space, compelling the model to learn only the biophysically relevant structures common to both domains and emphasizing that true translational power derives from integrating diverse, multi-centre data rather than from isolated single-site optimisation.

Limitations and future directions

Despite its strong performance, the present study has limitations. The observation that removing fibre orientation features improved model performance was initially counterintuitive. Our analysis suggests two primary reasons for this finding. First, the current representation of fibre orientation may contain inherent noise or inaccuracies that interfere with the model’s ability to learn the true underlying electrophysiological dynamics. Second, the model appears to be leveraging the spatial and geometric priors provided by the coordinates and surface area features more effectively. These geometric inputs may implicitly contain sufficient information about the overall anatomical structure, making the explicit fibre orientation data redundant or even counterproductive in its current form.

Beyond the feature representation, limitations also arise from the use of the UAC system. The UAC system is designed to map the closed, complex 3D left atrial surface onto a flat, 2D grid. To achieve this flattening, a continuous line must be cut on the 3D surface. Depending on the chosen coordinate orientation, this seam manifests as a central line on the 2D map. This seam creates a spatial discontinuity for the learning model, as points that are physically adjacent on the 3D surface are severed across this line in the 2D domain. Consequently, the neural operator struggles to accurately interpolate the LAT function across this boundary, resulting in a consistent band of large error centred precisely along this necessary central UAC seam.

The additional large localised errors, which are more distributed, are plausibly linked to areas of extreme anatomical complexity and steep electrophysiological gradients. These clusters primarily occur around the highly patient-specific structures like the PV junctions and LAA. The UAC projection cannot fully remove the residual geometric variability present in these complex regions, increasing the model’s uncertainty. Furthermore, activation wavefronts often exhibit the steepest LAT gradients in these areas. The model is challenged to precisely localise and reproduce these sharp functional variations across diverse anatomies, leading to the small, perceptible clusters of higher error in these critical physiological locations. These difficulties suggest that while our model excels at capturing global activation patterns, the fine-grained,

localised details influenced by complex anatomical geometries remain a challenging frontier. Future work could explore incorporating a more fine-grained representation of these critical regions to further enhance local accuracy.

Finally, the framework was trained exclusively on high-fidelity simulated datasets. Although these simulations spanned a wide range of anatomies and electrophysiological parameters, direct validation against clinical electroanatomic mapping data is essential to confirm robustness under real-world conditions. The current implementation focuses on the left atrium; extending the approach to whole-heart or biventricular modelling will require adapting the anatomical projection and scaling the data generation process. Future directions include combining hybrid simulation–clinical datasets, refining anatomical embeddings, and exploring adaptive spatial resolutions to capture fine-scale electrophysiological detail without compromising inference speed.

Methods

We developed a unified, scalable computational framework that overcomes the computational bottlenecks of conventional EP simulations (Fig. 3). This framework is designed to circumvent the computational bottleneck inherent in conventional EP simulations. It leverages a diverse, high-fidelity dataset generated by a custom high-performance FEM solver to enable the training of a neural operator that can predict LAT maps with accuracy comparable to ground truth, but within milliseconds.




Fig. 3 The overview of our proposed computational framework for operator learning for electrophysiology (EP) simulations. **a**, Preparation of left atrium (LA) mesh, mapping the 3D vertex coordinates defining the mesh to a 2D grid by Universal Atrial Coordinates (UAC), and identifying the pacing locations. **b**, The application of our GPU-based FEM solver to produce a large amount of simulated local activation time (LAT) maps, which are also projected to 2D grids by UAC. **c**, The training of neural operators using the structured simulation data.

Image data processing

Datasets description

The full dataset comprised 147 patient-specific left LA anatomies derived from two large-scale clinical cohorts. These anatomies served both as inputs to the EP simulations and as training and evaluation data for the neural operator models. To support both internal evaluation and cross-domain generalisation analysis, the 147 anatomies were partitioned into two distinct cohorts (Cohort A and Cohort B).

Cohort A, consisting of 100 AF LA samples (43 paroxysmal, 41 persistent, and 16 long-standing persistent), was obtained from the AF Recurrence Cohort [40], which was recruited in London, UK. LA geometries were segmented from late gadolinium-enhanced cardiac magnetic resonance (LGE-CMR) images acquired at end-diastole, yielding three-dimensional anatomical surfaces including the pulmonary veins and mitral annulus. Fibrotic tissue distribution was quantified using standard LGE intensity thresholding, providing patient-specific fibrosis maps that informed the EP simulations. This cohort served as the controlled domain for model development and internal validation.

Cohort B, containing 47 LA samples from AF patients recruited from Amsterdam, The Netherlands, was sourced from an independent clinical study [41]. The LA geometries in this cohort underwent the same segmentation and fibrosis quantification pipeline as those in Cohort A to maintain structural consistency.

As LGE-CMR does not resolve myofibre orientation, fibre direction vectors for anisotropic conduction were generated for all 147 anatomies using a model based on high-resolution ex-vivo Diffusion Tensor Magnetic Resonance Imaging (DTMRI) data [42]. These fibre fields were subsequently incorporated into an atrial fibre atlas and projected onto each patient-specific geometry using the UAC system, as described by [43].

Notably, Cohorts A and B were derived from two distinct and independent studies conducted by separate research groups operating within independent healthcare systems. This independence in data acquisition and analysis supports robust cross-domain evaluation. Moreover, despite adhering to equivalent segmentation and processing workflows [44, 45], the two cohorts exhibit meaningful domain shifts. Variations arise from MRI acquisition parameters (scanner vendor, pulse sequences, and field strength), from different practices in anatomical trimming (particularly of pulmonary veins and the LA appendage), and from inter-operator differences in anatomical landmarking and boundary definitions. These protocol-level and anatomical choices shape the resulting geometries and fibrosis patterns, such that models trained exclusively on Cohort A typically encounter Cohort B as an out-of-distribution domain. These differences are not physiological changes. These differences are not physiological. They are artefacts of data preparation, yet they change the geometric domain on which the operator is defined.

UAC projection

We projected the spatial coordinates and fibre directions on the mesh, along with the resulting LAT map, to structured grids of fixed size by UAC. The UAC system [22]

defines a bijective transformation

$$\Phi : \mathcal{M} \rightarrow [0, 1]^2,$$

where $\mathcal{M} \subset \mathbb{R}^3$ denotes the atrial surface manifold, and $[0, 1]^2$ is a normalised two-dimensional coordinate space. The map Φ is constructed using anatomical landmarks, typically the pulmonary vein junctions and the mitral annulus, as anchor points to define physiologically meaningful coordinates. These coordinates correspond to a geodesic distance-based unfolding of the atrial surface, such that anatomical regions (e.g., anterior wall, posterior wall, roof) consistently map to aligned regions of the 2D domain across patients.

The standardised representation allows to map topological entities from different anatomy to corresponding regions. Let $f : \mathcal{M} \rightarrow \mathbb{R}$ be a scalar function defined on the atrial mesh (e.g., LAT values), and let $(\alpha_i, \beta_i) = \Phi(\mathbf{x}_i)$ be the UAC coordinates of a mesh vertex $\mathbf{x}_i$. We then construct a function $\tilde{f} : [0, 1]^2 \rightarrow \mathbb{R}$ via spatial interpolation:

$$\tilde{f}(\alpha, \beta) = \mathcal{I} \left(\{(\alpha_i, \beta_i, f_i)\}_{i=1}^N \right),$$

where $\mathcal{I}(\cdot)$ denotes a suitable interpolation operator.

In practice, we use a hybrid of linear and nearest-neighbour interpolation to preserve continuity in densely sampled regions while maintaining robustness in sparsely sampled areas. The 3D vertex coordinates (x, y, z) for a mesh are projected directly to the 2D grid of shape $(3, 50, 50)$ where the first dimension (3) represents the channels corresponding to the x , y , and z coordinates, and the latter two dimensions (50×50) define the resolution of the structured grid in the UAC system.

Fibre orientations, however, are originally defined per mesh element rather than per vertex. To obtain vertex-wise fibre directions suitable for UAC projection, we convert element-level vectors to nodal vectors using area-weighted averaging. For each vertex, we accumulate the fibre vectors from all incident triangles, scaling each vector by the area of its corresponding triangle so that larger elements contribute proportionally. The resulting vectors are normalised to unit length to preserve directional information. These vertex-level fibre directions are then mapped to the 2D grid using the same UAC interpolation procedure as these coordinates, resulting in a 2D grid of shape $(3, 50, 50)$.

Identify pacing locations

To ensure the generalisability of neural operators and to capture a wide range of conduction dynamics, the EP simulations were initiated using a distributed set of seven clinically relevant pacing sites in the LA: the Left Atrial Appendage (LAA), the Left Superior and Inferior Pulmonary Veins (LSPV, LIPV), the Right Superior and Inferior Pulmonary Veins (RSPV, RIPV), the Coronary Sinus (CS), and the Roof which is defined as the superior wall of the left atrium connecting the LSPV and RSPV [46–48].

Pacing locations are defined across all patient anatomies using the UAC system. For each anatomical region, a 2D bounding box within the UAC domain was specified, delineating a subset of mesh vertices. The physical pacing centre is computed as the mean position of all vertices within that UAC-defined box. The stimulation region is

subsequently defined by all mesh vertices located within a fixed Euclidean radius of 2.0 mm from this central point, ensuring comparable stimulation volumes across all seven sites. The spatial distribution of these seven sites on a representative LA mesh is mapped onto the UAC coordinate system in Fig. 4.




Fig. 4 Marking pacing sites on an LA. **a**, Anterior and posterior views of an LA mesh where the seven pacing locations have been marked. **b**, Spatial distribution of these pacing sites in UAC coordinates. The blue regions indicate the anatomical openings (pulmonary veins and mitral valve).

GPU-based FEM EP simulation

To generate a large-scale, comprehensive dataset for training, we simulated LA electrical propagation using the widely adopted monodomain equation [49] (further implementation details are provided in Supplementary Section 1). This biophysical model captures the electrical conduction through anisotropic cardiac tissue, balancing capacitive and ionic currents with diffusive current flow, governed by the longitudinal (σ_l) and transverse (σ_t) conductivities.

The primary and clinically significant output extracted from these simulations is the Local Activation Time (LAT) map, which defines the time t at which the transmembrane potential V first crosses a predefined threshold (V_{th} , typically 0 mV) at any location x in the cardiac domain:

$$\text{LAT}(x) = \min \{t > 0 \mid V(x, t) > V_{th}\},$$

For each patient, we varied the longitudinal (σ_l) and transverse (σ_t) conductivities across seven clinically relevant pacing locations. At each pacing site, 300 combinations of σ_l and the anisotropy ratio σ_l/σ_t were sampled using Latin hypercube sampling [24]. The chosen range of σ_l (0.1–0.4 S/m) is consistent with experimentally measured conduction velocities and previously reported myocardial conductivity values [50], while anisotropy ratios between 4 and 10 reflect the substantially faster fibre-aligned conduction observed in atrial tissue [51]. These values imply corresponding σ_t in the range 0.01–0.1 S/m.

Each simulation was run for 600 ms to capture the full activation sequence, and the local activation time (LAT) map was recorded during runtime. Each simulation

required approximately one minute on an NVIDIA A100 GPU, enabling the generation of 210,000 simulations for Cohort A and 98,700 simulations for Cohort B, yielding a combined total of 308,700 simulations. Executing these simulations across eight A100 GPUs reduced the total wall-clock time from an estimated 214 days (single GPU) to approximately 27 days.

Neural operator learning

We formulate LAT map prediction as a neural operator learning problem, in which the goal is to approximate a mapping from patient-specific anatomical and electrophysiological parameters to the corresponding spatial distribution of activation times. Formally, we aim to learn an operator

$$\mathcal{G} : (A, \sigma, p) \mapsto \text{LAT}(\mathbf{x}),$$

where A denotes the atrial geometry, σ the conductivity tensor, p the pacing site location, and $\text{LAT}(\mathbf{x})$ the activation time at position $\mathbf{x} \in \Omega$ where Ω is the computational c. This framework requires the generation of large and diverse training datasets that pair input configurations (A, σ, p) with ground-truth LAT maps obtained from high-fidelity EP simulations.

Model architecture

Our proposed architecture integrates a coordinate embedding module for point cloud projection with a Vision Transformer-based encoder and a convolutional decoder, enabling high-fidelity spatial field prediction from both sparse and structured physiological data. The model is designed to process heterogeneous inputs, including sparse point clouds (e.g., pacing coordinates) and spatially distributed scalar fields (e.g., tissue conductivity, surface area), by transforming them into a unified, fixed-resolution tensor representation suitable for transformer-based spatial modelling (Fig. 5).

Point-to-grid projection.

To encode pacing (stimulation) locations provided as sparse, unordered two-dimensional point clouds of shape $(N, 2)$ where N is the number of vertices in the pacing region, we introduce an embedding module that maps the point set into a 2D grid representation (Fig. 5b). This module follows the architectural design of PointNet [52], applying a multi-layer perceptron (MLP) independently to each input point, followed by a permutation-invariant mean pooling operation to obtain a global descriptor of the stimulation pattern. This descriptor is then decoded via a learnable mapping to a 2D spatial grid of shape $(3, 50, 50)$. Unlike PointNet, which is optimised for classification and segmentation tasks, our adaptation is explicitly designed to produce spatially structured embeddings that are compatible with convolutional and transformer-based feature extractors.




Fig. 5 Model architecture. **a**, The construction of the input and output for neural operator learning, where patient-specific anatomical and physiological parameters are projected onto a structured grid via the Universal Atrial Coordinate system. The pacing location is converted from a point cloud representation into a spatial distribution on the same grid, enabling consistent integration with other input modalities. **b**, A neural network module that maps a point cloud of shape $(N, 2)$, representing spatial input features such as pacing sites or anatomical landmarks, into a structured 2D grid of shape $(1, 50, 50)$. Here, N refers to the number of vertices where stimulus (pacing) is applied. **c**, The Vision Transformer block used in (d), which is the original ViT model proposed in [23]. **d**, Our full encoder-decoder architecture, which uses convolutional patch embedding and a Vision Transformer encoder to process the input grid, and then upsamples the latent features using a transposed convolution and projects them into the final LAT map using a convolutional output layer.

Multi-modal feature integration.

The input to the neural operator is a multi-modal, multi-channel grid that encodes patient-specific mesh properties and simulation parameters into a fixed-size ten-channel representation (Fig. 5a), which can be divided into two groups. The first group comprises the seven geometric and anatomical channels, including the UAC-projected 3D vertex coordinates (x, y, z) and the 3D fibre direction vectors (six channels in total, each forming a $(3, 50, 50)$ feature map). This is supplemented by the total surface area of the left atrium, a global anatomical feature that is broadcast across the spatial domain to form a $(1, 50, 50)$ constant grid. The second group comprises the two parametric features including the two conduction velocities $(\sigma_l$ and $\sigma_t)$, which are also broadcast across the spatial domain to form a $(2, 50, 50)$ feature map, and the

encoded pacing location in a $(1, 50, 50)$ feature map. The output of the neural network is trained against the target LAT map of shape $(1, 50, 50)$.

ViT-Based Spatial Representation Learning.

To model spatial dependencies across the input domain, we employ a Vision Transformer architecture that encodes the multi-channel input grid into a sequence of contextualised latent representations. A convolutional patch embedding layer [53] first divides the input grid into non-overlapping patches and projects each patch into a shared embedding space, producing a sequence of latent tokens. Learnable positional embeddings are added to retain spatial structure within the tokenised representation. This sequence is then processed by a stack of standard ViT blocks, each comprising multi-head self-attention and feed-forward sublayers, which enables the integration of global context and the capture of long-range spatial interactions (Fig. 5c).

Following the ViT-based encoder, the latent tokens are reshaped into a compact 2D feature map. To recover full spatial resolution, we employ a lightweight decoder consisting of a transposed convolution layer for upsampling, followed by a convolutional projection head that produces the final output map. The transposed convolution operation, originally introduced in deep learning by [54] for adaptive deconvolutional networks and later popularised for CNN visualisation by [55], enables learnable upsampling by applying the transpose of the convolution operator. This decoding step ensures that the global features encoded by the transformer are effectively translated back into a high-resolution spatial field for pixel-wise prediction.

Sobolev Loss with Spatial Smoothness Regularisations

To ensure our model produces physiologically consistent LAT maps, we employ an objective function grounded in the relative Sobolev H^1 norm. This choice enforces consistency not only in the predicted values but also in their spatial derivatives, a property essential for accurately modelling continuous and smooth conduction patterns.

To further refine the solution and mitigate non-physical, high-frequency artefacts while preserving sharp wavefront transitions, we incorporate two standard regularisation terms. The first is Total Variation (TV) regularisation $\mathcal{L}_{\text{TV}}$, which penalises the L^1 norm of the image gradient. The second is a Laplacian regularisation $\mathcal{L}_{\text{Laplacian}}$, which promotes harmonicity by penalising spatial curvature (full definitions are provided in Supplementary Section 4).

The overall objective function is defined as a weighted sum:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{H^1}(u, v) + \lambda \mathcal{L}_{\text{TV}}(u) + \lambda \mathcal{L}_{\text{Laplacian}}(u),$$

where $\mathcal{L}_{H^1}(u, v)$ is the primary loss between the predicted LAT map (u) and the ground truth (v), and the hyperparameter λ controls the strength of the regularisation. This combination ensures that the model’s predictions are both highly accurate and physiologically coherent.

Declarations

Funding

This work is supported by the NIHR Imperial Biomedical Research Centre (BRC) and by the British Heart Foundation Centre of Research Excellence (RE/24/130023). SAN is supported by NIH R01-HL152256 and R01-HL162260, ERC PREDICT-HF453 (864055), BHF (RG/20/4/34803), EPSRC (EP/P01268X/1, EP/Z531297/1) and by the Technology Missions Fund under the EPSRC Grant EP/X03870X/1 & The Alan Turing Institute. CR receives funding from the British Heart Foundation (RG/20/4/34803). We acknowledge computational resources and support provided by the Imperial College Research Computing Service.

Data availability

Dataset A used in this study is publicly available on Figshare under the permanent identifier [10.6084/m9.figshare.30712559](https://doi.org/10.6084/m9.figshare.30712559). Dataset B contains sensitive information and cannot be publicly released due to ethical and privacy restrictions. Access to Dataset B may be requested from the corresponding author and will be provided subject to institutional approval and a formal data-sharing agreement.

Code availability

All code used to generate the results in this study is available at: https://github.com/sagebei/unified_framework_operator_learning.

References

- [1] Kornej, J., Börschel, C. S., Benjamin, E. J. & Schnabel, R. B. Epidemiology of atrial fibrillation in the 21st century: novel methods and new insights. *Circulation research* **127**, 4–20 (2020).
- [2] Cheng, S. *et al.* Global burden of atrial fibrillation/atrial flutter and its attributable risk factors from 1990 to 2021. *Europace* **26**, euae195 (2024).
- [3] Li, X. *et al.* Global burden and health inequality of atrial fibrillation/atrial flutter from 1990 to 2021. *Frontiers in Cardiovascular Medicine* **12**, 1585980 (2025).
- [4] Lai, Y. *et al.* Causes of death in patients with atrial fibrillation in the uk: a nationwide electronic health record study. *European heart journal open* **5**, oeae103 (2025).
- [5] Heijman, J., Sutanto, H., Crijns, H. J., Nattel, S. & Trayanova, N. A. Computational models of atrial fibrillation: achievements, challenges, and perspectives for improving clinical care. *Cardiovascular Research* **117**, 1682–1699 (2021).
- [6] Boyle, P. M. *et al.* Computationally guided personalized targeted ablation of persistent atrial fibrillation. *Nature biomedical engineering* **3**, 870–879 (2019).

- [7] Roney, C. H. *et al.* Time-averaged wavefront analysis demonstrates preferential pathways of atrial fibrillation, predicting pulmonary vein isolation acute response. *Frontiers in Physiology* **12**, 707189 (2021).
- [8] Krogh-Madsen, T., Abbott, G. W. & Christini, D. J. Effects of electrical and structural remodeling on atrial fibrillation maintenance: a simulation study. *PLoS computational biology* **8**, e1002390 (2012).
- [9] Göktepe, S. & Kuhl, E. Computational modeling of cardiac electrophysiology: a novel finite element approach. *International journal for numerical methods in engineering* **79**, 156–178 (2009).
- [10] Rocha, B. M. *et al.* A macro finite-element formulation for cardiac electrophysiology simulations using hybrid unstructured grids. *IEEE Transactions on Biomedical Engineering* **58**, 1055–1065 (2010).
- [11] Jaffery, O. A., Melki, L., Slabaugh, G., Good, W. W. & Roney, C. H. A review of personalised cardiac computational modelling using electroanatomical mapping data. *Arrhythmia & Electrophysiology Review* **13**, e08 (2024).
- [12] Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics* **378**, 686–707 (2019).
- [13] Kovachki, N. *et al.* Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research* **24**, 1–97 (2023).
- [14] Lu, L., Jin, P., Pang, G., Zhang, Z. & Karniadakis, G. E. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature machine intelligence* **3**, 218–229 (2021).
- [15] Li, Z. *et al.* Fourier neural operator for parametric partial differential equations. *International Conference on Learning Representations* (2021).
- [16] Li, Z., Huang, D. Z., Liu, B. & Anandkumar, A. Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research* **24**, 1–26 (2023).
- [17] Li, Z. *et al.* Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems* **36**, 35836–35854 (2023).
- [18] Ahmad, Z. *et al.* Diffeomorphic latent neural operators for data-efficient learning of solutions to partial differential equations. *arXiv preprint arXiv:2411.18014* (2024).

- [19] Trayanova, N. A., Popescu, D. M. & Shade, J. K. Machine learning in arrhythmia and electrophysiology. *Circulation research* **128**, 544–566 (2021).
- [20] Yin, M. *et al.* A scalable framework for learning the geometry-dependent solution operators of partial differential equations. *Nature computational science* **4**, 928–940 (2024).
- [21] Loeffler, S. E. *et al.* Graph fourier neural kernels (g-funk): Learning solutions of nonlinear diffusive parametric pdes on multiple domains. *arXiv preprint arXiv:2410.04655* (2024).
- [22] Roney, C. H. *et al.* Universal atrial coordinates applied to visualisation, registration and construction of patient specific meshes. *Medical image analysis* **55**, 65–75 (2019).
- [23] Dosovitskiy, A. *et al.* An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations* (2021).
- [24] McKay, M. D., Beckman, R. J. & Conover, W. J. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **42**, 55–61 (2000).
- [25] Healy, J. & McInnes, L. Uniform manifold approximation and projection. *Nature Reviews Methods Primers* **4**, 82 (2024).
- [26] Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**, 600–612 (2004).
- [27] Tripura, T. & Chakraborty, S. Wavelet neural operator for solving parametric partial differential equations in computational mechanics problems. *Computer Methods in Applied Mechanics and Engineering* **404**, 115783 (2023).
- [28] Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. *International Conference on Medical image computing and computer-assisted intervention* 234–241 (2015).
- [29] He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition* 770–778 (2016).
- [30] Brown, A. L., Liu, J., Ennis, D. B. & Marsden, A. L. Cardiac mechanics modeling: recent developments and current challenges. *arXiv preprint arXiv:2509.07971* (2025).

- [31] Hao, Z. *et al.* Gnot: A general neural operator transformer for operator learning. *International Conference on Machine Learning* 12556–12569 (2023).
- [32] Li, Z., Meidani, K. & Farimani, A. B. Transformer for partial differential equations’ operator learning. *Transactions on Machine Learning Research* (2023).
- [33] Shih, B., Peyvan, A., Zhang, Z. & Karniadakis, G. E. Transformers as neural operators for solutions of differential equations with finite regularity. *Computer Methods in Applied Mechanics and Engineering* **434**, 117560 (2025).
- [34] Ovadia, O. *et al.* Vito: Vision transformer-operator. *Computer Methods in Applied Mechanics and Engineering* **428**, 117109 (2024).
- [35] Gaeta, S., Bahnson, T. D. & Henriquez, C. High-resolution measurement of local activation time differences from bipolar electrogram amplitude. *Frontiers in Physiology* **12**, 653645 (2021).
- [36] Narayan, S. M. & John, R. M. Advanced electroanatomic mapping: current and emerging approaches. *Current Treatment Options in Cardiovascular Medicine* **26**, 69–91 (2024).
- [37] Vaswani, A. *et al.* Attention is all you need. *Advances in neural information processing systems* **30** (2017).
- [38] Wang, C. *et al.* Convolutional embedding makes hierarchical vision transformer stronger. *European conference on computer vision* 739–756 (2022).
- [39] Luo, W., Li, Y., Urtasun, R. & Zemel, R. Understanding the effective receptive field in deep convolutional neural networks. *Advances in neural information processing systems* **29** (2016).
- [40] Roney, C. H. *et al.* Predicting atrial fibrillation recurrence by combining population data and virtual cohorts of patient-specific left atrial models. *Circulation: Arrhythmia and Electrophysiology* **15**, e010253 (2022).
- [41] Lee, A. W. *et al.* Is regional atrial strain a useful surrogate of regional atrial fibrosis in atrial cardiomyopathy? *European Heart Journal-Imaging Methods and Practice* **3**, qyaf068 (2025).
- [42] Pashakhanloo, F. *et al.* Myofiber architecture of the human atria as revealed by submillimeter diffusion tensor imaging. *Circulation: arrhythmia and electrophysiology* **9**, e004133 (2016).
- [43] Roney, C. H. *et al.* Constructing a human atrial fibre atlas. *Annals of biomedical engineering* **49**, 233–250 (2021).
- [44] Solís-Lemus, J. A. *et al.* Evaluation of an open-source pipeline to create patient-specific left atrial models: a reproducibility study. *Computers in Biology and*

Medicine **162**, 107009 (2023).

- [45] Razeghi, O. *et al.* Cemrgapp: an interactive medical imaging application with image processing, computer vision, and machine learning toolkits for cardiovascular research. *SoftwareX* **12**, 100570 (2020).
- [46] Romero, J., Natale, A. & Di Biase, L. Atrial fibrillation ablation beyond pulmonary veins: the role of left atrial appendage. *Revista Portuguesa de Cardiologia* **36**, 31–41 (2017).
- [47] Moreira, W., Timmermans, C., Wellens, H. J. & Rodriguez, L.-M. Atrial tachycardia originating from the pulmonary vein: importance of the stimulation sites. *Texas Heart Institute Journal* **35**, 356 (2008).
- [48] Elmariah, S. *et al.* Coronary sinus pacing for the management of right ventricular and atrial infarction with isolated right ventricular pulsus alternans. *Texas Heart Institute Journal* **40**, 497 (2013).
- [49] Geselowitz, D. B. & Miller, W. A bidomain model for anisotropic cardiac muscle. *Annals of biomedical engineering* **11**, 191–206 (1983).
- [50] Fu, Z. *et al.* Progress of conductivity and conduction velocity measured in human and animal hearts. *Reviews in Cardiovascular Medicine* **25**, 364 (2024).
- [51] Kotadia, I. *et al.* Anisotropic cardiac conduction. *Arrhythmia & Electrophysiology Review* **9**, 202 (2020).
- [52] Qi, C. R., Su, H., Mo, K. & Guibas, L. J. Pointnet: Deep learning on point sets for 3d classification and segmentation. *Proceedings of the IEEE conference on computer vision and pattern recognition* 652–660 (2017).
- [53] Wu, H. *et al.* Cvt: Introducing convolutions to vision transformers. *Proceedings of the IEEE/CVF international conference on computer vision* 22–31 (2021).
- [54] Zeiler, M. D., Taylor, G. W. & Fergus, R. Adaptive deconvolutional networks for mid and high level feature learning. *2011 international conference on computer vision* 2018–2025 (2011).
- [55] Zeiler, M. D. & Fergus, R. Visualizing and understanding convolutional networks. *European conference on computer vision* 818–833 (2014).