

JAX-EP: A Fully Differentiable Monodomain Solver for Cardiac Electrophysiology

Ovais Ahmed Jaffery

`o.jaffery@qmul.ac.uk`

Queen Mary's Digital Environment Research Institute (DERI)

Mahmoud Ehresh

Queen Mary's Digital Environment Research Institute (DERI)

H Valli

Royal Brompton & Harefield Hospitals, Guy's and St Thomas' NHS Foundation Trust

Gregory Slabaugh

Queen Mary's Digital Environment Research Institute (DERI)

Edward J. Vigmond

IHU Liryc, The Heart Rhythm Disease Institute, Fondation Bordeaux UniversitéUniversitéUniversité

Shouvik Haldar

Royal Brompton & Harefield Hospitals, Guy's and St Thomas' NHS Foundation Trust

Caroline H. Roney

Queen Mary's Digital Environment Research Institute (DERI)

Article

Keywords:

Posted Date: August 25th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10646136/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

JAX-EP: A Fully Differentiable Monodomain Solver for Cardiac Electrophysiology

Ovais Ahmed Jaffery^{1,*}, Mahmoud Ehresh², H Valli¹, Gregory Slabaugh¹, Edward J. Vigmond⁴, Shouvik Haldar^{2,3}, and Caroline H. Roney¹

¹Queen Mary's Digital Environment Research Institute (DERI), London, UK

²Cardiovascular Imaging Department, Royal Brompton & Harefield Hospitals, Guy's and St Thomas' NHS Foundation Trust, London, UK

³National Heart and Lung Institute, Imperial College London, London, UK

⁴IHU Liryc, The Heart Rhythm Disease Institute, Fondation Bordeaux Université, Pessac, France

*o.jaffery@qmul.ac.uk

Abstract

Cardiac electrophysiology digital twins integrate patient-specific anatomy, tissue properties and clinical measurements into mechanistic models with transformative potential for predicting arrhythmia mechanisms and guiding ablation therapy. Realising this vision demands solvers that are physiologically faithful, GPU-scalable, and fully differentiable, enabling gradient-based inference of electrophysiological parameters directly from clinical recordings. Here we present JAX-EP, the first fully differentiable monodomain cardiac electrophysiology solver, in which exact gradients propagate end-to-end through ionic membrane kinetics, anisotropic tissue conductivity, and fibre orientation via forward mode autodiff. Implemented entirely in JAX as a single JIT-compiled, checkpoint regularised time marching loop, JAX-EP couples the modified Mitchell–Schaeffer ionic model to an anisotropic Crank–Nicolson diffusion operator on triangulated surface meshes, operating identically across one-, two-, and three-dimensional cardiac geometries. Applied to a patient-specific left atrial surface mesh (132,803 nodes), the solver achieved $38\text{--}40\times$ GPU acceleration over CPU execution; autodiff delivered exact gradients for all five ionic and conductivity parameters in five forward passes, verified against finite difference reference to within 0.75%, and the spatial sensitivity of activation times with respect to longitudinal conductivity across all LA nodes was computed $37\times$ faster than serial finite difference evaluation. `jax.vmap` vectorised batch simulation yielded $8,426\text{--}10,433\times$ speedup over CPU sequential execution, enabling parameter search on a 3D manifold using a single GPU. Driven by simulated high density omnipolar electrograms from a two-beat S1–S2 protocol, JAX-EP recovered five parameters across three physiologically distinct regimes with mean absolute parameter error of 0.06–0.10%, reducing activation–recovery interval restitution error from 30.5% at initialisation to 0.5% following two phase optimization. By introducing an automatic differentiation-enabled cardiac electrophysiology solver, JAX-EP opens new avenues for gradient-based parameter learning and functional personalisation of cardiac electrophysiology digital twins.

Introduction

Advances in computational cardiac modelling are driving a shift from mechanistic research tools towards clinically deployable technologies for personalised cardiovascular care. Computational electrophysiology (EP) is at the forefront of this transition, enabling patient-specific EP digital twins (DT) for the diagnosis and treatment of arrhythmias^{1,2}. A faithful representation of cardiac EP relies on the combined personalization of anatomical structure, tissue propagation and cellular function. Beyond anatomical reconstruction, functional digital twinning requires computationally efficient calibration strategies that preserve essential biophysical features such as fibre orientation, tissue anisotropy, and an accurate representation of cellular dynamics³. For this reason, methods for generating these DTs are evolving from a static activation time-based description towards functional calibration targets such as features derived from morphology of the electrogram (EGM), restitution dynamics and the effective refractory period (ERP) of the cardiac tissue^{4,5}. Contemporary pipelines increasingly integrate imaging and EP based multimodal datasets^{6,7}, combining LGE MRI with EAM data to jointly inform structural and functional properties, while relying on reduced order emulators^{8,9} to maintain computational tractability.

The monodomain model, derived from Tung's 1978 bidomain framework¹⁰, couples a reaction–diffusion PDE for the transmembrane potential with a system of ionic ODEs describing membrane kinetics¹¹. As the computational foundation of contemporary EP digital twins, its clinical utility ultimately depends on the accurate recovery of patient specific electrophysiological properties. Despite considerable advances in modeling of cardiac DTs^{3,12,13}, the electrophysiological properties governing atrial excitation, including AP kinetics, ionic time constants, and intracellular conductivity (G_{il}), cannot be reliably recovered from clinical EP data¹⁴. This fundamental limitation remains a major barrier to the clinical translation of cardiac EP




Figure 1. JAX-EP overview. *Top:* The JAX-EP solver pipeline: patient MRI and fibre atlas define the LA geometry; ionic profiles and tissue conductivity parameterise the forward model. Four JAX primitives (`@jax.jit`, `lax.scan`, `sparse.linalg.cg`, GPU execution) implement the monodomain solver as a single Accelerated Linear Algebra (XLA)-compiled kernel, producing transmembrane voltage V_m and HD grid omnipolar EGMs (OEGMs). The entire pipeline is end-to-end differentiable via forward mode autodiff (curved arrow). *Bottom-left (1) Differentiable inversion:* `jax.jacfwd` propagates five tangent vectors alongside the primal simulation ($N_T = 11,000$ timesteps), delivering exact gradients $\partial \mathcal{L} / \partial \theta$ in five forward passes at constant memory $\mathcal{O}(M \times N_p)$. *Bottom-centre (2) Parameter learning:* a two phase pipeline (gradient-free Nelder–Mead warm start followed by gradient-based L-BFGS-B refinement) recovers ionic and conductivity parameters θ^* from nine clique OEGMs with sub-percent accuracy (MSE = 1.0×10^{-9}). *Bottom-right (3) GPU acceleration and batch simulation:* `jax.vmap` evaluates B parameter sets as a single fused GPU kernel, yielding $10,433 \times$ batch speedup over CPU sequential execution; forward simulation achieves $38\text{--}40 \times$ GPU acceleration, enabling the full pipeline on a single commodity GPU in minutes. *Bottom right panel:* Functional validation of the recovered parameters θ^* : the top panel shows near perfect overlay of recovered (orange) and ground-truth (black) OEGM waveform morphology at the centre clique; the bottom panel shows the activation–recovery interval (ARI) restitution curve, confirming that recovered parameters faithfully reproduce the dynamic electrophysiological behaviour of the tissue across the full range of diastolic intervals (100–250 ms), with individual point errors below 1%.

digital twins. Without reliable estimation of patient specific electrophysiological parameters, mechanistic models cannot fully characterise the arrhythmogenic substrate, investigate drivers of atrial fibrillation (AF), support personalised catheter ablation planning, or ultimately predict long term procedural outcomes. In current practice, tissue conductivity is typically estimated indirectly by calibrating simulated and measured conduction velocities, whereas cellular electrophysiological properties are commonly assigned population averaged values representing healthy or fibrotic tissue states^{15–17}. Consequently, current cardiac EP DTs remain only partially personalised¹², as tissue conductivity and cellular electrophysiological parameters are not jointly recovered from patient specific clinical electrophysiology data. Existing inversion approaches address individual aspects of this challenge but have yet to achieve joint recovery of tissue and cellular electrophysiological properties within a

unified mechanistic framework. For instance, recent ECG calibrated atrial DT frameworks have advanced automated patient specific model generation, yet calibration of the full EP parameter space remains computationally infeasible and insufficiently constrained by current clinical data¹⁸. Eikonal and PIEMAP-based approaches efficiently estimate conduction patterns from EAM data but neglect ionic dynamics¹⁹. LGE-MRI-informed atrial DTs generally rely on fixed ionic parameter sets^{17,20}. Physics-informed neural networks embed EP constraints into neural architectures but often sacrifice numerical fidelity²¹, while Gaussian-process and ensemble-based methods provide efficient calibration but commonly assume isotropic tissue representations^{9,22}. Although reaction eikonal frameworks have enabled large scale personalisation studies^{23,24}, direct recovery of tissue and cellular parameters from clinical EP data within a unified mechanistic framework remains elusive. Consequently, direct inversion of the monodomain equations at clinical mesh resolution has not yet been demonstrated²⁵.

The obstacle is computational: inverting the monodomain requires gradients of simulation outputs with respect to model parameters, which for a long trajectory non linear PDE demands either a manually derived adjoint or prohibitively expensive finite difference (FD) approximation^{26,27}. Differentiable programming eliminates this barrier through automatic differentiation (autodiff)²⁸. JAX is an open source numerical computing framework that combines high performance execution with composable automatic differentiation, enabling exact gradients to be obtained directly from the forward solver²⁹. JAX-Fluids³⁰, JAX-FEM³¹, and TORAX³² each demonstrated that expressing a forward solver as a composition of JAX primitives within `jax.lax.scan` yields end-to-end differentiability through thousands of timesteps directly from the forward solver^{28,29}. Cardiac EP presents a harder challenge than any of these: the ionic gating variable introduces steep Jacobian gradients over long trajectories, the surface FEM on an unstructured clinical mesh requires anisotropic cotangent-weight stiffness operators, and the observable is a nonlinear far field projection of the voltage field. To the best of our knowledge, no existing cardiac EP solver exposes these gradients.

We present JAX-EP: the first fully differentiable monodomain EP solver on a patient-specific left atrial (LA) surface mesh (Fig. 1). By expressing the time-marching loop, cotangent-weight FEM operators, and smooth modified Mitchell–Schaeffer (mMS) ionic model as JAX primitives within `jax.lax.scan`, the entire LA manifold is simulated end-to-end differentially with respect to ionic and conductivity parameters without adjoint derivation or any additional code beyond the forward solver. `jax.jacfwd` delivers exact forward mode gradients for all five mMS parameters in five passes, verified against an independent FD reference to within 0.75%, and drives a parameter estimation pipeline that recovers ionic and conductivity parameters from HD grid omnipolar EGMs with sub-percent accuracy. `jax.vmap` delivers batch simulation speedups of 8,426–10,433× over CPU sequential execution, enabling parameter search on a single GPU without HPC infrastructure. The spatial Jacobian $\partial AT / \partial G_{ij}$ across all 132,803 LA nodes is computed in two forward passes (37× GPU acceleration), and forward mode autodiff is identified as the structurally correct differentiation mode for cardiac EP. Fig. 1 presents the solver architecture and its three principal contributions: parameter learning, differentiable gradients, and GPU acceleration, demonstrated on a patient-specific LA geometry. Broadly, this work establishes differentiable cardiac electrophysiology as a practical framework for translating mechanistic digital twins from computational models into clinically actionable decision support tools.

Results

JAX-EP is a fully differentiable monodomain EP solver in which differentiability arises from the implementation architecture rather than from additional mathematical derivation. As a consequence, JAX-EP exposes exact gradients through the coupled ionic and tissue dynamics, enabling joint recovery of cellular and tissue electrophysiology and contributing towards the longer-term goal of patient-specific atrial EP digital twins in clinical practice. By expressing the time-marching loop as a `jax.lax.scan` primitive and all spatial operators as differentiable JAX functions, JAX-EP simulates the entire LA manifold end-to-end while remaining differentiable with respect to the biophysical parameter vector θ via forward mode autodiff; `jax.jacfwd` propagates tangent vectors alongside the primal state, delivering exact gradients for all parameters in forward passes without a backward pass. Results are presented in six categories: patient-specific clinical pacing LA forward simulation; end-to-end differentiability and gradient verification; JAX-specific computational capabilities; ionic and conductivity parameter learning; robustness to clinical-type EGM distortion, a key requirement for translation to intracardiac recordings acquired during routine EP studies; and solver convergence.

LA forward simulation



A patient-specific endocardial surface mesh was reconstructed from clinical MRI and EP study data acquired as part of the HEAT-AF trial¹ at Royal Brompton & Harefield Hospitals, Guy’s & St Thomas’ NHS Foundation Trust, London, UK. Virtual catheter geometry and dimensions were derived from the clinical HD Grid Mapping Catheter (Advisor™ HD Grid, Abbott), and pacing sites followed the clinical protocol used during patient data acquisition. The JAX-EP monodomain solver, coupled with the mMS ionic model, was applied to simulate clinically representative pacing from the coronary sinus (CS) and left atrial appendage (LAA) on the patient-specific LA surface mesh. Spatially varying fibre orientations were assigned from a human

CS Pacing | AT = [13, 217] ms | CV = 0.384 m/s | Left Lateral View

LAA Pacing | AT = [13, 262] ms | CV = 0.382 m/s | Left Lateral View

A

B



★ Pacing site ● HD Grid Site 1 (white) ● HD Grid Site 2 (cyan) Left Lateral View

C Nine-Clique Omnipolar EGMs — HD Grid Site 1 (CS pacing: blue | LAA pacing: orange dashed)




Figure 2. LA forward simulation: (A–B) LAT maps for CS (panel A - left) and LAA (panel B - right) pacing. Stars mark pacing sites; white circles and yellow diamonds mark HD grid recording sites 1 and 2 respectively. (C) Nine-clique omnipolar EGMs computed via spatial gradient method from the 16 electrode HD grid (site 1). CS pacing (blue) and LAA pacing (orange dashed) overlaid showing physiologic morphology of simulated EGMs. Centre clique C5 highlighted (green background).

Table 1. Forward simulation parameters and results.

Parameter	CS pacing	LAA pacing
Mesh nodes / elements	132,803 / 264,821	132,803 / 264,821
Pacing mask (5 mm radius)	729 nodes	880 nodes
Protocol	$2 \times S1$, BCL = 600 ms, $N_T = 24,200$, $\Delta t = 0.05$ ms	
Ionic model (set III)	mMS: $\tau_{in} = 0.300$, $\tau_{out} = 5.000$, $\tau_{open} = 120.0$, $\tau_{close} = 150.0$	
Conductivity	$G_{il} = 0.350$ mS cm ⁻¹ , $G_{il} : G_{it} = 4 : 1$	
HD grid (per site)	16 electrodes, 2×8 arrangement, 3 mm electrode spacing, 9 omnipolar cliques	
AT range (ms)	[13, 223]	[13, 263]
Total AT duration (ms)	210	250
CV (m s ⁻¹)	0.384	0.382
Activated nodes	132,803 (100%)	132,803 (100%)
Wall time (CPU)	2,608.2 s	2,738.6 s
Wall time (GPU, T4)	69.1 s	68.8 s
GPU speedup	38×	40×

atrial fibre atlas³³ to achieve physiological spread of activation. Simulation parameters and mesh details are summarised in Table 1 for LA forward simulation shown in (Fig. 2).

CS pacing yielded a total activation sequence of 210 ms and conduction velocity $CV \approx 0.384$ m s⁻¹; LAA pacing yielded 250 ms and $CV \approx 0.382$ m s⁻¹ (Fig. 2A,B). The 40 ms increase in total activation time during LAA pacing reflects the longer conduction pathway from the LAA to the distal CS. Local activation times (LATs) were computed as the first crossing of the membrane voltage threshold $V_{gate} = 0.13$ following the stimulus, yielding physiologically consistent values of ~ 13 ms at the pacing site, in agreement with the 10 ms stimulus onset. GPU execution reduced forward simulation time by approximately 40-fold relative to a single CPU core, producing identical activation maps.

Two simulated 16-electrode HD grid catheters placed at recording sites on the LA posterior wall yielded 32 unipolar EGMs per pacing configuration, from which omnipolar EGMs (OEGMs) were derived at nine 2×2 electrode cliques via the spatial gradient method³⁴ (Fig. 2C). The OEGMs exhibited physiologically consistent morphology under both pacing configurations, with sharp activation deflections resolved at all nine clique positions, demonstrating that JAX-EP reproduces a physiologically plausible signal chain from transmembrane voltage to catheter-recorded waveform. The centre clique (C5) served as the primary reference OEGM in subsequent analyses.

End-to-end differentiability and gradient verification

FD gradients of the objective loss function $\mathcal{L}(\theta)$ defined as the mean squared error between the simulated and target EGMs were computed at two high definition grid sites under CS and LAA pacing, yielding 20 independent evaluations across all five parameters $\theta = [\tau_{in}, \tau_{out}, \tau_{open}, \tau_{close}, G_{il}]$ ($\epsilon = 10^{-3}$,^{30,31}). All 20 were finite, non-zero, and physically consistent: loss rose with τ_{in} , fell with G_{il} , and sign reversals between sites reflect the geometry dependent sensitivity of each HD grid location to the propagating wavefront (Fig. 3). GPU reduced FD computation $6.3 \times$ relative to CPU (18.9 s versus 67.9–119.7 s). `jax.grad` executed on the 2,894-node HD-grid foot print patch in 22.5 s on GPU, confirming end-to-end differentiability by construction. Backpropagating through $N_T = 11,100$ timesteps, however, caused gradient explosion: Jacobian eigenvalues compound multiplicatively in reverse, turning the time axis into a hard memory wall of $\mathcal{O}(N_T \times N_p)$ peak storage, or 52 GB on the full LA mesh.

`jax.jacfwd` resolved this by propagating $M = 5$ tangent vectors *forward* chronologically alongside the primal state, keeping memory at $\mathcal{O}(M \times N_p) \approx 6.6 \times 10^5$ values which is independent of N_T entirely and well within 15 GB VRAM. FD gradients served as an independent reference: ratios $\partial \mathcal{L} / \partial \theta_{jacfwd} / \partial \mathcal{L} / \partial \theta_{FD}$ fell within 0.75% of unity for all parameters (Fig. 3B), confirming exact agreement at machine precision. Gradient computation completed in 11.7 s on GPU versus 19.9 s for FD, delivering exact gradients in half the forward passes. In cardiac EP, N_T mostly comprise large number of time steps: resolving the 300–400 ms action potential at $\Delta t \approx 0.05$ ms demands $N_T \sim 10^4$ on any geometry, a hard constraint of ion channel physics. For reverse-mode autodiff, this constitutes an inherent computational bottleneck, whereas it is largely irrelevant for forward-mode autodiff (Fig. 10). Therefore, `jax.jacfwd` provides the structurally appropriate differentiation strategy for cardiac EP, and forward-mode autodiff gradients were adopted throughout the Phase 2 parameter estimation framework presented in the subsequent sections. The spatial Jacobian $\partial AT / \partial G_{il}$ across all 132,803 LA nodes required only two forward passes, completing in 138.6 s on GPU versus 5,096.2 s on CPU ($37 \times$). A non-differentiable solver would need 132,803 perturbation runs; which are computationally prohibitive at clinical mesh resolution. Jacobian values were uniformly negative (CS: $[-325, 0]$ ms; LAA: $[-375, 0]$ ms per unit G_{il}): conductivity increase uniformly accelerates activation, with




Figure 3. Gradient verification and JAX-EP computational capabilities. **(A)** Log-magnitude of FD gradients $|\partial \mathcal{L} / \partial \theta_i|$ for all five parameters across four configurations (CS and LAA pacing, HD grid Site 1 and Site 2). All 20 gradients are finite and non-zero, confirming physically consistent parameter sensitivities. **(B)** autodiff (autodiff) verification: `jax.jacfwd` exact gradients compared against independent FD reference for all five parameters. Ratios close to unity confirm exact autodiff gradients in $M = 5$ passes versus $2M = 10$ for FD. Agreement was further confirmed across a broader sample of 50 randomly drawn parameter sets (250 paired gradient comparisons) using a Bland–Altman analysis, which showed negligible bias (-3.91×10^{-4}) and a mean absolute relative error of 0.23% ($R^2 = 0.994$). **(C)** Gradient computation time per configuration: FD (10 passes, CPU and GPU) versus `jax.jacfwd` (5 passes, GPU). GPU reduces FD time 6.3 $\times$; `jax.jacfwd` further halves the pass count while eliminating approximation error. **(D)** `jax.vmap` batch simulation: simultaneous evaluation of the Nelder–Mead simplex yields 10,433 $\times$ (CS) and 8,426 $\times$ (LAA) speedups over CPU sequential execution and $\sim 276\times$ over GPU sequential execution, at zero additional implementation cost.

greater sensitivity under LAA pacing reflecting the longer propagation path. To assess agreement beyond these representative configurations, we performed a Bland–Altman analysis comparing `jax.jacfwd` and FD gradients across 50 randomly sampled parameter sets (250 paired comparisons), which showed negligible bias (-3.91×10^{-4}), narrow 95% limits of agreement (-1.78×10^{-2} to $+1.71 \times 10^{-2}$), and a mean absolute relative error of 0.23% ($R^2 = 0.994$), with a small number of outliers (max 21.6%) concentrated in τ_{in} , τ_{out} , and τ_{close} .

JAX implementation: compilation, scan, and batch simulation

Four JAX primitives underpin the JAX-EP solver. `jax.jit` traces the forward solver once and caches the compiled XLA kernel, eliminating Python overhead on subsequent calls; wall times of 69.1 s (CS) and 68.8 s (LAA) on a Tesla T4 GPU represent 38–40× acceleration over a single CPU core for $N_T = 24,200$ timesteps, with GPU and CPU activation maps identical to within 0.05 ms (Pearson correlation 1.000). `jax.lax.scan` expresses the full timestep loop as a single JAX primitive, enabling compiler-level optimisation of the entire trajectory and providing the loop structure through which `jax.jacfwd` propagates forward mode tangent vectors. `jax.checkpoint` reduces peak activation memory from $\mathcal{O}(N_T \times N_p)$ to $\mathcal{O}(\sqrt{N_T} \times N_p)$ by rematerialising intermediate states on demand during the backward pass³⁵, though `jax.jacfwd` renders this unnecessary for the forward mode path by eliminating the backward pass entirely. `jax.vmap` vectorises the forward solver over parameter batches at no additional implementation cost, yielding speedups of 10,433× (CS) and 8,426× (LAA) over CPU sequential execution (Fig. 3D). Together these four primitives, `jit`, `scan`, `jacfwd`, and `vmap`, deliver a differentiable, memory efficient, and batch capable EP solver from a single forward implementation, without explicit GPU programming or adjoint derivation.

Parameter learning from EGM waveforms

Ionic and tissue parameter learning was performed on the HD grid LA boundary patch (2,894 nodes) for three parameter sets spanning distinct regions of a bounded physiological parameter space^{5,22}. An in silico HD grid recording protocol comprising a simulated 16-electrode HD grid catheter on the posterior atrial wall under CS and LAA pacing was employed (Fig. 2 A, B). OEGMs were computed at all nine 2×2 electrode cliques via the spatial gradient method³⁶. The optimisation target comprised simulated OEGMs recorded during a two-beat S1 S2 protocol: a conditioning stimulus (S1) at $t = 10$ ms followed by a test stimulus (S2) at a coupling interval of 500 ms. This protocol was chosen to minimise forward simulation cost while retaining sufficient waveform information for joint recovery of ionic membrane kinetics (τ_{in} , τ_{out} , τ_{open} , τ_{close}) and longitudinal conductivity (G_{ij}).

A two phase learning strategy was applied. Phase 1 employed Nelder–Mead search on a normalised nine clique S2 EGM shape loss to identify a physiologically consistent parameter region. Phase 2 refined the solution via `jax.jacfwd` L-BFGS-B minimising a joint S1+S2 waveform MSE across all nine cliques, with S1 weighted at 0.5 relative to S2. Parameters were initialised at the midpoint of physiological bounds³⁷, corresponding to a mean initialisation MAPE of 15.7–20.6% across the three sets. Ground truth EGMs were generated independently and withheld from the optimiser, ensuring a strictly GT free learning pipeline with no data leakage between the ground truth generator and the optimiser, consistent with differentiable physics based approaches.^{30,31}

Learned parameter accuracy yielded mean MAPE of 0.09%, 0.10%, and 0.06% for Sets 1-3 respectively, with all five parameters recovered to sub percent accuracy across all three sets (Fig. 5). Learned parameter sets were validated by independent forward simulation, confirming physiological equivalence via ARI restitution, CV, and nine clique EGM waveform overlays (Fig. 4). Wall times on a single CPU core were 43.3, 143.4, and 75.1 min for Sets 1, 2, and 3 respectively; GPU acceleration (Tesla T4) reduced these to 7.1, 17.0, and 4.1 min (6.5–7.9× speedup).

Robustness to clinical-type EGM distortion

To assess robustness of the parameter estimation pipeline to clinically realistic signal corruption, EGM targets were distorted at three severity levels combining additive Gaussian noise and far-field contamination. Far-field bleed was simulated by adding a weighted contribution of the mean signal from all other cliques to each clique, mimicking cross-talk from distant cardiac sources. Gaussian noise was added per clique with power normalised to a specified SNR. Three distortion levels were evaluated:

- **L1 (low):** SNR = 40 dB, far-field $\alpha = 0.08$, mean signal distortion $D \approx 4\%$
- **L2 (medium):** SNR = 30 dB, far-field $\alpha = 0.25$, mean signal distortion $D \approx 14\%$
- **L3 (high):** SNR = 20 dB, far-field $\alpha = 0.60$, mean signal distortion $D \approx 35\%$

where $\mathbf{e}$ denotes the clique EGM time series and signal distortion is defined as $D = \|\mathbf{e}_{\text{distorted}} - \mathbf{e}_{\text{clean}}\|_2 / \|\mathbf{e}_{\text{clean}}\|_2 \times 100\%$, averaged across all nine cliques. No preprocessing was applied to the distorted targets prior to parameter estimation, and the same two-phase optimisation pipeline was applied without modification across all distortion levels. Figs. 11 and 12 (Appendix B) illustrate the distorted EGMs at each distortion level and the effect on parameter recovery. Table 2 summarises the recovered parameters and MAPE.




Figure 4. JAX-EP ionic and tissue parameter learning on the LA patch geometry (2,894 nodes). **(A)** ARI restitution curves computed from independent forward simulations at ground-truth (black), initialisation (blue dashed), and learned (orange) parameters across four S2 coupling intervals (BCL = 350–500 ms), demonstrating convergence from an initialisation mean absolute percentage error (MAPE) of 30.5% to a learned MAPE of 0.5%. **(B)** Centre-clique (C5) omnipolar EGM at BCL = 500 ms for ground truth, initialization, and learned parameters; the learned waveform overlaps the ground truth with MSE reduced from the initialization value. **(C)** Wall time comparison between CPU (single core, grey) and GPU (Tesla T4, green) for Phase 1 (Nelder–Mead) and Phase 2 (JAX FD L-BFGS-B), and total, demonstrating 6.5 $\times$ acceleration for this parameter set. **(D)** Normalised parameter learning chart: bars show the ratio of initialisation (blue) and learned (orange) parameters to ground truth (black dashed line at 1.0), with values annotated above each bar. Near exact recovery is achieved for all five parameters (τ_{in} , τ_{out} , τ_{open} , τ_{close} , G_{il}), with a mean parameter MAPE of 0.09%.

At L1 ($D \approx 4\%$), all three parameter sets were recovered with MAPE below 0.50%, representing a less than 6 fold degradation from the clean baseline despite combined noise and far-field contamination. At L2 ($D \approx 14\%$), MAPE remained below 1.40% across all sets, with τ_{in} , τ_{open} and G_{il} well preserved and the primary degradation confined to τ_{out} and τ_{close} , which govern repolarisation dynamics and are more sensitive to morphological distortion of the EGM recovery waveform. At L3 ($D \approx 35\%$), MAPE increased to 3.35% for Sets 1 and 2, whilst Set 3 remained below 1.20%, reflecting varying sensitivity of different parameter combinations to high-level distortion. The consistent preservation of τ_{in} , τ_{open} and G_{il} across all distortion levels confirms that activation timing and upstroke kinetics, the parameters most directly observable from EGM morphology, are robustly recovered even at clinically challenging SNR levels. The progressive degradation of τ_{out} and τ_{close} with increasing distortion is consistent with the identifiability analysis: these parameters govern slow repolarisation dynamics that are disproportionately affected by noise and far-field contamination in the EGM recovery window. Overall, these results establish that the parameter estimation pipeline maintains accuracy (MAPE < 1.5%) up to signal distortion levels of approximately 14%, and degrades gradually beyond this threshold rather than failing catastrophically.



Figure 5. Ionic and tissue parameter learning results for three parameter sets spanning distinct physiological regimes. **(A)** Normalised parameter learning chart showing the ratio of initial (faded bars) and learned (solid bars) parameter values to ground truth for Set 1 (blue), Set 2 (red), and Set 3 (green). The dashed horizontal line at unity denotes exact recovery; bars close to this line indicate near-perfect parameter learning. The five parameters are the mMS ionic time constants (τ_{in} , τ_{out} , τ_{open} , τ_{close}) and the longitudinal tissue conductivity (G_{II}). **(B)** Summary table reporting, for each parameter set and each row type (Initial, Learned, Ground truth), the mean ARI restitution (ms) and mean CV restitution ($m s^{-1}$) averaged across four S2 coupling intervals (BCL = 350–500 ms), the Phase 1 (Nelder–Mead) and Phase 2 (JAX FD L-BFGS-B) final loss values, and the total wall time on CPU (single core) and GPU (Tesla T4). Parameters were initialised at the midpoint of the physiological bounds³⁷; mean initialisation MAPE ranged from 16 to 21% across the three sets, and was reduced to 0.06 to 0.10% following two phase learning.

Solver convergence and spatial accuracy

Spatial and temporal convergence of the JAX-EP surface FE solver was assessed on a 20×7 mm flat plate geometry, discretised using an unstructured Delaunay triangulation at three spatial resolutions ($\Delta x \approx 0.5, 0.3, 0.2$ mm; 615–3,636 nodes) and three temporal resolutions ($\Delta t = 0.10, 0.05, 0.025$ ms). The solver employed the identical cotangent-weight FE assembly, Crank–Nicolson diffusion operator, and mMS ionic model³⁷ as the LA forward simulation and parameter learning pipeline, ensuring that the convergence study directly characterises the deployed solver. A corner stimulus of $200 \mu A cm^{-2}$ was applied for 2 ms within a 1.5 mm radius, with longitudinal conductivity $G_{II} = 0.350 mS cm^{-1}$ and transverse-to-longitudinal anisotropy ratio of 4:1. Complete activation was achieved at all resolutions. LAT maps and isochrone contours are shown in Fig. 6A,C, confirming the expected elliptical wavefront driven by anisotropic conductivity.

Spatial convergence was quantified using the L_2 norm of the local activation time (LAT) error relative to the ($\Delta x = 0.2$ mm) reference solution, yielding errors of 0.925 ms and 0.131 ms for $\Delta x = 0.5$ mm and $\Delta x = 0.3$ mm, respectively, while the conduction velocity converged to $0.471 m s^{-1}$ (Fig. 6B). Temporal refinement at fixed $\Delta x = 0.2$ mm produced L_2 errors of 0.087 ms and 0.032 ms for $\Delta t = 0.10$ ms and $\Delta t = 0.05$ ms, respectively, with the CV remaining essentially unchanged (0.470 – $0.471 m s^{-1}$) across all time steps (Fig. 6D). The monotonic reduction in both spatial and temporal errors with mesh and timestep refinement demonstrates the expected convergence behaviour of the operator-split Crank–Nicolson formulation on unstructured meshes³⁸.

Discussion

JAX-EP demonstrates that a fully differentiable monodomain cardiac EP solver can be constructed from standard JAX primitives without adjoint derivation, specialist automatic differentiation libraries, or modifications to the forward solver beyond

Table 2. Parameter recovery under combined Gaussian noise and far-field contamination at three distortion levels (L1–L3) for three mMS parameter sets. GT parameters: Set 1 [0.340, 5.580, 192.0, 120.0, 0.200]; Set 2 [0.260, 3.460, 148.0, 176.0, 0.200]; Set 3 [0.300, 5.000, 120.0, 150.0, 0.200]. All time constants in ms; G_{il} in $S\ m^{-1}$.

Condition	Set	τ_{in}	τ_{out}	τ_{open}	τ_{close}	G_{il}	MAPE
<i>Clean</i>							
	Set 1	0.340	5.565	191.9	120.2	0.200	0.09%
	Set 2	0.260	3.470	148.0	175.7	0.200	0.10%
	Set 3	0.300	5.008	120.0	149.9	0.200	0.06%
<i>L1 — SNR = 40 dB, $\alpha = 0.08$, $D \approx 4\%$</i>							
	Set 1	0.341	5.649	192.0	119.3	0.200	0.44%
	Set 2	0.260	3.509	148.1	174.6	0.200	0.49%
	Set 3	0.300	5.045	120.0	149.4	0.200	0.31%
<i>L2 — SNR = 30 dB, $\alpha = 0.25$, $D \approx 14\%$</i>							
	Set 1	0.342	5.799	192.0	117.9	0.201	1.37%
	Set 2	0.261	3.422	147.5	178.0	0.201	0.69%
	Set 3	0.301	5.148	119.9	148.0	0.201	1.03%
<i>L3 — SNR = 20 dB, $\alpha = 0.60$, $D \approx 35\%$</i>							
	Set 1	0.349	6.042	191.4	116.6	0.206	3.35%
	Set 2	0.264	3.786	148.3	168.2	0.203	3.35%
	Set 3	0.305	5.079	119.1	150.6	0.203	1.18%

Table 3. Spatial and temporal convergence of the JAX-EP surface FEM solver on the 20×7 mm flat plate geometry. Reference solutions: $\Delta x = 0.2$ mm (spatial), $\Delta t = 0.025$ ms (temporal).

	Nodes	L2 AT error (ms)	CV ($m\ s^{-1}$)	Wall time (s)
<i>Spatial convergence ($\Delta t = 0.05$ ms fixed)</i>				
$\Delta x = 0.5$ mm	615	0.925	0.450	1.3
$\Delta x = 0.3$ mm	1,632	0.131	0.463	2.3
$\Delta x = 0.2$ mm	3,636	ref	0.471	5.0
<i>Temporal convergence ($\Delta x = 0.2$ mm fixed)</i>				
$\Delta t = 0.100$ ms	3,636	0.087	0.470	3.0
$\Delta t = 0.050$ ms	3,636	0.032	0.471	6.4
$\Delta t = 0.025$ ms	3,636	ref	0.471	13.5

replacing the hard gate switch with a smooth sigmoid approximation. Reverse-mode autodiff via `jax.grad` generates the vector-Jacobian products (VJPs) required for gradient computation automatically from the forward implementation, eliminating the need for a hand-coded adjoint PDE solver. The key architectural insight is that expressing the entire N_T -step simulation as a single `jax.lax.scan` primitive makes the trajectory differentiable by construction, with gradient memory managed by `jax.checkpoint` and batch simulation enabled by `jax.vmap` at no additional implementation cost. This design pattern, demonstrated previously for fluid mechanics and plasma physics^{30–32}, here extends to cardiac EP on clinical resolution unstructured surface meshes for the first time. By making gradient-based monodomain inversion computationally practical, JAX-EP provides a foundation for future clinical calibration of patient-specific atrial EP digital twins from intracardiac EGM recordings.

Differentiability and gradient stability. Reverse mode autodiff (`jax.grad`) produces gradient explosion through long scan chains, a fundamental consequence of compounding Jacobian eigenvalues across $N_T \gg 1$ timesteps. Forward mode autodiff (`jax.jacfwd`) resolves this by propagating tangent vectors alongside the primal simulation, yielding exact gradients for all five parameters with ratios within 0.75% of the FD reference and halving the number of forward passes required. This identifies forward mode autodiff as the structurally correct differentiation mode for cardiac EP, where $N_T \gg M$ is an inescapable physical constraint. The spatial Jacobian $\partial AT / \partial G_{il}$ remains accessible via two forward passes ($37\times$ GPU acceleration), demonstrating practical utility of the differentiable architecture independent of the autodiff mode. Consequently, stable gradient propagation is not merely a numerical requirement but a prerequisite for robust recovery of patient-specific electrophysiology, enabling faithful

characterisation of the arrhythmogenic substrate to support patient stratification and possibly pre-procedure ablation planning. Future work may exploit the differentiable architecture with physiologically detailed atrial and ventricular ionic models³⁹ by combining Rush–Larsen or exponential integrators for gate stability with `jax.jacfwd` for exact gradient propagation, enabling end-to-end differentiable personalisation of whole heart DTs directly from clinical EGM recordings.

Parameter learning from EGM morphology. Recovery of five ionic and conductivity parameters from simulated omnipolar EGMs achieved mean MAPE of 0.06–0.10% across three physiologically distinct parameter sets, reducing ARI restitution error from 30.5% at initialisation to 0.5% after convergence. This demonstrates that a two-beat S1–S2 OEGM protocol retains sufficient waveform information to jointly constrain ionic kinetics and tissue conductivity without restitution curves or multi-beat protocols, consistent with the clinical motivation for catheter guided parameter estimation⁴⁰. Existing approaches to atrial EP personalisation recover CV fields from activation time maps¹⁹, fit ionic parameters to restitution curves⁵, or calibrate reduced models to scalar EP features^{9,22}. These methods do not invert the forward monodomain model directly and cannot simultaneously recover ionic kinetics and conductivity from the same waveform. JAX-EP recovers both jointly by fitting OEGM morphology across all nine HD grid cliques (Fig. 7), with forward physics enforced exactly through the differentiable monodomain solver, and without dependence on LGE-MRI substrate maps⁶.

Computational performance. The $38\text{--}40\times$ GPU speedup over single-core CPU execution for the LA forward simulation is consistent with GPU acceleration reported for cardiac EP solvers employing explicit CUDA implementations^{41,42}, noting that ionic model complexity (two state variables in the mMS model versus up to 43 in physiologically detailed formulations) and mesh size will influence absolute speedup figures across implementations. The $6.5\text{--}7.9\times$ speedup for parameter learning reflects




Figure 6. Convergence study of the JAX-EP FE solver on a 20×7 mm flat plate geometry (Delaunay triangulation, mMS ionic model, $G_{ij} = 0.350\text{ mS cm}^{-1}$). (A) LAT map at the spatial reference resolution ($\Delta x = 0.2$ mm, 3,636 nodes), with white isochrone contours at 5 ms intervals. The elliptical wavefront reflects 4:1 anisotropic conductivity. S denotes the corner stimulus site. (B) Log-log spatial convergence of L2 activation time error with $\mathcal{O}(h)$ and $\mathcal{O}(h^2)$ reference lines; CV converges to 0.554 m s^{-1} . (C) LAT map at the temporal reference resolution ($\Delta t = 0.025$ ms), confirming negligible change from Panel A. (D) Log-log temporal convergence of L2 activation time error; L2 error falls below 0.05 ms at $\Delta t = 0.05$ ms.




Figure 7. Ground truth versus recovered omnipolar EGMs across all nine HD grid cliques (C1–C9) for three physiologically distinct parameter sets (Set 1, Set 2, Set 3). Black: ground truth EGMs simulated with GT ionic parameters. Red dashed: EGMs re-simulated using the recovered parameters from the two phase `jax.jacfwd` L-BFGS-B pipeline. Both S1 and S2 beats are shown. Waveform agreement across all cliques and all three sets confirms sub-percent parameter recovery (mean MAPE 0.06–0.10%) and demonstrates that the differentiable monodomain solver faithfully reproduces the target EGM morphology from the recovered biophysical parameters.

the hybrid CPU/GPU workload of the Nelder–Mead Phase 1 (sequential function evaluations) and FD L-BFGS-B Phase 2 (10 forward passes per gradient step). The `jax.vmap` batch speedup of 8,426–10,433 $\times$ over CPU sequential execution demonstrates that ensemble simulation is computationally tractable on a single GPU, with direct implications for uncertainty quantification, sensitivity analysis, and population cohort studies. Total parameter learning wall times of 4.1–17.0 min on a Tesla T4 for the LA patch represent a promising step towards intra-procedural application, though extension to the complete LA mesh and multiple HD grid sites would increase this cost proportionally. Ensemble simulations on GPU via `jax.vmap` can also significantly reduce the computational burden of AF inducibility and AF burden computational studies, where large parameter sweeps across substrate conditions are required to assess arrhythmia vulnerability^{6,43}.

Solver convergence and accuracy. Spatial and temporal convergence of the surface FEM solver was confirmed on an unstructured Delaunay triangulation of a 20×7 mm flat plate using the identical solver as the full LA simulation. L2 activation

time errors decreased monotonically with mesh refinement and timestep reduction, with conduction velocity converging to 0.554 m s^{-1} across all resolutions. The convergence rates are consistent with first-order behaviour on unstructured meshes, as expected for the operator-split Crank-Nicolson scheme³⁸. The use of a flat plate geometry with Delaunay triangulation ensures that the convergence study exercises the same cotangent-weight FEM assembly and Crank-Nicolson solver as the patient-specific LA simulation, in contrast to structured FD benchmarks that would test a different discretisation.

Benchmark validation Existing community benchmarks target detailed ventricular monodomain and bidomain solvers^{44,45}; accordingly, validation against CARP using the identical monodomain formulation and modified Mitchell–Schaeffer model represents the most appropriate benchmark for the present study. JAX-EP demonstrated close agreement with the established cardiac EP simulator CARP at both the cellular and tissue levels.

0D cellular benchmark. At the single-cell level, the mMS action potential morphology produced by JAX-EP was compared against the CARP `bench.pt` reference implementation under identical ionic parameters ($\tau_{\text{in}} = 0.3$, $\tau_{\text{out}} = 5$, $\tau_{\text{open}} = 120$, $\tau_{\text{close}} = 150 \text{ ms}$). JAX-EP reproduced the action potential plateau and repolarisation phases with a mean absolute error of 9.90×10^{-3} and a root-mean-square error of 7.75×10^{-2} in normalised voltage units, confirming consistent ionic model implementation. The minor upstroke discrepancy is attributable to the smooth sigmoid gate approximation adopted in JAX-EP to ensure end-to-end differentiability, in contrast to the hard threshold gate used in CARP.




Figure 8. 0D single-cell benchmark: JAX-EP vs CARP. (A) Full action potential (AP) morphology over 500 ms; JAX-EP (blue dashed) overlaid on the CARP `bench.pt` reference (black solid), showing close agreement in the plateau and repolarisation phases. (B) Upstroke alignment zoom (10–40 ms); the minor discrepancy reflects the smooth sigmoid gate approximation in JAX-EP versus the hard-threshold gate in CARP, a deliberate design choice to ensure end-to-end differentiability. (C) Numerical residual (JAX-EP – CARP); mean absolute error = 9.90×10^{-3} and RMSE = 7.75×10^{-2} in normalised voltage units.

2D tissue benchmark. At the tissue level, JAX-EP demonstrated close agreement with CARP for homogeneous tissue propagation on an identical $20 \times 7 \text{ mm}$ flat plate geometry (Fig. 9). Both solvers reproduced the prescribed pacing protocol and achieved complete tissue activation. JAX-EP yielded an activation duration of 47.1 ms (activation interval: 10.1–57.2 ms) compared with 45.2 ms (10.0–55.2 ms) for CARP, with corresponding conduction velocities of 0.478 m/s and 0.504 m/s, respectively. Direct node wise comparison of local activation times yielded an L_2 error of 1.366 ms, a maximum absolute error of 2.296 ms, and a mean bias of 1.339 ms (JAX-EP minus CARP). The observed inter-solver variability compares favourably




Figure 9. Benchmark validation of JAX-EP against CARP on a 20×7 mm flat plate geometry. (A) Local activation time (LAT) map computed by JAX-EP ($\Delta x = 0.2$ mm, $\Delta t = 0.05$ ms, $\beta = 140$ mm $^{-1}$, mMS ionic model). (B) Corresponding LAT map produced by CARP under identical mesh, fibre, ionic, and stimulus conditions. (C) Point-wise LAT difference (JAX-EP – CARP), with L2 error of 1.37 ms and maximum deviation of 2.30 ms. (D) Node wise scatter of JAX-EP versus CARP LATs; the dashed line denotes the identity. (E) Distribution of LAT errors across all 14,271 nodes, with mean bias of 1.34 ms and conduction velocity agreement within 2.59%, consistent with inter-solver variability reported in the N-version benchmark⁴⁴.

with the tolerances established in the N-version benchmark of Niederer et al.⁴⁴, confirming the numerical fidelity of the JAX-EP implementation across both cellular and tissue scales.

Biological identifiability. Within the mMS model and the parameter sets evaluated here, gradient-based inversion recovers all five parameters with sub-percent accuracy, confirming that the computational identifiability problem is well-posed for this model at this scale. Extension to biophysically richer ionic models — such as the Courtemanche human atrial or Ten Tusscher ventricular formulations — will introduce stronger parameter interdependencies and a more complex loss landscape where identifiability from EGM morphology alone cannot be assumed. Translation to clinical recordings introduces a further challenge: EGM morphology reflects the combined influence of ionic dynamics, fibrosis, fibre orientation, wall thickness, and catheter contact, none of which are independently observable. The extent of meaningful personalisation will ultimately be a clinical question as much as a computational one.

Translational vision. Beyond parameter learning, JAX-EP establishes a differentiable computational framework for the next generation of cardiac EP solvers. By combining automatic differentiation with GPU-accelerated forward simulation, it enables gradient-based optimisation on clinical-resolution meshes within practical runtimes, making large-scale parameter estimation, sensitivity analysis, uncertainty quantification, and virtual cohort studies computationally feasible on commodity GPU hardware.

The proposed inversion framework is likewise not restricted to unipolar EGM waveform fitting. Because gradients are propagated through the complete forward monodomain model, any differentiable objective can be incorporated into the optimisation, including LAT maps, conduction velocity, restitution curves, activation–recovery intervals, electrogram-derived features, or composite multi-objective loss functions. This flexibility enables diverse clinical and experimental measurements

to be assimilated within a single mechanistic framework, rather than calibrating individual electrophysiological properties in isolation.

Ultimately, the clinical opportunity extends well beyond reproducing activation maps or EGM time series. A differentiable cardiac EP solver provides a computational foundation for functionally calibrated patient-specific digital twins capable of recovering the underlying tissue and cellular electrophysiology, thereby enabling quantitative characterisation of the arrhythmogenic substrate. Such models may provide a mechanistic framework for evaluating structural imaging biomarkers such as LGE-MRI, investigating AF inducibility and AF burden, assessing pharmacological response, and prospectively evaluating alternative catheter ablation strategies before lesions are delivered.

More broadly, differentiable cardiac EP establishes a computational engine upon which future optimisation, uncertainty-aware inference, and clinical data assimilation frameworks can be built, bringing cardiac EP digital twins closer to prospective, patient-specific decision support.

Limitations and future directions. Several limitations should be acknowledged. First, parameter learning was demonstrated on simulated rather than clinical OEGM data. The primary future direction is validation against real HD grid recordings from patients undergoing EP study, where measurement noise, catheter contact variability, and ionic model mismatch will present additional challenges not captured in the current in silico framework.

Second, the mMMS ionic model, while sufficient for capturing action potential morphology and restitution, does not represent the full complexity of human atrial EP, including rate dependent remodelling and fibrosis mediated conduction disturbances. Extension to biophysically detailed ionic models³⁹ is a natural next step, though this increases the number of free parameters M , directly scaling the memory and computational cost of forward mode autodiff, which requires M forward passes and $\mathcal{O}(M \times N_p)$ memory (Fig. 10). Reverse mode autodiff, which would reduce gradient computation to two passes regardless of M , remains infeasible due to gradient explosion across the long time horizon N_T ³⁸. Future work may circumvent this through continuous adjoint formulations or implicit differentiation of the ionic ODE, which bypass the unstable backward unrolling entirely while retaining the two pass efficiency of reverse mode autodiff.

Third, the monodomain solver employs an iterative conjugate gradient (CG) solve at each timestep. No preconditioning is currently applied; an algebraic multigrid or incomplete Cholesky preconditioner within the JAX CG loop would reduce iteration count and wall time, particularly for fine clinical meshes with high condition number. Fourth, the lead field model adopts the infinite medium approximation, setting $\sigma_r = 1$ and neglecting boundary effects from tissue heterogeneity. Extension to clinical data will require a physically consistent bidomain partitioning, $\sigma_r = G_{il}/(G_{il} + \sigma_e)$, updated at each optimisation iteration as G_{il} is recovered. Furthermore, because JAX-EP operates on the atrial endocardial surface rather than computing extracellular potentials at the body surface, the monodomain formulation does not capture the far-field and scar-conduction characteristics that a bidomain model would resolve; a bidomain extension would be required primarily if body-surface ECG signals, rather than atrial EGMs, were to be simulated, at substantially greater computational cost.

Finally, parameter learning was performed on a localised 2,894-node patch of the LA surface beneath the HD grid recording site. Scaling to spatially heterogeneous parameter fields across the full LA surface would require distributed GPU memory management and region specific optimisation strategies beyond the scope of the current implementation. Additionally, the identifiability of individual parameters from a single S1–S2 recording has not been formally characterised: τ_{in} and G_{il} are strongly constrained by activation timing, whilst τ_{open} and τ_{close} are more weakly identifiable from the two-beat S1–S2 protocol and may benefit from multi-beat or restitution protocol targets.

The current validation is conducted entirely in silico within the mMMS model. Biological identifiability, the ability to uniquely recover tissue parameters from clinical EGMs in the presence of fibrosis, anisotropy, and catheter variability, remains an open question that will require adapted optimisation targets and clinically motivated constraints beyond the scope of the present study.

The deterministic parameter learning demonstrated here represents an important first step towards clinical translation. In practice, patient-specific electrophysiological calibration should account for measurement noise, catheter contact variability, spatial registration uncertainty, and model mismatch, requiring optimisation over distributions of plausible solutions rather than a single waveform. Coupling JAX-EP with established probabilistic inference and uncertainty quantification frameworks therefore represents a natural way forward towards robust and clinically interpretable parameter recovery.

Future work will prioritise clinical validation against real HD grid recordings, extension to spatially heterogeneous parameter fields, integration with LGE-MRI fibrosis maps for substrate informed initialisation⁶, and application to the electro-optical flow pipeline⁴⁰ for atrial arrhythmia ablation planning. Ensemble simulations via `jax.vmap` can also substantially reduce the computational burden of atrial fibrillation (AF) inducibility studies, where large parameter sweeps across substrate conditions are required to assess arrhythmia vulnerability⁶. The differentiable solver architecture established here provides the computational foundation for each of these extensions. More broadly, we envisage this framework integrating with complementary patient-specific datasets, including LGE-MRI-derived fibrosis, electroanatomic voltage mapping, conduction velocity, activation timing, and continuous rhythm monitoring (e.g. implantable loop recorder AF burden), towards progressively more personalised digital

twins.

Ethical approval

The clinical data used in this study were obtained from the HEAT-AF investigation, which was approved by the UK National Research Ethics Service South Central, Hampshire A Research Ethics Committee (REC Reference: 20/SC/0071; IRAS Reference: 263083). All procedures were conducted in accordance with the standards set by the Declaration of Helsinki. Written informed consent was obtained from each participant prior to enrolment. The present study was not prospectively registered in a clinical trials database because it comprised an observational computational modelling study utilising data collected as part of routine clinical care and an existing approved clinical investigation (HEAT-AF)⁴⁶.

Methods

The Methods section is structured as follows. We first introduce the monodomain PDE and the mMS ionic model in Section *Governing equations*. Section *Spatial discretisation* derives the cotangent weight FE stiffness matrix and its anisotropic conductivity extension, followed by Section *Time integration*, which presents the Strang–Crank–Nicolson time integration scheme. Section *Differentiable programming in JAX* forms the core methodological contribution, detailing the differentiable computational graph, the JAX transformation pipeline, forward- and reverse-mode autodiff, gradient checkpointing for memory management, forward sensitivity analysis, and gradient verification. Section *Computation of observables* describes the computation of extracellular OEGMs and differentiable activation-time maps, while Section *Solver verification* presents numerical convergence and differentiability analyses.

Governing equations

JAX-EP solves the monodomain model⁴⁷ of cardiac electrophysiology on a two-dimensional Riemannian surface $\Gamma \subset \mathbb{R}^3$. The state of the system at time t is described by the vector of state variables $\mathbf{q}(\mathbf{x}, t) = [V(\mathbf{x}, t), h(\mathbf{x}, t)]^\top$, where $V(\mathbf{x}, t) \in [0, 1]$ is the normalised transmembrane voltage and $h(\mathbf{x}, t) \in [0, 1]$ is the inactivation gating variable. The governing equations are:

$$\beta C_m \frac{\partial V}{\partial t} = \nabla_\Gamma \cdot (\mathbf{D}(\mathbf{x}; G_{il}) \nabla_\Gamma V) + I_{ion}(V, h; \boldsymbol{\theta}_{ion}) + I_{stim}(\mathbf{x}, t), \quad (1)$$

$$\frac{\partial h}{\partial t} = \mathcal{G}(V, h; \tau_{open}, \tau_{close}), \quad (2)$$

where ∇_Γ denotes the surface gradient operator, β (cm^{-1}) is the membrane surface-to-volume ratio, and C_m ($\mu\text{F cm}^{-2}$) is the membrane capacitance. Equation (1) is a reaction–diffusion equation describing how the transmembrane voltage evolves in time: the left-hand side represents the capacitive current stored across the cell membrane, while on the right-hand side the diffusion term $\nabla_\Gamma \cdot (\mathbf{D} \nabla_\Gamma V)$ governs the passive spread of electrical activity across the tissue surface (with the conductivity tensor $\mathbf{D}$ setting the direction and speed of propagation), the reaction term I_{ion} captures the nonlinear ionic currents flowing through the cell membrane, and I_{stim} is an externally applied stimulus current used to initiate excitation¹⁰. Equation (2) is an ordinary differential equation governing the gating variable h , whose function $\mathcal{G}$ switches h between its open and closed states on the timescales τ_{open} and τ_{close} , thereby controlling the recovery and re-excitability of the tissue.

Following Corrado et al.⁵, the gate threshold $V_g = 0.13$ and the inward current threshold $a_c = 0.13$ are fixed at physiologically estimated values for human atrial tissue, leaving four ionic parameters and the longitudinal conductivity as the identifiable parameter vector:

$$\boldsymbol{\theta} = [\tau_{in}, \tau_{out}, \tau_{open}, \tau_{close}, G_{il}] \in \mathbb{R}^5. \quad (3)$$

Here τ_{in} and τ_{out} are the time constants of the inward (depolarising) and outward (repolarising) ionic currents that set the upstroke and duration of the action potential, while τ_{open} and τ_{close} govern the rates at which the gating variable recovers and inactivates, and thus control the refractory period; G_{il} is the longitudinal conductivity setting conduction velocity along the fibre direction.

JAX-EP is designed such that scalar, vector, and matrix quantities arising in the discretisation are differentiable functions of $\boldsymbol{\theta}$ under JAX autodiff. The only departure from the standard mMS formulation is the replacement of the piecewise constant ionic gate switch (6) with the smooth sigmoid approximation (8), ensuring compatibility with both reverse mode autodiff (autodiff) for spatial sensitivity analysis and central FD gradient evaluation for parameter learning.

Anisotropic conductivity tensor

Spatially varying fibre directions $\hat{\mathbf{f}}(\mathbf{x})$ were mapped onto the patient-specific LA mesh using the human atrial fibre atlas proposed by Roney et al.³³, which constructs a mean fibre field from high resolution diffusion tensor weighted MRI (DTMRI) data and transfers it to patient-specific anatomies via a universal atrial coordinate (UAC) system⁴⁸. The conductivity tensor encodes fibre aligned anisotropy:

$$\mathbf{D}(\mathbf{x}; G_{il}) = \frac{G_{il}}{\beta C_m} \hat{\mathbf{f}}(\mathbf{x}) \otimes \hat{\mathbf{f}}(\mathbf{x}) + \frac{G_{it}}{\beta C_m} (\mathbf{P}_\Gamma(\mathbf{x}) - \hat{\mathbf{f}}(\mathbf{x}) \otimes \hat{\mathbf{f}}(\mathbf{x})), \quad (4)$$

where $\hat{\mathbf{f}}(\mathbf{x})$ is the local unit fibre direction, $\mathbf{P}_\Gamma = \mathbf{I} - \hat{\mathbf{n}} \otimes \hat{\mathbf{n}}$ projects onto the surface tangent plane ($\hat{\mathbf{n}}$ is the outward unit normal), G_{il} (mS cm⁻¹) is the longitudinal conductivity, and $G_{it} = G_{il}/4$ the transverse conductivity (characterized by a 4:1 conductivity anisotropy ratio based on established histological⁴⁹, *in vivo*⁵⁰, and computational^{51,52} studies). The diffusion tensor $\mathbf{D}$ depends linearly on G_{il} , allowing the derivative $\partial \mathbf{D} / \partial G_{il}$ to be evaluated analytically.

Ionic model

The ionic current follows the modified Mitchell–Schaeffer (mMS) two-current model of Corrado and Niederer⁵, which extends the original MS model⁵³ to eliminate spurious pacemaker behaviour by gating the outward current with $(1 - h)$ and shifting the inward current threshold to V_g :

$$I_{\text{ion}}(V, h; \boldsymbol{\theta}_{\text{ion}}) = \frac{hV(V - V_g)(1 - V)}{\tau_{\text{in}}} - \frac{(1 - h)V}{\tau_{\text{out}}}, \quad (5)$$

where $V_g = 0.13$ is the dimensionless gate threshold. The original mMS gate dynamics use a piecewise constant switch:

$$\frac{\partial h}{\partial t} = \begin{cases} \frac{1 - h}{\tau_{\text{open}}} & V < V_g, \\ \frac{-h}{\tau_{\text{close}}} & V \geq V_g. \end{cases} \quad (6)$$

This switch is non-differentiable at $V = V_g$ as the derivative is zero almost everywhere and undefined at the transition, producing NaN or vanishing gradients through h under automatic differentiation. JAX-EP replaces (6) with a smooth sigmoid gate:

$$\mathcal{G}(V, h; \tau_{\text{open}}, \tau_{\text{close}}) = \frac{1 - h}{\tau_{\text{open}}} (1 - \sigma_\lambda(V)) - \frac{h}{\tau_{\text{close}}} \sigma_\lambda(V), \quad (7)$$

where:

$$\sigma_\lambda(V) = (1 + e^{-\lambda(V - V_g)})^{-1}, \quad \lambda = 150. \quad (8)$$

The sigmoid σ_λ provides a smooth approximation to the Heaviside switch, satisfying $\sigma_\lambda(V) < 10^{-6}$ for $V < V_g - \delta$ and $\sigma_\lambda(V) > 1 - 10^{-6}$ for $V > V_g + \delta$, where $\delta = \lambda^{-1} \ln(10^6) \approx 0.092$ with $\lambda = 150$. The approximation error $|\sigma_\lambda(V) - \mathcal{H}(V - V_g)|$ exceeds 10^{-2} only within the interval $|V - V_g| < \lambda^{-1} \ln 99 \approx 0.031$, a narrow voltage interval encountered only briefly during the action potential (AP) upstroke and therefore without physiological consequence. Unlike the Heaviside, σ_λ is infinitely differentiable on $\mathbb{R}$, with $d\sigma_\lambda/dV = \lambda \sigma_\lambda(V) (1 - \sigma_\lambda(V)) > 0$ everywhere, enabling unobstructed gradient flow through the gate variable h under reverse mode autodiff. This smooth substitution constitutes the *only* departure from the mMS formulation of Corrado and Niederer³⁷; the ionic current equation (5) is retained without alteration, and the parameter set $\boldsymbol{\theta}_{\text{ion}}$ carries identical physiological meaning in both formulations.

Spatial discretisation

Cotangent weight FE operator

The surface Γ is approximated by a triangulated mesh $\mathcal{T} = (V_h, E_h, F_h)$, where V_h , E_h , and F_h are the sets of vertices, edges, and triangular faces, with cardinalities N_p , N_e , and N_f respectively.

We discretise the Laplace–Beltrami operator using the cotangent weight scheme⁵⁴, which arises from the Galerkin weak formulation with piecewise linear basis functions on $\mathcal{T}$; it provides a sparse, symmetric positive semi-definite stiffness matrix. In this scheme, the contribution of each mesh edge (i, j) to the discrete operator is weighted by the sum of the cotangents of the two angles opposite to that edge in the adjacent triangles, encoding the local geometry of the surface triangulation into the diffusion operator. For a nodal field $\mathbf{v} \in \mathbb{R}^{N_p}$:

$$(\mathbf{K}_0 \mathbf{v})_i = \sum_{j \in \mathcal{N}(i)} w_{ij} (v_j - v_i), \quad w_{ij} = \frac{1}{2} (\cot \alpha_{ij}^{(1)} + \cot \alpha_{ij}^{(2)}), \quad (9)$$

where $\mathcal{N}(i)$ is the one-ring of node i and $\alpha_{ij}^{(1,2)}$ are the angles opposite edge (i, j) in the two triangles sharing it. $\mathbf{K}_0 \in \mathbb{R}^{N_p \times N_p}$ is the isotropic (unit-conductivity) cotangent stiffness matrix. The full derivation from the weak formulation to equation (9), including the lumped mass approximation and the resulting semi discrete system, is given in Appendix A.

Anisotropic stiffness matrix

The isotropic cotangent weights w_{ij} are scaled by the effective conductivity along each mesh edge. This conductivity is obtained by projecting the anisotropic conductivity tensor (equation (4)) onto the local edge direction $\hat{\mathbf{e}}_{ij}$. For edge (i, j) with unit direction $\hat{\mathbf{e}}_{ij}$ and surface-projected fibre direction $\hat{\mathbf{f}}_{ij}$, the effective conductivity weighted coefficient is:

$$\tilde{w}_{ij}(G_{\text{il}}) = |w_{ij}| \left[\frac{G_{\text{il}}}{\beta C_m} \cos^2 \varphi_{ij} + \frac{G_{\text{it}}}{\beta C_m} (1 - \cos^2 \varphi_{ij}) \right], \quad (10)$$

where $\varphi_{ij} = \arccos(\hat{\mathbf{e}}_{ij} \cdot \hat{\mathbf{f}}_{ij})$. The geometric quantities $\{w_{ij}, \hat{\mathbf{e}}_{ij}, \hat{\mathbf{f}}_{ij}, \varphi_{ij}\}$ are mesh dependent constants precomputed once before simulation. The only parameter dependent term in equation (10) is G_{il} , and $\partial \tilde{w}_{ij} / \partial G_{\text{il}}$ admits a closed form derivative for all $G_{\text{il}} > 0$. The anisotropic stiffness matrix $\mathbf{K}(\boldsymbol{\theta}) \in \mathbb{R}^{N_p \times N_p}$ assembles as:

$$K_{ij}(\boldsymbol{\theta}) = \begin{cases} -\tilde{w}_{ij}(G_{\text{il}}) & j \in \mathcal{N}(i), \\ \sum_{k \in \mathcal{N}(i)} \tilde{w}_{ik}(G_{\text{il}}) & i = j, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

The lumped mass matrix $\mathbf{M} \in \mathbb{R}^{N_p \times N_p}$ is diagonal⁵⁵ with:

$$M_{ii} = C_m \max(A_i \times 10^{-8}, A_{\min}), \quad (12)$$

where A_i is the Voronoi nodal area and $A_{\min} = P_5(\{A_j\})$ is the 5th percentile of the nodal area distribution, introduced to prevent near zero diagonal entries arising from degenerate boundary nodes or poorly conditioned surface triangles. Lumping preserves the sparsity of the linear system⁵⁶ and avoids global mass matrix solves, which are computationally inefficient within JAX's compiled `jax.lax.scan` time stepping framework.

Time integration

Strang operator splitting

We advance equations (1)–(2) using second-order Strang operator splitting⁵⁷, which separates the stiff linear diffusion operator from the nonlinear ionic source:

$$(V^{k+1}, h^{k+1}) = \mathcal{S}_{\Delta t/2} \circ \mathcal{D}_{\Delta t} \circ \mathcal{S}_{\Delta t/2}(V^k, h^k), \quad (13)$$

Here the superscript k indexes the global timestep, and the superscript (l) (introduced in the ionic sub-step below) indexes the forward-Euler sub-step within a single global step Δt . where $\mathcal{S}_\tau$ is the ionic flow map (equations (5)–(8) over time τ) and $\mathcal{D}_\tau$ is the diffusion propagator (equation (1) with $I_{\text{ion}} = 0$ over τ). Strang splitting is second order accurate in Δt and decomposes each timestep into separate diffusion and ionic substeps, both of which are fully differentiable under JAX autodiff. Detailed solver parameters are contained in Table 4.

Ionic sub-step

The ionic flow map $\mathcal{S}_{\Delta t/2}$ integrates the voltage equation using n_{sub} explicit forward-Euler sub-steps of size $\delta t = \Delta t / (2n_{\text{sub}})$:

$$V^{(l+1)} = \text{clip}(V^{(l)} + \delta t I_{\text{ion}}(V^{(l)}, h^{(l)}; \boldsymbol{\theta}_{\text{ion}}), 0, 1). \quad (14)$$

The gate variable h is updated using the Rush–Larsen scheme⁵⁸, which solves the linear gate ODE analytically over each sub-step:

$$h^{(l+1)} = h_{\infty}(V^{(l)}) + (h^{(l)} - h_{\infty}(V^{(l)})) \exp\left(-\frac{\delta t}{\tau_h(V^{(l)})}\right), \quad (15)$$

where $h_{\infty}(V) = \alpha_h(V) / (\alpha_h(V) + \beta_h(V))$ and $\tau_h(V) = 1 / (\alpha_h(V) + \beta_h(V))$ are the steady-state gate value and time constant, with $\alpha_h = (1 - \sigma_{\lambda}(V)) / \tau_{\text{open}}$ and $\beta_h = \sigma_{\lambda}(V) / \tau_{\text{close}}$, following the sigmoid gate approximation of equation (8). The Rush–Larsen scheme is unconditionally stable for the gate ODE regardless of sub-step size, and its exponential form is continuously differentiable with respect to τ_{open} and τ_{close} for all $V \in [0, 1]$. Clip operations enforce physical bounds and are differentiable except on the measure-zero boundary $\{0, 1\}$, where sub-gradients are used.

Differentiability of the ionic sub-step. Reverse mode autodiff through the full sequence of N_T timesteps suffers from gradient explosion due to the accumulation of Jacobian eigenvalues across the scan loop, a well documented instability for long time integrations⁵⁹. We therefore employ forward mode autodiff (`jax.jacfwd`), which propagates a tangent vector alongside the primal computation at each timestep without requiring a backward pass. This eliminates the exponential gradient accumulation and yields numerically stable, exact gradients for all five ionic parameters θ_{ion} in five forward passes which is equivalent in cost to half the number of FD evaluations while avoiding approximation error.

Crank-Nicolson diffusion step

The diffusion propagator $\mathcal{D}_{\Delta t}$ is implemented with the Crank-Nicolson implicit scheme⁶⁰, which is unconditionally stable and second order accurate in Δt :

$$\underbrace{(\mathbf{M} + \frac{\Delta t}{2} \mathbf{K}(\boldsymbol{\theta}))}_{\mathcal{A}} \mathbf{V}^{k+1} = \underbrace{(\mathbf{M} - \frac{\Delta t}{2} \mathbf{K}(\boldsymbol{\theta}))}_{\mathcal{B}} \mathbf{V}^k. \quad (16)$$

The operator $\mathcal{A}$ is symmetric positive definite (since $\mathbf{M} \succ 0$ and $\mathbf{K} \succeq 0$) and is solved by the conjugate-gradient (CG) method (`jax.scipy.sparse.linalg.cg`) with tolerance $\epsilon_{\text{CG}} = 10^{-6}$. Both $\mathcal{A}$ and $\mathcal{B}$ depend on G_{il} through $\mathbf{K}(\boldsymbol{\theta})$, so gradients with respect to G_{il} must propagate through the linear solve. Under forward mode autodiff, JAX propagates the tangent through the CG iteration directly; under reverse-mode autodiff the adjoint of the solve $\mathbf{V}^{k+1} = \mathcal{A}^{-1} \mathbf{r}$ is the symmetric adjoint linear solve $\boldsymbol{\lambda} = \mathcal{A}^{-\top} \mathbf{g} = \mathcal{A}^{-1} \mathbf{g}$, adding exactly one additional CG solve per timestep in the backward pass. Here $\mathbf{r} = \mathcal{B} \mathbf{V}^k$ is the Crank–Nicolson right-hand side, $\mathbf{g}$ is the incoming cotangent (output-gradient) vector supplied by reverse-mode autodiff, and $\boldsymbol{\lambda}$ is the corresponding adjoint solution. Both paths are handled automatically by JAX without manual adjoint derivation.

Parameter	Symbol	Value	Unit
Global timestep	Δt	0.05	ms
Ionic sub-steps	n_{sub}	2	—
CG tolerance	ϵ_{CG}	10^{-6}	—
CG max iterations	—	50	—
Surface-to-volume	β	140	cm^{-1}
Membrane capacitance	C_m	1	$\mu\text{F cm}^{-2}$
Sigmoid sharpness	λ	150	—
Stimulus amplitude	I_{stim}	200	$\mu\text{A cm}^{-2}$
Stimulus duration	Δt_s	2	ms

Table 4. JAX-EP numerical parameters.

Computation of observables

Omnipolar EGM computation

Unipolar extracellular potentials at all $N_e = 16$ HD grid electrodes are computed from the transmembrane voltage field via the infinite medium approximation to the quasi-static lead field model⁶¹:

$$\phi_e(t) = [\mathbf{W} \mathbf{V}(t)]_e, \quad W_{ei} = \frac{1}{4\pi \|\mathbf{x}_i - \mathbf{x}_e\|}, \quad (17)$$

where $\boldsymbol{\Phi}$ is the vector of unipolar extracellular potentials ϕ_e across all electrodes and $\mathbf{W} \in \mathbb{R}^{N_e \times N_p}$ is a fixed geometric lead-field matrix depending only on electrode-to-node distances. The conductivity scaling factor $\sigma_r = G_{\text{il}} / (G_{\text{il}} + 0.2)$, arising from bidomain partitioning, is set to unity under the assumption of a homogeneous extracellular medium. where σ_e (here $\sigma_e = 0.2 \text{ mS cm}^{-1}$) is the extracellular conductivity.

The implications of this approximation are discussed in the limitations section. Since $\mathbf{W}$ is held fixed and applied identically to both target and simulated OEGMs, amplitude scaling cancels in the waveform MSE loss and does not affect parameter recovery. OEGMs are formed at each 2×2 clique as the spatial Laplacian of the unipolar potentials⁶²:

$$\text{OEGM}_c(t) = \frac{1}{4} (\phi_{c_1}(t) - \phi_{c_2}(t) - \phi_{c_3}(t) + \phi_{c_4}(t)) \quad (18)$$

This signed four electrode combination suppresses far field contributions and approximates the second spatial derivative of the extracellular potential within the clique, consistent with the HD grid spatial gradient approach³⁶. where $c_1, \dots, c_4$ denote

the four electrodes of the 2×2 HD grid clique. As a fixed linear map, $\mathbf{W}$ constitutes a trivially differentiable layer in the computational graph (21), with $\partial \mathcal{L} / \partial \mathbf{V}^k = \mathbf{W}^\top (\partial \mathcal{L} / \partial \Phi)$. All nine clique OEGMs are fitted simultaneously during parameter learning.

Differentiable activation-time map

Activation time (AT) at mesh node i is the time of maximum depolarisation rate. The standard definition $t_i^{\text{AT}} = \text{argmax}_t (\partial V_i / \partial t)$ is non-differentiable due to the hard argmax. We replace it with a soft temporal attention map accumulated in the `scan` carry, requiring no trajectory storage:

$$\omega_i^k = \sigma_{\beta_{\text{AT}}} (V_i^k - V_g), \quad (19)$$

$$\hat{t}_i^{\text{AT}} = \frac{\sum_{k=1}^{N_T} t_k \omega_i^k}{\sum_{k=1}^{N_T} \omega_i^k + \varepsilon}, \quad (20)$$

with $\beta_{\text{AT}} = 50$ and $\varepsilon = 10^{-6}$. where $\sigma_{\beta_{\text{AT}}}$ is the logistic sigmoid of Eq. (8) with steepness β_{AT} , and ω_i^k is the resulting soft attention weight at node i and timestep k .

The soft AT map is an attention weighted centroid of the time axis, with attention scores peaked near the activation threshold V_g . It is continuously differentiable in all entries of $\boldsymbol{\theta}$ by construction, enables $\partial \hat{t}_i^{\text{AT}} / \partial \boldsymbol{\theta}$ via `jax.grad`, and requires only two $\mathbb{R}^{N_p}$ running accumulators in the scan carry with $\mathcal{O}(N_p)$ memory overhead, making it suitable for gradient-based optimisation directly on activation times in future applications.

Differentiable programming in JAX

This section is the core methodological contribution of JAX-EP. We describe: (1) the construction of the simulation as a single differentiable computational graph; (2) the JAX transformation pipeline (`jit`, `scan`, `grad`, `checkpoint`, `vmap`); (3) forward mode and reverse mode differentiation; (4) forward sensitivity analysis; and (5) gradient verification.

The simulation as a differentiable computational graph

The full JAX-EP simulation constitutes a pure function $\mathcal{F} : \mathbb{R}^5 \rightarrow \mathbb{R}^{N_e \times N_T}$ mapping the parameter vector $\boldsymbol{\theta}$ to the matrix of extracellular electrograms $\Phi = [\phi_e(t_k)]_{e,k}$, where $e \in \{1, \dots, N_e\}$ indexes the recording electrodes and $k \in \{1, \dots, N_T\}$ indexes the timesteps. The graph is a composition:

$$\boldsymbol{\theta} \xrightarrow[\text{(10)}]{\text{assemble}} \mathbf{K}(\boldsymbol{\theta}) \xrightarrow[\text{(13)}]{\text{time-march}} \{\mathbf{V}^k\}_{k=1}^{N_T} \xrightarrow[\text{(17)}]{\mathbf{W}^\cdot} \Phi \xrightarrow[\text{(20)}]{\text{soft argmax}} \hat{\mathbf{t}}^{\text{AT}}, \quad (21)$$

where each stage is a differentiable map. JAX traces this graph once using abstract evaluation and lowers the traced computation to an XLA computation graph, which is then compiled, device placed, and executed without Python overhead.

JIT compilation and lax.scan

The time marching loop over N_T timesteps is expressed as:

$$(\mathbf{V}^{N_T}, \mathbf{h}^{N_T}), \{\phi(t_k)\}_{k=1}^{N_T} = \text{jax.lax.scan}(\mathcal{M}_{\boldsymbol{\theta}}, (\mathbf{V}^0, \mathbf{h}^0), \{s_k\}_{k=1}^{N_T}), \quad (22)$$

where $\mathcal{M}_{\boldsymbol{\theta}}$ is the Strang-split step (13) and $s_k \in \{0, 1\}$ is the stimulus flag. `jax.lax.scan` lowers the time loop into a single XLA `while_loop` with statically known iteration count, avoiding Python level interpreter overhead at each time step; its vector-Jacobian product is the reverse scan over the adjoint step, rendering the entire loop natively differentiable.

The complete simulation, including all N_T conjugate gradient solves, is compiled to a single XLA executable through `jax.jit`. Here $N_T = \lceil T_{\text{sim}} / \Delta t \rceil$ denotes the total number of time steps over the simulation duration T_{sim} , and N_p denotes the number of nodes in the patient specific surface mesh.

We distinguish between *static* and *dynamic* parameters. Static parameters (mesh topology, N_T , N_p) define the structure of the XLA graph and require recompilation if changed. Dynamic parameters (the entries of $\boldsymbol{\theta}$, initial conditions, pacing schedule) are *traced* through the graph and can change between calls without recompilation. This design follows the TORAX convention³² and ensures that the gradient $\nabla_{\boldsymbol{\theta}} \mathcal{L}$ is computed with zero compilation overhead after the first call.

1. **Omnipolar EGMs** $\Phi(t) = \mathbf{WV}(t)$ followed by $\text{OEGM}_c(t) = \frac{1}{4}(\Phi_{c_1} - \Phi_{c_2} - \Phi_{c_3} + \Phi_{c_4})$: $\partial \mathcal{L} / \partial \mathbf{V}^k = \mathbf{W}^\top (\partial \mathcal{L} / \partial \Phi)$; $\mathcal{O}(N_e N_p)$; $\mathbf{W}$ is a fixed geometric matrix ($\sigma_r = 1$, no G_{II} dependence); all nine HD grid clique OEGMs are fitted simultaneously.

2. **Crank-Nicolson diffusion** $\mathbf{V}^{k+1} = \mathcal{A}^{-1} \mathcal{B} \mathbf{V}^k$: the adjoint $\boldsymbol{\lambda}^k = \mathcal{A}^{-\top} \mathbf{g}^k$ requires one additional CG solve per timestep, where $\mathbf{g}^k = \partial \mathcal{L} / \partial \mathbf{V}^{k+1} + \mathcal{B}^\top \boldsymbol{\lambda}^{k+1}$ is propagated backwards through the scan. The gradient with respect to G_{il} accumulates as: $\partial \mathcal{L} / \partial G_{il} += (\boldsymbol{\lambda}^k)^\top (-(\partial \mathcal{A} / \partial G_{il}) \mathbf{V}^{k+1} + (\partial \mathcal{B} / \partial G_{il}) \mathbf{V}^k)$, where $\partial \mathcal{A} / \partial G_{il} = \frac{\Delta t}{2} \partial \mathbf{K} / \partial G_{il}$ is $\mathcal{O}(N_{\text{edges}})$ from equation (10).
3. **Ionic step**: JAX evaluates the VJP of the Strang-split ionic substep (14)–(8) pointwise at each node, exploiting the element wise independence of the ionic ODE; no explicit $2N_p \times 2N_p$ Jacobian is formed.
4. **Stiffness assembly**: the chain rule through equation (10) gives $\partial \tilde{w}_{ij} / \partial G_{il} = |w_{ij}| (\frac{3}{4} \cos^2 \phi_{ij} + \frac{1}{4}) / (\beta C_m)$, an $\mathcal{O}(1)$ expression per edge.

The complete backward pass costs $\mathcal{O}(N_T)$ additional CG solves plus $\mathcal{O}(N_T N_p)$ arithmetic for the ionic Jacobians.

Gradient checkpointing

`jax.checkpoint` is provided as an optional memory optimisation for reverse mode autodiff (`jax.grad`), reducing peak activation storage from $\mathcal{O}(N_T N_p)$ to $\mathcal{O}(\sqrt{N_T N_p})$ during the backward pass³⁵. For the forward mode autodiff path employed in this work (`jax.jacfwd`), no backward pass exists and `jax.checkpoint` is not required. It is retained in the implementation for future reverse mode applications or adjoint based gradient formulations.

Forward mode autodiff

Reverse mode autodiff (`jax.grad`) is ill suited to long time series cardiac PDE solvers because back propagating adjoint vectors through N_T implicit Crank-Nicolson solves requires unrolling the CG iterations into the backward graph, which accumulates ill conditioned gradient structures and causes gradient explosion or device out-of-memory failure. We instead exploit the low parameter dimensionality ($M = 5$) by employing forward mode autodiff (`jax.jacfwd`), which propagates M dual number tangent vectors alongside the primal simulation in a single left-to-right pass (Fig. 10). At each timestep k , the tangent update reduces to a sparse matrix-vector product of dimension $N_p \times M$, computed simultaneously with the primal state:

$$\dot{\mathbf{V}}^{k+1} = \frac{\partial \mathbf{V}^{k+1}}{\partial \boldsymbol{\theta}} \bigg|_{\mathbf{V}^k} \dot{\boldsymbol{\theta}}, \quad (23)$$

where $\dot{\boldsymbol{\theta}} \in \mathbb{R}^M$ is the parameter seed vector. Because forward mode follows the causal direction of time, no intermediate states are stored: the active memory footprint is $\mathcal{O}(M \times N_p)$, independent of N_T and well within 15 GB VRAM for $M = 5$. Each parameter requires one forward pass, yielding exact gradients in five simulations versus ten finite-difference evaluations while eliminating approximation error. Gradient stability was confirmed empirically: ratios of forward mode to finite-difference gradients were within 0.75% of unity for all five parameters across all tested configurations (Fig. 3B). This capability extends naturally to spatial sensitivity analysis: the spatial Jacobian $\partial \overline{AT} / \partial G_{il} \in \mathbb{R}^{N_p}$ over the full LA mesh is computed via two forward passes, identifying regions of high conductivity sensitivity that govern EGM waveform morphology.

Batch simulation via vmap

The forward pass $\mathcal{F}(\boldsymbol{\theta})$ is vectorised over a batch of B parameter vectors with zero code modification via `jax.vmap`:

$$\text{vmap}(\mathcal{F})\left(\{\boldsymbol{\theta}^{(b)}\}_{b=1}^B\right) = \{\mathcal{F}(\boldsymbol{\theta}^{(b)})\}_{b=1}^B. \quad (24)$$

`vmap` lifts $\mathcal{F}$ to operate on a leading batch axis, emitting vectorised XLA operations that run with a constant factor of B additional compute and not B serial calls. When composed with `jit`, the batched simulation is compiled once and executed as a single fused kernel. This is used in the inversion application to evaluate the Nelder–Mead simplex vertices in one batch.

Gradient verification

All autodiff gradients are verified against central FD:

$$g_i^{\text{FD}} = \frac{\mathcal{L}(\boldsymbol{\theta} + \varepsilon \mathbf{e}_i) - \mathcal{L}(\boldsymbol{\theta} - \varepsilon \mathbf{e}_i)}{2\varepsilon}, \quad \varepsilon = 10^{-3}, \quad (25)$$

where $\mathbf{e}_i$ is the i -th standard basis vector. Here $\boldsymbol{\theta}$ is the parameter vector defined in Eq. (3), and ε is a small finite perturbation applied to one parameter at a time to approximate the derivative of the loss $\mathcal{L}$ by central differences; its value (10^{-3}) is small enough to keep the truncation error low while remaining large enough to avoid floating-point cancellation. Agreement is measured by the ratio $r_i = g_i^{\text{autodiff}} / g_i^{\text{FD}}$; a ratio $r_i \approx 1$ for all five parameters confirms correct differentiation. Ratios substantially different from unity may indicate either: (a) a non-differentiable operation in the graph (e.g. a hard `sign` or `argmax`), or (b) numerical instability in the backward pass (e.g. gradient explosion through long ionic scan chains). We report FD gradient ratios for all parameters and all geometries in the Results section.




Figure 10. FD versus forward mode autodiff (`jax.jacfwd`) for exact gradient computation through the JAX-EP monodomain solver. *Left:* FD perturbs each of the $M = 5$ parameters by $\pm\epsilon$, requiring $2M = 10$ forward evaluations per gradient step; gradients are approximate and susceptible to truncation and round-off error. *Right:* `jax.jacfwd` augments the primal state V with a parameter tangent $dV/d\theta_i$ and advances both simultaneously in the causal direction, delivering exact gradients at machine precision in $M = 5$ passes with $\mathcal{O}(M \times N_p)$ constant memory and no backward pass.

Solver verification

JAX-EP is verified at two levels.

(i) **Numerical convergence.** Spatial and temporal convergence is assessed on a 20×7 mm flat plate geometry discretised with an unstructured Delaunay triangulation, using the same cotangent-weight surface FEM and Crank-Nicolson scheme as the full LA simulation. Activation time errors are quantified by the L2 norm relative to a reference solution at the finest resolution, across three spatial ($\Delta x = 0.5, 0.3, 0.2$ mm) and three temporal ($\Delta t = 0.10, 0.05, 0.025$ ms) resolutions.

(ii) **Gradient verification.** End-to-end differentiability is confirmed by comparing FD gradients of a scalar EGM loss with respect to all five parameters $\theta = [\tau_{in}, \tau_{out}, \tau_{open}, \tau_{close}, G_{il}]$, evaluated at two HD grid recording sites under CS and LAA pacing on the full LA mesh, yielding 20 independent gradient evaluations. Gradients are classified as correct if they are finite, non-zero, and physically consistent with the expected parameter sensitivities of the mMS ionic model.

(iii) **Benchmark against CARP.** JAX-EP is further benchmarked against CARP⁶³ on an identical 20×7 mm structured mesh ($\Delta x = 0.2$ mm) with matched ionic parameters, fibre orientation, and stimulus protocol, yielding an L2 activation time agreement of 1.37 ms and conduction velocity agreement within 2.59%, consistent with the inter-solver variability reported in the N-version benchmark of Niederer et al.⁴⁴.

Data and software availability

The computational data and software generated and used during the current study are publicly available through the JAX-EP GitHub repository (<https://github.com/OvJafri/JAX-EP>) and its permanent Zenodo archive (<https://doi.org/10.5281/zenodo.22078184>). The deposited raw computational data include CARP reference traces used for 0D cellular-scale validation, the 2D flat-plate mesh and pacing-site specifications used for CARP-based 2D validation, and simulated ground-truth omnipolar EGMs generated for three patient-specific cases and used for the 3D parameter-learning experiments. The CARP-generated .rgb simulation data from the 2D flat-plate validation are permanently archived in a separate Zenodo repository due to their file size (<https://doi.org/10.5281/zenodo.22060549>). The JAX-EP repository provides the computational code and reproducibility workflows for 0D and 2D solver validation, end-to-end differentiability demonstrations, 3D forward simulations using a representative cube with spatial and temporal resolution comparable to the patient-specific LA mesh, and parameter-learning experiments at the 0D cellular, 1D fibre, and 3D patient-derived LA patch scales. It also includes scripts for generating the reported results and figures, processing and visualising simulated omnipolar EGMs, representative examples, and visual demonstrations of the framework. JAX-EP is implemented in Python 3.9 with JAX 0.7.2²⁹ (XLA backend), NumPy 1.26, and SciPy 1.12, using double-precision arithmetic

(jax_enable_x64=True) for the 3D cube and LA patch simulations, and runs on both CPU and GPU through JAX's device-agnostic backend. The patient-specific LA mesh and HEAT-AF clinical data are not publicly available due to patient privacy and data-sharing restrictions. The GitHub repository serves as the actively maintained software and reproducibility resource, while Zenodo provides the permanent archival records of the code and associated computational data.

Acknowledgements

The authors gratefully acknowledge the participants and clinical teams involved in the HEAT-AF study for providing the clinical data used in this work. The authors acknowledge the use of the Queen Mary University of London, High Performance Computing (HPC) facilities, which supported the development and testing of the computational framework presented in this study. The authors thank the members of the Personalised Cardiac Modelling Group at Queen Mary University of London for valuable scientific discussions throughout this work. C.H.R. acknowledges support from the Future Leaders Fellowship funded by UK Research and Innovation (UKRI), grant number MR/W004720/1.

Author contributions statement

O.A.J. conceived the study, developed the JAX-EP framework, performed the computational experiments, analysed the results, and wrote the manuscript. M.E. contributed to the development of the in silico HD Grid catheter framework, simulated EGM generation and assessment. H.V. curated the clinical imaging and electrophysiological data and prepared the patient-specific datasets. G.S. provided guidance on the computational methodology and differentiable solver framework and contributed to the interpretation of the results. E.J.V. provided senior expertise in computational cardiac electrophysiology and contributed to the modelling methodology and interpretation of the electrophysiological findings. S.H. contributed to the conceptual development of the EGM-based calibration strategy, provided senior clinical guidance on the translational aspects of the study, and contributed to the interpretation of the findings. C.H.R. conceived and supervised the study, secured funding, and provided scientific leadership throughout the project. All authors reviewed, edited, and approved the final manuscript.

Additional information

Competing interests: The authors declare no competing interests.

References

1. Chrispin, J. *et al.* Digital twin-guided ablation for ventricular tachycardia. *New Engl. J. Medicine* DOI: [10.1056/NEJMc2517822](https://doi.org/10.1056/NEJMc2517822) (2026). TWIN-VT trial, Johns Hopkins University.
2. Schuijt, E., Scherr, D., Plank, G., Schotten, U. & Heijman, J. Evolution in electrophysiology 100 years after einthoven: translational and computational innovations in rhythm control of atrial fibrillation. *Europace* **27**, euae304 (2025).
3. Niederer, S. A., Lumens, J. & Trayanova, N. A. Computational models in cardiology. *Nat. Rev. Cardiol.* **16**, 100–111 (2019).
4. Potse, M. *et al.* Patient-specific modelling of cardiac electrophysiology in heart-failure patients. *EP Eur.* **16**, iv56–iv61, DOI: [10.1093/europace/euu257](https://doi.org/10.1093/europace/euu257) (2014).
5. Corrado, C. *et al.* A work flow to build and validate patient specific left atrium electrophysiology models from catheter measurements. *Med. image analysis* **47**, 153–163 (2018).
6. Ehresh, M. *et al.* Comparative multimodal calibration of patient-specific atrial fibrillation models: Impact of imaging and electrophysiology data on arrhythmogenic substrate identification. *The J. Physiol.* DOI: [10.1113/jp290765](https://doi.org/10.1113/jp290765) (2026). Published online June 2026.
7. Diaz, M. P. M. *et al.* Impact of effective refractory period personalization on prediction of atrial fibrillation vulnerability. *Europace* **25**, euad122–542 (2023).
8. Coveney, S. *et al.* Calibrating cardiac electrophysiology models using latent Gaussian processes on atrial manifolds. *Sci. Reports* **12**, 16572, DOI: [10.1038/s41598-022-20745-z](https://doi.org/10.1038/s41598-022-20745-z) (2022).
9. Mamajiwala, M. *et al.* Rapid calibration of atrial electrophysiology models using gaussian process emulators in the ensemble kalman filter. *Sci. Reports* (2026).
10. Tung, L. *A bi-domain model for describing ischemic myocardial dc potentials*. Ph.D. thesis, Massachusetts Institute of Technology (1978).

11. Clayton, R. *et al.* Models of cardiac tissue electrophysiology: progress, challenges and open questions. *Prog. biophysics molecular biology* **104**, 22–48 (2011).
12. Sakata, K. & Trayanova, N. A. Towards mechanisms-driven strategy for persistent atrial fibrillation ablation: Leveraging digital twins. *The J. Precis. Medicine: Heal. Dis.* 100036 (2026).
13. Waight, M. C. *et al.* Heart digital twins predict features of invasive reentrant circuits and ablation lesions in scar-dependent ventricular tachycardia. *Circ. Arrhythmia Electrophysiol.* **18**, e013660 (2025).
14. Bhagirath, P., Strocchi, M., Bishop, M. J., Boyle, P. M. & Plank, G. From bits to bedside: entering the age of digital twins in cardiac electrophysiology. *Europace* **26**, euae295 (2024).
15. Johnston, B. M. & Johnston, P. R. Approaches for determining cardiac bidomain conductivity values: progress and challenges. *Med Biol Eng Comput.* **58**, 2919–2935 (2020).
16. Roney, C. H. *et al.* Constructing bilayer and volumetric atrial models at scale. *Interface focus* **13**, 20230038 (2023).
17. Trayanova, N. A., Boyle, P. M. & Nikolov, P. P. Personalized imaging and modeling strategies for arrhythmia prevention and therapy. *Curr. opinion biomedical engineering* **5**, 21–28 (2018).
18. Zappon, E. *et al.* An efficient end-to-end computational framework for the generation of ecg calibrated volumetric models of human atrial electrophysiology. *Med. Image Analysis* **102**, 103514, DOI: [10.1016/j.media.2025.103514](https://doi.org/10.1016/j.media.2025.103514) (2026).
19. Grandits, T. *et al.* An inverse eikonal method for identifying ventricular activation sequences from epicardial activation maps. *J Comput. Phys* **419**, 109700 (2020).
20. Trayanova, N. A. & Prakosa, A. Up digital and personal: How heart digital twins can transform heart patient care. *Hear. Rhythm.* **21**, 89–99 (2024).
21. Sahli Costabal, F., Yang, Y., Perdikaris, P., Hurtado, D. E. & Kuhl, E. Physics-informed neural networks for cardiac activation mapping. *Front. Phys.* **8**, 42 (2020).
22. Coveney, S., Cantwell, C. & Roney, C. Atrial conduction velocity mapping: clinical tools, algorithms and approaches for understanding the arrhythmogenic substrate. *Med Biol Eng Comput.* 1–16 (2022).
23. Campos, F. O. *et al.* An automated near-real time computational method for induction and treatment of scar-related ventricular tachycardias. *Med. Image Analysis* **80**, 102483 (2022).
24. Qian, S. *et al.* Developing cardiac digital twin populations powered by machine learning provides electrophysiological insights in conduction and repolarization. *Nat. Cardiovasc. Res.* DOI: [10.1038/s44161-025-00650-0](https://doi.org/10.1038/s44161-025-00650-0) (2025).
25. Li, C. *et al.* Cryoballoon roof line and pulmonary vein isolation for persistent atrial fibrillation: A randomized trial. *Circ. Arrhythmia Electrophysiol.* **18**, e013826 (2025).
26. Kunisch, K. & Rund, A. Time optimal control of the monodomain model in cardiac electrophysiology. *IMA J. Appl. Math.* **80**, 1664–1683, DOI: [10.1093/imamat/hxv010](https://doi.org/10.1093/imamat/hxv010) (2015).
27. Nagaiah, C. & Kunisch, K. Higher order optimization and adaptive numerical solution for optimal control of monodomain equations in cardiac electrophysiology. *Appl. Numer. Math.* **61**, 53–65, DOI: [10.1016/j.apnum.2010.08.004](https://doi.org/10.1016/j.apnum.2010.08.004) (2011).
28. Newbury, R. *et al.* A review of differentiable simulators. *IEEE Access* **12**, 97581–97604, DOI: [10.1109/ACCESS.2024.3425448](https://doi.org/10.1109/ACCESS.2024.3425448) (2024).
29. Bradbury, J. *et al.* JAX: composable transformations of Python+NumPy programs (2018).
30. Bezgin, D. A., Bauer, S. J. & Adams, N. A. JAX-Fluids: a fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. *Comput. Phys. Commun.* **282**, 108527, DOI: [10.1016/j.cpc.2022.108527](https://doi.org/10.1016/j.cpc.2022.108527) (2023).
31. Xue, T. *et al.* JAX-FEM: a differentiable GPU-accelerated 3D finite element solver for automatic inverse design and mechanistic data science. *Comput. Phys. Commun.* **291**, 108802, DOI: [10.1016/j.cpc.2023.108802](https://doi.org/10.1016/j.cpc.2023.108802) (2023).
32. Citrin, J. *et al.* TORAX: a fast and differentiable Tokamak transport simulator in JAX. *arXiv preprint arXiv:2406.06718* (2024). ArXiv:2406.06718 [physics.plasm-ph].
33. Roney, C. H. *et al.* A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction. *Comput. Biol Med* **104**, 278–290 (2019).
34. Jaffery, O. A. *et al.* Subject-specific ablation of pathologic conduction patterns beyond the pulmonary veins: A personalised modelling approach. In *2023 Computing in Cardiology (CinC)*, vol. 50, 1–4 (IEEE, 2023).
35. Chen, Z. *et al.* Biophysical modeling predicts ventricular tachycardia inducibility and circuit morphology: a combined clinical validation and computer modeling approach. *J Cardiovasc. Electrophysiol* **27**, 851–860 (2016).

36. Deno, D. C. *et al.* Orientation-independent catheter-based characterization of myocardial activation. *IEEE Trans Biomed Eng* **64**, 1067–1077 (2016).
37. Corrado, C. & Niederer, S. A. A two-variable model robust to pacemaker behaviour for the dynamics of the cardiac action potential. *Math. Biosci.* **281**, 46–54 (2016).
38. Sundnes, J. *et al.* *Computing the electrical activity in the heart*, vol. 1 (Springer Science & Business Media, 2007).
39. Courtemanche, M., Ramirez, R. J. & Nattel, S. Ionic mechanisms underlying human atrial action potential properties: insights from a mathematical model. *Am. J. Physiol. Circ. Physiol.* **275**, H301–H321 (1998).
40. 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 & Electrophysiol. Rev.* **13** (2024).
41. Sachetto Oliveira, R. *et al.* Toward cardiac electrophysiology digital twins with an efficient open source scalable solver on GPU clusters. *Sci. Reports* **15**, DOI: [10.1038/s41598-025-33709-w](https://doi.org/10.1038/s41598-025-33709-w) (2025).
42. Sachetto Oliveira, R. *et al.* Performance evaluation of GPU parallelization, space-time adaptive algorithms and their combination for simulating cardiac electrophysiology. *Int. J. for Numer. Methods Biomed. Eng.* **34**, e2913, DOI: [10.1002/cnm.2913](https://doi.org/10.1002/cnm.2913) (2018).
43. Jaffery, O. A. *et al.* Automated generation of ablation lesion masks: a unison of electro and optic flow mapping for persistent af virtual cohorts. *Europace* **28**, euaf290 (2026).
44. Niederer, S. A. *et al.* Verification of cardiac tissue electrophysiology simulators using an n-version benchmark. *Philos. Transactions Royal Soc. A: Math. Phys. Eng. Sci.* **369**, 4331–4351 (2011).
45. Pathmanathan, P., Bernabeu, M., Niederer, S., Gavaghan, D. & Kay, D. Computational modelling of cardiac electrophysiology: explanation of the variability of results from different numerical solvers. *Int J Numer Method Biomed Eng* **28**, 890–903 (2012).
46. Valli, H. *et al.* High density evaluation of the arrhythmogenic substrate in atrial fibrillation (heat-af). *Hear. Rhythm. O2* (2025).
47. Keener, J. & Sneyd, J. *Mathematical Physiology*, vol. 8 of *Interdisciplinary Applied Mathematics* (Springer, New York, 1998).
48. Roney, C. H. *et al.* Universal atrial coordinates applied to visualisation, registration and construction of patient specific meshes. *Med. image analysis* **55**, 65–75 (2019).
49. Spach, M. S., Dolber, P. C. & Heidlage, J. F. Influence of the passive anisotropic properties on directional differences in propagation following modification of the sodium conductance in human atrial muscle. a model of reentry based on anisotropic discontinuous propagation. *Circ. research* **62**, 811–832 (1988).
50. Cantwell, C. D. *et al.* Techniques for automated local activation time annotation and conduction velocity estimation in cardiac mapping. *Comput. Biol Med* **65**, 229–242 (2015).
51. Kotadia, I. *et al.* Anisotropic cardiac conduction. *Arrhythmia & Electrophysiol. Rev.* **9**, 202 (2020).
52. Harrild, D. M. & Henriquez, C. S. A computer model of normal conduction in the human atria. *Circ. Res.* **87**, e25–e36 (2000).
53. Mitchell, C. C. & Schaeffer, D. G. A two-current model for the dynamics of cardiac membrane. *Bull. Math. Biol.* **65**, 767–793, DOI: [10.1016/S0092-8240\(03\)00041-7](https://doi.org/10.1016/S0092-8240(03)00041-7) (2003).
54. Pinkall, U. & Polthier, K. Computing discrete minimal surfaces and their conjugates. *Exp. Math.* **2**, 15–36, DOI: [10.1080/10586458.1993.10504266](https://doi.org/10.1080/10586458.1993.10504266) (1993).
55. Hinton, E., Rock, T. & Zienkiewicz, O. C. A note on mass lumping and related processes in the finite element method. *Earthq. Eng. Struct. Dyn.* **4**, 245–249, DOI: [10.1002/eqe.4290040305](https://doi.org/10.1002/eqe.4290040305) (1976).
56. Hughes, T. J. R. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis* (Dover Publications, Mineola, NY, 2000).
57. Strang, G. On the construction and comparison of difference schemes. *SIAM J. on Numer. Analysis* **5**, 506–517, DOI: [10.1137/0705041](https://doi.org/10.1137/0705041) (1968).
58. Rush, S. & Larsen, H. A practical algorithm for solving dynamic membrane equations. *IEEE Transactions on Biomed. Eng.* **25**, 389–392, DOI: [10.1109/TBME.1978.326270](https://doi.org/10.1109/TBME.1978.326270) (1978).

59. Pascanu, R., Mikolov, T. & Bengio, Y. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 1310–1318 (2013).
60. Crank, J. & Nicolson, P. A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type. *Math. Proc. Camb. Philos. Soc.* **43**, 50–67, DOI: [10.1017/S0305004100023197](https://doi.org/10.1017/S0305004100023197) (1947).
61. Geselowitz, D. B. On bioelectric potentials in an inhomogeneous volume conductor. *Biophys. J.* **7**, 1–11, DOI: [10.1016/S0006-3495\(67\)86573-3](https://doi.org/10.1016/S0006-3495(67)86573-3) (1967).
62. Deno, D. C. *et al.* High-resolution, live, directional mapping. *Hear. Rhythm.* **17**, 1621–1628 (2020).
63. Plank, G. *et al.* The opencarp simulation environment for cardiac electrophysiology. *Comput. methods Programs Biomed.* **208**, 106223 (2021).
64. Thomée, V. *Galerkin Finite Element Methods for Parabolic Problems* (Springer, Berlin, 2006), 2nd edn.

Appendix A: Cotangent-weight finite-element discretisation

The cotangent weight stiffness matrix (9) used in JAX-EP is derived from the Galerkin weak formulation of the surface monodomain equation (1). Multiplying equation (1) by a test function $\phi \in H^1(\Gamma)$ and integrating over Γ , with integration by parts on the closed manifold eliminating the boundary term:

$$\int_{\Gamma} \beta C_m \frac{\partial V}{\partial t} \phi \, d\Gamma = - \int_{\Gamma} \nabla_{\Gamma} V \cdot \mathbf{D} \nabla_{\Gamma} \phi \, d\Gamma - \int_{\Gamma} \beta I_{\text{ion}} \phi \, d\Gamma + \int_{\Gamma} \beta I_{\text{stim}} \phi \, d\Gamma, \quad (26)$$

consistent with the sign convention of equation (1), in which I_{ion} enters with a negative sign (outward current reduces voltage). Introducing piecewise-linear basis functions $\{\varphi_i\}_{i=1}^{N_p}$ with $\varphi_i(\mathbf{x}_j) = \delta_{ij}$ and approximating $V \approx \sum_i V_i \varphi_i$, the global stiffness and mass matrix entries are:

$$K_{ij} = \int_{\Gamma} \nabla_{\Gamma} \varphi_i \cdot \mathbf{D} \nabla_{\Gamma} \varphi_j \, d\Gamma, \quad (27)$$

$$M_{ij} = \int_{\Gamma} \beta C_m \varphi_i \varphi_j \, d\Gamma, \quad (28)$$

where β is retained inside the mass matrix consistent with its definition in equation (1) and is not absorbed elsewhere in the implementation.

For the isotropic case $\mathbf{D} = d \mathbf{P}_{\Gamma}$ and piecewise-linear elements, the gradient of each basis function is constant within each triangle. Evaluating the elemental stiffness integrals and assembling the contributions from the two triangles sharing edge (i, j) yields the classical cotangent formula⁵⁴:

$$K_{ij} = -\frac{d}{2} (\cot \alpha_{ij}^{(1)} + \cot \alpha_{ij}^{(2)}) = -d w_{ij}, \quad i \neq j, \quad (29)$$

where $\alpha_{ij}^{(1)}$ and $\alpha_{ij}^{(2)}$ are the angles in the two triangles opposite edge (i, j) , and $K_{ii} = \sum_{j \in \mathcal{N}(i)} d w_{ij}$ by the row-sum condition. The anisotropic conductivity tensor (4) introduces an edge-dependent effective conductivity d_{ij} (see equation (10)), yielding the anisotropic stiffness matrix $\mathbf{K}(\boldsymbol{\theta})$ used throughout JAX-EP.

For each triangular element T , the local consistent mass matrix entries are $M_{ij}^T = \beta C_m |T|/12$ for $i \neq j$ and $M_{ii}^T = \beta C_m |T|/6$, which upon assembly give the global consistent mass matrix. JAX-EP instead adopts the lumped-mass approximation, concentrating the nodal mass at the diagonal:

$$M_{ii} = \beta C_m A_i, \quad A_i = \frac{1}{3} \sum_{T \ni i} |T|, \quad (30)$$

where A_i is the barycentric nodal area (one-third of the sum of areas of all triangles incident to node i). Lumping preserves the $\mathcal{O}(h^2)$ spatial accuracy of the scheme on quasi-uniform meshes⁶⁴, eliminates the global mass-matrix solve at each timestep, and is compatible with JAX's static memory model for `jax.lax.scan`. The resulting semi-discrete system is:

$$\mathbf{M} \dot{\mathbf{V}} = -\mathbf{K}(\boldsymbol{\theta}) \mathbf{V} - \mathbf{f}_{\text{ion}}(\mathbf{V}, \mathbf{h}; \boldsymbol{\theta}_{\text{ion}}) + \mathbf{f}_{\text{stim}}(t), \quad (31)$$

where $\mathbf{f}_{\text{ion}}$ is the assembled ionic current vector entering with a negative sign, consistent with the PDE convention of equation (1), and $\mathbf{f}_{\text{stim}}$ is the stimulus current vector, which is advanced in time by the Strang-split Crank Nicolson scheme described in the Time integration section above.

Appendix B: Clinical-Type EGM Distortion and Parameter Recovery




Figure 11. Clinical-type EGM distortion envelope for Set 1 across levels L1–L3. Shaded bands span the signal range; individual traces and mean overlaid on clean ground truth (black dashed). (A,B) Far-field contamination ($\alpha = 0.08$ – 0.60), centre (C5) and edge (C1) cliques. (C,D) Gaussian noise (SNR = 40–20 dB). (E,F) Combined distortion used in parameter estimation. (G) Time-resolved distortion magnitude per type (C5); distortion is concentrated in the activation window.




Figure 12. Effect of clinical-type EGM distortion on parameter recovery across three parameter sets. **(H)** Parameter MAPE as a function of signal distortion level D ; the shaded region denotes $\text{MAPE} < 1\%$ and the dotted vertical line marks the $D \approx 14\%$ threshold below which all three sets are recovered within this target. **(I)** Per-parameter recovery ratio (recovered / ground truth) averaged across the three parameter sets; τ_{in} , τ_{open} and G_{il} remain within $\pm 2\%$ of ground truth across all distortion levels, whilst τ_{out} and τ_{close} degrade progressively with increasing distortion, consistent with their sensitivity to EGM repolarisation morphology.