Reconstructing Conduction Velocity Maps from Sparse Measurements in Atrial Fibrillation

Alexander J Sharp[1,2], Timothy R Betts[2], Abhirup Banerjee[1]

[1]Department of Engineering Science, University of Oxford, Oxford, UK
[2]Department of Cardiology, Oxford University Hospitals NHS Foundation Trust, Oxford, UK

Abstract

High-resolution mapping of conduction velocity (CV) during atrial fibrillation (AF) provides important insight into arrhythmogenic substrate, but is rarely available in routine clinical practice, where measurements are sparse and spatially incomplete. We propose a statistical framework to reconstruct dense CV maps from limited observations using a principal component analysis (PCA)-based statistical appearance model with posterior inference, trained on charge density mapping data from 49 patients with persistent AF. Sparse sampling was simulated, and full maps were reconstructed using a posterior formulation conditioned on observed values, with performance evaluated using leave-one-out cross-validation. Our findings show that sparse sampling of ~20-30% of the atrial surface is sufficient to achieve useful reconstruction, with model-derived uncertainty closely reflecting reconstruction error. This approach enables estimation of whole-chamber CV from limited data, providing a practical route to recover spatially resolved electrophysiological information for mechanistic analysis and in silico modelling of AF.

1. Introduction

Understanding spatial distributions of conduction velocity (CV) is central to characterising the arrhythmogenic substrate that supports atrial fibrillation (AF). CV reflects the combined effects of structural and functional remodelling, including fibrosis, anisotropy, and ion channel alterations, and plays a key role in re-entry dynamics, disease progression, and ablation outcomes [1-3]. Importantly, regional CV slowing measured during AF has been shown to better correlate with propagation patterns representing potential AF drivers than measurements obtained in sinus rhythm [4].

Despite its clinical relevance, acquiring high-resolution CV maps during AF remains challenging [5]. The irregular and chaotic nature of AF complicates conventional contact mapping, where electrograms are acquired sequentially and spatial coverage is inherently limited. Global chamber mapping approaches can overcome these limitations, but availability is limited [6]. As a result, CV measurements in routine clinical settings are sparse and spatially incomplete.

Statistical appearance models (SAMs) provide a framework to capture population-level variability in spatial CV distributions [7]. However, their use has largely focused on analysing variability or generating synthetic data, with limited investigation into their use for reconstructing dense spatial fields from sparse observations.

Posterior statistical models offer a principled solution to this problem. By conditioning a statistical model on partial observations, posterior formulations compute the conditional distribution of the full field given sparse measurements, yielding both a most probable reconstruction and an associated uncertainty estimate [8].

In this study, we investigate whether sparse sampling of CV, combined with a SAM and a posterior inference framework, can enable accurate reconstruction of full CV maps. We assess how reconstruction accuracy varies with sampling density and model complexity, and evaluate both prediction error and model-derived uncertainty. This approach aims to bridge the gap between limited clinical measurements and the dense CV representations required for mechanistic analysis and *in silico* modelling of AF.

2. Methods

2.1. Data acquisition and preprocessing

Data were obtained from 49 patients with persistent AF undergoing global chamber charge density mapping, as previously described [7]. For each patient, an intrachamber ultrasound-derived left atrial (LA) geometry was reconstructed, with an associated spatial CV map.

CV was computed at each mesh vertex as the median value across multiple wavefronts over a 20-second recording of AF. This duration was selected based on prior work demonstrating spatiotemporal stability of AF propagation patterns over this timescale [9].

To enable statistical analysis across subjects, dense surface correspondence was established using a particle-based approach (ShapeWorks [10]), yielding 576 correspondence points per subject. CV values were assigned to each correspondence point using nearest-neighbour interpolation from the original mesh, resulting in a CV appearance vector $\boldsymbol{s} \in \mathbb{R}^{576}$ for each subject.

2.2. Statistical appearance model

A SAM was constructed using principal component analysis (PCA). The sample mean vector $\boldsymbol{\mu}$ and covariance matrix were computed across subjects, and PCA was applied to obtain a low-dimensional representation of CV variability:

$$\boldsymbol{s} = \boldsymbol{\mu} + \boldsymbol{Q}\boldsymbol{\alpha}$$

where $\boldsymbol{Q}$ contains the principal components and $\boldsymbol{\alpha}$ are latent coefficients describing subject-specific variation along each mode.

2.3. Posterior reconstruction

To reconstruct full CV maps from sparse measurements, a posterior formulation of the SAM was employed [8]. Given a subset of observed CV values $\boldsymbol{s}_g$, the model estimates the most probable full CV distribution conditioned on these observations. The reconstructed CV vector is given by the posterior mean:

$$\boldsymbol{\mu}_c = \boldsymbol{\mu} + \boldsymbol{Q}\left(\boldsymbol{Q}_g^T\boldsymbol{Q}_g + \sigma^2\boldsymbol{I}_n\right)^{-1}\boldsymbol{Q}_g^T(\boldsymbol{s}_g - \boldsymbol{\mu}_g)$$

where $\boldsymbol{Q}_g$ and $\boldsymbol{\mu}_g$ denote the principal components and mean restricted to the observed indices, the parameter σ^2 models deviation of observed values from the model subspace and acts as a regularisation term controlling the trade-off between data fidelity and model prior, and the matrix $\mathbf{I}_n$ denotes the $n \times n$ identity matrix where n is the number of retained principal components.

In addition to the conditioned mean, the model provides a posterior covariance:

$$\mathbf{R}_c = \sigma^2\boldsymbol{Q}(\boldsymbol{Q}_g^T\boldsymbol{Q}_g + \sigma^2\boldsymbol{I}_n)^{-1}\boldsymbol{Q}^T$$

which describes the uncertainty in the prediction. Per-point uncertainty was quantified as the square root of the diagonal elements of $\mathbf{R}_c$, yielding a spatial map of standard deviation.

2.4. Evaluation framework

Model performance was evaluated using a leave-one-out cross-validation framework. For each subject, the model was constructed from the remaining cohort, and sparse observations were sampled from the excluded subject. The posterior formulation was then used to reconstruct the full CV map.

Reconstruction accuracy was quantified using root-mean-square error (RMSE), computed exclusively over unobserved points to reflect true predictive performance. A mean-only baseline was also evaluated, where the predicted CV map was the training set mean.

To assess spatial patterns of performance, per-point errors were aggregated across cross-validation folds and mapped onto the mean atrial geometry. Posterior standard deviation values were computed in the same manner to generate uncertainty maps. The relationship between prediction error and model-derived uncertainty was assessed using Pearson correlation between spatial RMSE and posterior standard deviation.

2.5. Parameter selection

Several parameters influence reconstruction performance, including sampling density, model complexity, and the regularisation term σ^2.

For each subject, sparse observations were simulated by generating a single random ordering of correspondence points, from which nested subsets were defined so that each sampling density included all points from lower densities. Sampling densities ranging from 1% to 90% were evaluated.

Model complexity was controlled by the number of retained principal components. The default model retained components explaining 95% of cumulative variance, consistent with common practice in statistical modelling. To assess the effect of model complexity, additional models capturing 80% and 99% of variance were also evaluated.

The parameter σ^2 controls the balance between adherence to observed values and reliance on the statistical model, acting as a regularisation term in the posterior formulation. For each outer cross-validation fold, σ^2 was selected using an inner 5-fold cross-validation on the training set. A logarithmically spaced grid of candidate values (10^{-6} to 10^{4}) was evaluated, and the value minimising reconstruction RMSE on held-out inner folds was selected.

3. Results

3.1. Reconstruction accuracy

Reconstruction performance improved consistently with increasing sampling density for all models (Figure 1). RMSE decreased most steeply at low sampling densities, with the largest gains observed between 1% and ~20-30% surface coverage. Beyond this range,

performance plateaued, with only marginal improvements at higher sampling densities. Across all conditions, posterior reconstructions consistently outperformed the mean-only baseline, demonstrating the value of incorporating sparse observations into the statistical model.

Model complexity influenced reconstruction behaviour. The simplest (80% variance) model performed best at low sampling densities, whereas the most complex (99% variance) model performed best at higher sampling densities. The 95% variance model provided a balance between these extremes, achieving strong performance across all sampling densities.

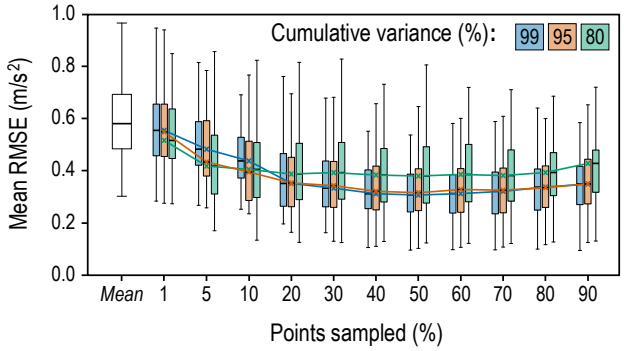

Figure 1. Reconstruction accuracy as a function of sampling density and model complexity. RMSE: root mean square error.

A qualitative example for the 95% variance model (Figure 2) illustrates these trends, demonstrating progressive recovery of spatial detail with increasing sampling density.

3.2. Spatial distributions of reconstruction error and uncertainty

Spatial distributions of reconstruction error and posterior uncertainty are shown in Figure 3 for the 95% variance model at 30% sampling density. Regions of consistently higher RMSE were observed across subjects, particularly around the left atrial appendage, pulmonary veins, and mitral valve annulus.

Posterior standard deviation maps exhibited similar spatial patterns. A strong positive correlation was observed between per-point RMSE and posterior standard deviation (Pearson $r = 0.89$, $p < 0.001$), indicating that regions of higher prediction error correspond to higher model uncertainty.

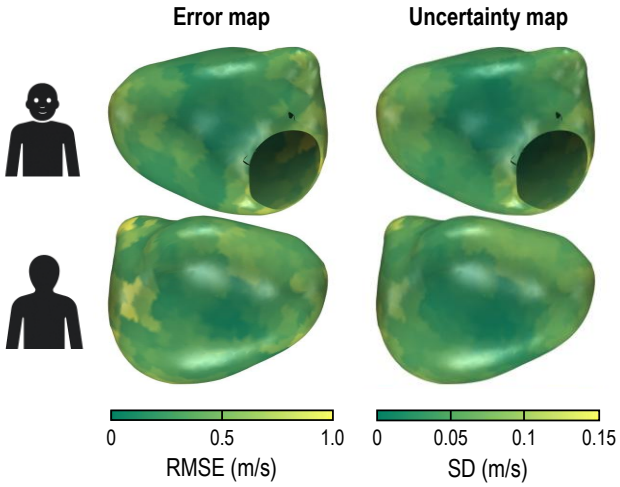

Figure 3. Spatial distributions of reconstruction error and posterior uncertainty. RMSE: root mean square error; SD: standard deviation.

4. Discussion and conclusion

4.1. Sparse sampling and model complexity

A key finding of this study is that most gains in reconstruction performance are achieved using a relatively small fraction of observed surface points. The early plateau suggests that CV distributions exhibit consistent spatial structure, such that sparse observations are sufficient to constrain the underlying pattern and enable reconstruction of the full field, with diminishing benefit from additional sampling.

The 95% variance model provides a practical balance,

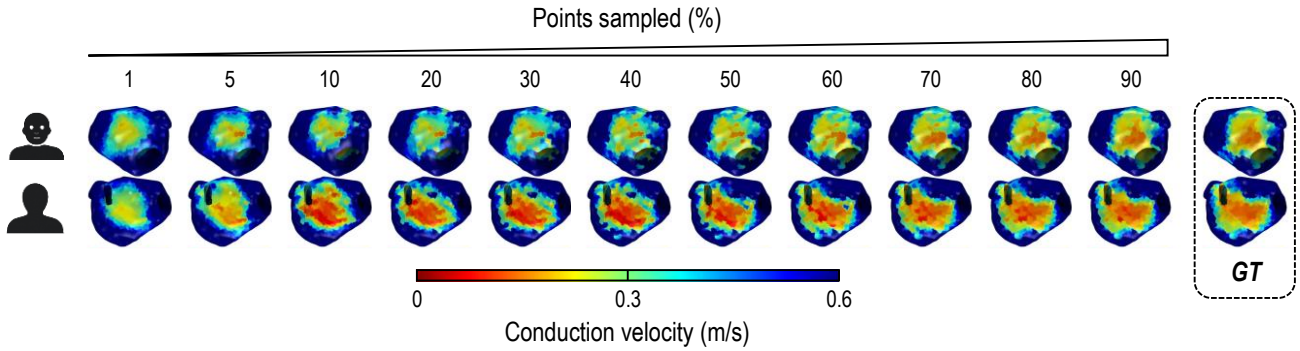

Figure 2. Qualitative reconstruction of CV maps from sparse observations. GT: ground truth.

retaining sufficient variability to support accurate reconstruction while remaining robust to limited sampling.

4.2. Accuracy, spatial variation, and uncertainty

Reconstruction error remained at approximately 0.3 m/s, indicating that while the model captures dominant patterns of variation, it is less able to recover fine-scale, case-specific details. This likely reflects intrinsic beat-to-beat variability in CV during AF, which limits the precision with which spatial CV patterns can be reconstructed. Importantly, the observed error is small relative to the magnitude of this variability (0.40-0.72 m/s) [11], indicating that the model operates within physiologically meaningful limits.

Reconstruction performance also varied spatially. Regions around the pulmonary veins, left atrial appendage, and mitral valve annulus exhibited consistently higher errors, likely reflecting increased anatomical variability and, in the case of the pulmonary veins and mitral valve, non-conducting boundaries that make local CV more difficult to infer.

The strong correlation between reconstruction error and posterior standard deviation demonstrates that the model provides meaningful uncertainty estimates. Regions of higher uncertainty closely align with areas of increased prediction error, supporting the interpretation of posterior covariance as a measure of confidence.

4.3. Clinical implications

The central motivation for this work is whether a model trained on high-resolution global mapping data could be used to infer whole-chamber CV patterns from sparse clinical measurements. The results suggest that this is feasible, with reconstruction performance improving substantially even with limited observations.

However, several important considerations remain. In this study, sampling was random and does not reflect clinical acquisition strategies. In practice, targeted sampling of specific regions may further improve reconstruction by more effectively constraining the posterior solution and represents an important direction for future work. Additionally, no conventional contact mapping data were used, and translation to such settings will require validation in real-world acquisition scenarios.

Despite these limitations, this work represents an important first step towards integrating statistical models with sparse clinical measurements. If validated, such an approach could enable reconstruction of dense, spatially resolved CV maps from limited contact data, bridging the gap between high-resolution mapping techniques and routine clinical practice.

Acknowledgments

This work was supported by a Heart Research UK Novel and Emerging Technologies grant (NET25-100009). AB is supported by a Royal Society University Research Fellowship (URF\R1\221314). We thank the DISCOVER investigators and centres for their data contributions and the ShapeWorks team for developing of the software used in parts of the analysis.

References

[1] S. Honarbakhsh et al., “Structural remodeling and conduction velocity dynamics in the human left atrium: Relationship with reentrant mechanisms sustaining atrial fibrillation,” *Heart Rhythm*, vol. 16, pp. 18-25, 2019.

[2] C. H. Roney et al., “Challenges associated with interpreting mechanisms of AF,” *Arrhythm Electrophysiol Rev*, vol. 8, pp. 273-284, 2019.

[3] E. Invers-Rubio et al., “Regional conduction velocities determined by non-invasive mapping are associated with arrhythmia-free survival after atrial fibrillation ablation,” *Heart Rhythm*, vol. 21, pp. 1570-1580, 2024.

[4] A. J. Sharp et al., “Identifying extra pulmonary vein targets for persistent atrial fibrillation ablation: bridging advanced and conventional mapping techniques,” *Europace*, vol. 27, no. euaf048, 2025.

[5] S. Coveney et al., “Atrial conduction velocity mapping: clinical tools, algorithms and approaches for understanding the arrhythmogenic substrate,” *Med Biol Eng Comput*, vol. 60, pp. 2463-2478, 2022

[6] T. R. Betts et al., “Treatment of pathophysiologic propagation outside of the pulmonary veins in retreatment of atrial fibrillation patients: RECOVER AF study,” *Europace*, vol. 25, no. euad097, 2023

[7] A. J. Sharp et al., “Statistical shape and appearance model of left atrial geometry and conduction velocity in atrial fibrillation,” *Computing in Cardiology*, vol. 52, 2025.

[8] T. Albrecht et al., “Posterior shape models,” *Med Image Anal*, vol. 17, pp. 959-973, 2013.

[9] M. T. B. Pope et al., “Spatial and temporal variability of rotational, focal, and irregular activity: practical implications for mapping of atrial fibrillation,” *J Cardiovasc Electrophysiol*, vol. 32, pp. 2393-2403, 2021.

[10] J. Cates et al., “Shapeworks: particle-based shape correspondence and visualization software,” *Statistical Shape and Deformation Analysis*, pp. 257-298, 2017.

[11] K. T. Konings et al., “High-density mapping of electrically induced atrial fibrillation in humans,” *Circulation*, vol 89, pp. 1665-1680, 1994.

Address for correspondence:

Alexander James Sharp
Institute of Biomedical Engineering, Old Road Campus Research Building, Oxford, OX3 7DQ, United Kingdom
alexander.sharp@eng.ox.ac.uk