



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/145442/>

Version: Accepted Version

Article:

Coveney, S., Corrado, C., Roney, C. et al. (2020) Probabilistic interpolation of uncertain local activation times on human atrial manifolds. *IEEE Transactions on Biomedical Engineering*, 67 (1). pp. 99-109. ISSN: 0018-9294

<https://doi.org/10.1109/TBME.2019.2908486>

© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/ republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works. Reproduced in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Probabilistic Interpolation of Uncertain Local Activation Times on Human Atrial Manifolds

Sam Coveney, Cesare Corrado, Caroline H Roney, Richard D Wilkinson, Jeremy E Oakley, Finn Lindgren, Steven E Williams, Mark D O'Neill, Steven A Niederer, and Richard H Clayton

Abstract—Objective: Local activation time (LAT) mapping of the atria is important for targeted treatment of atrial arrhythmias, but current methods do not interpolate on the atrial manifold and neglect uncertainties associated with LAT observations. In this paper we describe novel methods to (i) quantify uncertainties in LAT arising from bipolar electrogram analysis and assignment of electrode recordings to the anatomical mesh, (ii) interpolate uncertain LAT measurements directly on left atrial manifolds to obtain complete probabilistic activation maps, and (iii) interpolate LAT jointly across both the manifold and different S1-S2 pacing protocols. **Methods:** A modified center of mass approach was used to process bipolar electrograms, yielding a LAT estimate and error distribution from the electrogram morphology. An error distribution for assigning measurements to the anatomical mesh was estimated. Probabilistic LAT maps were produced by interpolating on a left atrial manifold using Gaussian Markov random fields, taking into account observation errors and characterizing LAT predictions by their mean and standard deviation. This approach was extended to interpolate across S1-S2 pacing protocols. **Results:** We evaluated our approach using recordings from three patients undergoing atrial ablation. Cross-validation showed consistent and accurate prediction of LAT observations both at different locations on the left atrium and for different S1-S2 intervals. **Significance:** Interpolation of scalar and vector fields across anatomical structures from point measurements is a challenging problem in biomedical engineering, compounded by uncertainties in measurements and meshes. New methods and approaches are required, and in this paper we have demonstrated an effective method for probabilistic interpolation of uncertain LAT.

Index Terms—Bipolar Electrogram, Local Activation Time, Atrial Fibrillation, Uncertainty, Probabilistic Interpolation

I. INTRODUCTION

LOCAL activation time (LAT) measurements, which indicate the onset of cardiac tissue depolarization, are part of routine clinical care for the diagnosis and treatment of atrial arrhythmias using radio frequency catheter ablation. LAT is typically determined from unipolar or bipolar electrograms recorded at different locations using an electro-anatomic mapping system, and this information is used to build a map of LAT over the atrial surface [1], [2].

This work was funded by the UK Engineering and Physical Sciences Research Council through grants EP/P010741/1 and EP/P01268X/1.

Sam Coveney and Richard Clayton are with the Insigneo Institute for in-silico Medicine and Department of Computer Science at the University of Sheffield, United Kingdom. (s.coveney@sheffield.ac.uk)

Cesare Corrado, Caroline Roney, Steven Williams, Mark O'Neill and Steven A Niederer are with the Division of Imaging Sciences and Biomedical Engineering at King's College London, United Kingdom.

Richard Wilkinson and Jeremy Oakley are with the School of Mathematics and Statistics at the University of Sheffield, United Kingdom

Finn Lindgren is with the School of Mathematics at the University of Edinburgh, United Kingdom

Electrograms are recorded using a multipolar catheter which records electrical activity on the endocardial surface at one or more locations, sometimes for multiple pacing protocols and catheter positions. LAT maps can be used to calculate conduction velocity maps [3], and different pacing protocols allow the dynamic response of the electro-anatomical substrate to be inferred [4]. Although advances in catheter design have increased the number of recordings that can be made simultaneously, longer procedure times pose a risk to the patient and a cost for the health care system, which places limits on the spatial density of the mapping observations, as well as how many pacing protocols and pacing locations can be used. Data recording is therefore often sparse.

To create a LAT map, observations must be interpolated. Current interpolation techniques are predominantly linear, although a Radial Basis Function method [5] and a spline method [6] have been proposed. However, current methods do not interpolate directly on the atrial manifold, and to our knowledge no methods have been proposed to interpolate between pacing protocols. Furthermore, current methods do not properly account for observation uncertainty, nor for interpolation uncertainty. Estimating the uncertainties in conduction velocities, and therefore uncertainties in LAT maps, has been identified as an important challenge [7].

There are many potential sources of uncertainty that could influence the best estimate of LAT at a particular location. These include uncertainties in the assignment of LAT to electrograms, especially those with a complex morphology, and mismatch between catheter electrode locations and anatomical mesh for a particular recording. There have been few attempts to assign confidences to LAT observations, with some exceptions focusing on electrograms [8], [9] and uncertainty based on catheter to surface distance [10]. Besides observation uncertainty there is additional uncertainty arising from the interpolation procedure, which should be quantified so that all LAT predictions are assigned a confidence.

In this paper, we demonstrate a novel technique that accounts for these different sources of uncertainty. Our approach involves probabilistic global interpolation of heterogeneously uncertain LATs over endocardial surfaces using Gaussian Markov Random Fields (GMRFs) in a Bayesian hierarchical framework. We assign uncertainties to LAT arising from bipolar electrogram morphology by extending an existing LAT assignment method. We then estimate the uncertainties resulting from projection of measurement locations onto the left atrial surface. Using the GMRF method we construct a probabilistic LAT map over the entire left atrial surface, which

accounts for heterogeneous uncertainties in LAT observations and quantifies uncertainty on LAT predictions. We extend this method to also interpolate between different pacing protocols. We apply our methods to recordings from 3 clinical cases, and demonstrate good performance using cross-validation.

II. METHODS

A. Electrogram recordings

The recordings analyzed in this paper were part of a study involving patients with paroxysmal atrial fibrillation and undergoing first-time atrial fibrillation ablation, which is detailed elsewhere [4]. We analyzed recordings from three patients. Ethical approval was granted by the UK National Research Ethics Service (10/H0802/77) and all patients gave written informed consent for inclusion in the study. The research conformed to the principles described in the Declaration of Helsinki. None of the patients had ischemic heart disease, cardiac surgery or structural heart disease. Following femoral access and trans-septal puncture, two 8.5 French SR0 long sheaths and a PentaRay mapping catheter (Biosense Webster, CA, 1mm electrode size, 4-4-4mm spacing) were advanced into the left atrium. Decapole (St Jude Medical, MN) and pentapole (Bard Electrophysiology, MA) catheters were positioned in the coronary sinus (CS) and high right atrium (HRA), respectively. Left atrial electrophysiology was investigated using an S1-S2 programmed pacing protocol [11] consisting of a 2-beat drivetrain with a cycle length S1=470 ms, followed by a single premature extra stimulus S2<S1. The pacing protocol was delivered using a custom-built, institutionally-approved stimulator that reduced the S1-S2 interval continuously and without operator interference from 343 ms to 200 ms (or loss of capture) by reducing the current S1-S2 interval by 2% of its value (rounded to nearest ms). All pacing stimuli were delivered at a voltage of at least twice threshold and with a pulse width of 2 ms, from either the coronary sinus or high right atrium.

Bipolar electrograms (EGMs) were recorded at a sampling rate of 4 kHz using the 10 electrodes of the PentaRay catheter for each S1-S2 interval. The catheter was then moved to another location on the left atrial endocardium, and the process repeated for up to 15 locations (i.e. maximally 150 traces per S1-S2 interval). For each S1-S2 interval, we discarded any EGM traces not containing a discernible activation complex. The maximum number of discernible traces across all patient cases for any S1-S2 interval was 99.

B. Electrogram processing

A theoretical link exists between the maximal downslope of unipolar electrograms and depolarization of cardiac cells [12]. However, there is no theoretical link nor accepted best method for assigning LAT for bipolar electrograms, which are commonly used in clinical practice because they measure only local activity. Several LAT assignment methods for bipolar electrograms (BEGMs) have been reported that are sufficiently accurate for using LAT as fiducial markers for conduction velocity calculations [7]. However, the differences between these methods may be relatively unimportant compared to the

intrinsic uncertainty in determining LAT for bipolar signals. We model the relation between the true local activation time (LAT_{EGM}) and the LAT assignment (t_{obs}) made from an *observed* EGM (derived using a specific method), as:

$$\begin{aligned} t_{obs} &= LAT_{EGM} + \epsilon_{EGM} \\ \epsilon_{EGM} &\sim \mathcal{N}(0, \sigma_{EGM}^2) \end{aligned} \quad (1)$$

where ϵ_{EGM} is an error assumed to have a normal distribution with standard deviation σ_{EGM} for tractability.

We can be more confident of our LAT assignment for signals with clean single deflections than for noisy multiple deflections, so the magnitude of σ_{EGM} should depend on the signal morphology in a robust and consistent way. Furthermore, the assessment of σ_{EGM} should be meaningfully related to the measurement yielding t_{obs} . For these reasons, in this paper we determine t_{obs} from bipolar electrograms using the ‘Center-of-Mass’ (CoM) method [13], [14], because this has been demonstrated to accurately assign LATs to bipolar signals based on comparison with expert assignment of corresponding unipolar electrograms. In the original publications, the beginning and end of the signal were carefully defined, the signal was filtered, differentiated and rectified, and LAT was defined as t_{50} : the time of 50% of the cumulative area under the rectified differentiated signal between the signal beginning and end. We extend this method in order to also give a consistent definition of σ_{EGM} .

We summarize our modified Center-of-Mass method as four stages below, illustrated in Figure 1, and refer the reader to Supplementary Material for a schematic and more precise details for step (b). Where the original method used a 4-state machine algorithm to bracket the EGM activation complex in a very particular way, we use an amplitude threshold instead. This simplification is justified since t_{50} is only a *noisy observation*; increasing the precision of this measurement does not reduce the *intrinsic* uncertainty about how t_{50} differs from LAT_{EGM} . Also, where the original method applied frequency filters to smooth signals, we use either a sliding-window average (step (b)) or a Gaussian Process (step (c)).

(a) **Regularize signal** (Figure 1(a)); subtract the mean and normalize the amplitude of the signal $V(t)$ to remove trends. No modifications to smoothness or morphology are made:

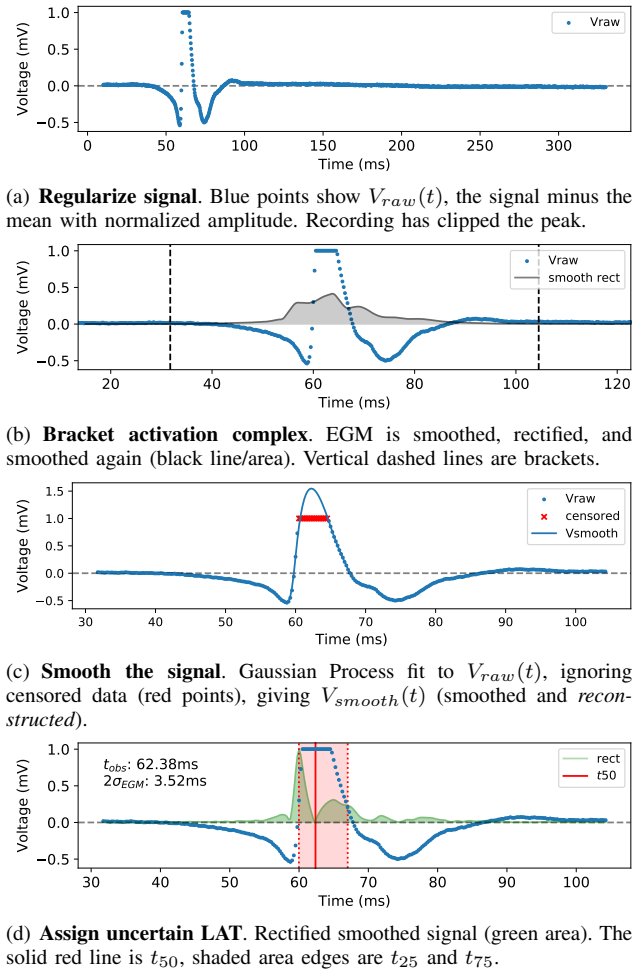
$$V_{raw}(t) = (V(t) - \overline{V(t)}) / \max(|V(t) - \overline{V(t)}|), \quad (2)$$

where $\overline{V(t)}$ is the average of $V(t)$ calculated for a time window of up to 350 ms following the pacing stimulus.

(b) **Bracket activation complex** (Figure 1(b)). We smooth $V_{raw}(t)$, then differentiate (central difference), rectify, and normalize it, and then smooth again to obtain ‘smooth rect’, which we used to bracket the complex by observing where the amplitude falls to 2.5% of the maximum. We discard $V_{raw}(t)$ outside the brackets. If bracketing failed then we did not proceed to (c) and no LAT observation was recorded.

(c) **Smooth signal** (Figure 1(c)). Before taking the rectified difference for area calculation in part (d), the signal must be smoothed. In one of three patients presented here, nearly half of the EGMs were partly ‘clipped’ due to thresholding by the recording equipment. Therefore we smoothed and

Fig. 1. Assignment of Local Activation Time observations and standard deviations to bipolar electrograms, using a modified Center-of-Mass method.



reconstructed $V_{raw}(t)$ to produce $V_{smooth}(t)$ using a Gaussian Process (GP) instead of a simpler smoothing method. We “censored” the clipped data and fitted a GP with Matérn 3/2 kernel (this choice prevents oscillatory artefacts in censored region) with a fixed nugget 0.001 (for smoothing) using the software GPy [15]. Since the clipping was not too severe, this enabled all of the clipped EGMs to be used, and improved the area calculation in (d). The effectiveness of the reconstruction was tested by artificially clipping complete signals; the reconstructed signals were a good match to the original.

(d) **Assign uncertain LAT** (Figure 1(d)). We differentiated, rectified and normalized the smoothed (and reconstructed) signal from (c), and calculated the cumulative area under this signal. For the observation of LAT, we define:

$$\begin{aligned} t_{obs} &= t_{50} \\ \sigma_{EGM} &= (t_{75} - t_{25})/4 \end{aligned} \quad (3)$$

This particular choice (t_{25} and t_{75}) allows plausible features for LAT_{EGM} to be within the $\approx 95\%$ confidence interval implied by the error distribution in Eq. (1). Mostly importantly, this extension of the Centre-of-Mass method is motivated by our aim to have self-consistent uncertainty quantification in EGMs, such that the uncertainty is related to the EGM feature used

for the observation of Local Activation Time. We return to this in the discussion (section IV).

C. Mesh processing

We obtained left atrial anatomies during the clinical procedures using an electro-anatomical mapping system (St Jude Velocity). To obtain a smoothed domain suitable for the GMRF method described in section II-D, we proceeded as follows. First, we smoothed each surface by applying a Poisson filter implemented in the MESHLAB library, [16]. Next, we determined the center location of the ostium of each pulmonary vein (PV) as the point presenting an abrupt increase of the PV diameters; we then truncated each PV at a distance of 3mm from the centre of the PV ostium. Finally, we generated a uniform triangulation with a characteristic edge length $h = 2.3$ mm. This procedure used the tools provided in the VMTK library, [17]. Figure 2 shows a smoothed anatomy obtained by applying this procedure.




Fig. 2. Smooth Left Atrial mesh, with the original mesh inset (small). MV: mitral valve, L(R)PV: left (right) pulmonary veins, LAA: left atrial appendage. The smoothed mesh ($\sim 10^4$ vertices) is suitable for the GMRF method.

The coordinates at which electrograms were recorded (barycentres of the bipolar electrodes on the catheter) did not lie perfectly on the original (non-smoothed) left atrial mesh. Thus there was additional uncertainty in the location of LAT observations on the mesh, which we modeled as:

$$\begin{aligned} LAT_{EGM} &= LAT_{mesh} + \epsilon_{pos} \\ \epsilon_{pos} &\sim \mathcal{N}(0, \sigma_{pos}^2) \end{aligned} \quad (4)$$

In order to register the electrode positions with the mesh, we projected each measurement coordinate \mathbf{x} onto the nearest coordinate on the *smoothed* mesh \mathbf{x}^* using a user-contributed MATLAB routine called `point2trimesh` [18]. Figure 3 shows this projection.

In order to meaningfully relate the uncertainty quantification to the measurement, we assume that larger projection distance $d = |\mathbf{x} - \mathbf{x}^*|$ ought to correspond to larger σ_{pos} [10]. However, since the activation wavefront travels over the manifold, we need to relate uncertainty in time to a distance on the manifold. Therefore, the signal measured by the electrode was assumed to lie within an arc of length $d \times \theta$ centered on \mathbf{x}^* (see inset of Figure 3), such that there is an associated LAT difference Δt between the mesh and the catheter electrode given by:

$$\Delta t = \frac{d \times \theta}{\text{conduction velocity}} \quad (5)$$

where we assume a conduction velocity of 350mm/s and fairly broad arc angle $\theta = 2\pi/10$ in order not to underestimate σ_{pos} . Assuming 95% confidence that LAT_{EGM} is within $\pm\Delta t/2$ of LAT_{mesh} , we can write:

$$2\sigma_d = \Delta t/2. \quad (6)$$

For the GMRF method discussed below, these measurements were then assigned to the nearest mesh vertex \mathbf{x}^{**} , but since this distance was small and usually already included within the arc length defined above we did not include additional uncertainty from this assignment.

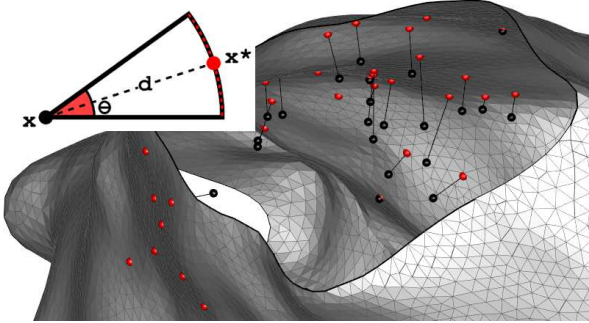


Fig. 3. Projection from measurement location x (black) to the nearest coordinate on the manifold x^* (red). The inset shows the arc $d \times \theta$ centered on x^* used to estimate the uncertainty arising from the projection.

Upon user inspection, the nearest-point projection resulted in a few incorrect assignments on the mesh, mostly for observations with large projection distances. This happened when measurements were from parts of the original mesh that had been removed, or when a more accurate projection method was needed to distinguish between two different mesh locations of similar projection distance. Some immediately neighboring electrode locations had completely inconsistent (but clear) electrograms, demonstrating a problem with identifying the catheter position in the procedure. We labeled these inaccurate coordinates and ignored any corresponding LAT observations in Section II-D (see Table I).

D. Gaussian Markov Random Field method

A popular technique for spatial interpolation is Kriging (e.g. [19]), or (almost) equivalently Gaussian process regression [20]. To use this method we need to specify a function to give the covariance between LATs at any two vertices on the mesh. A ‘standard’ Gaussian process regression approach would define the covariance function to be a decreasing function of Euclidean distance between input (vertex) locations. This would not be appropriate here because Euclidean distance may not reflect the actual distance traveled by an electrical wave over the curved left atrial manifold.

To account for the mesh surface in our interpolation, we use Gaussian Markov random fields, proposed in [21], which are equivalent to Gaussian processes under some circumstances, but which allow for more complex domains and more efficient computation. We give an outline of the method here, and refer the reader to [21] for the full details. We implement this method in R [22], using the package R-INLA [23], [24].

This approach works by representing the Gaussian field on the atrial domain, denoted $g(\cdot)$, by a decomposition defined on a triangular domain/mesh:

$$g(x) = \sum_{k=1}^n w_k \psi_k(x), \quad (7)$$

where each basis function ψ_k is defined to be piecewise linear for any location \mathbf{x} in the domain, taking the value 1 at the k -th vertex in the mesh, and 0 at any other vertex. The weight w_k gives the value of g at the k -th vertex. Uncertainty about the field g is then described by uncertainty in the weights w_1, \dots, w_n , which are assumed to have a multivariate normal distribution.

We have noisy observations of a subset of the weights, i.e., LAT observations at a subset of mesh vertices. Conditional on various parameters (to be described shortly), it is then straightforward to estimate g at the remaining vertices conditional upon the noisy observations using standard properties of multivariate normal distributions. We now have to specify a mean vector and variance matrix for the weights. We use a constant prior expectation β for all weights. For the variance, in [21] it is shown how to construct a (sparse) *precision* matrix Q (the inverse of the variance matrix) for the weights, which is parameterised in terms of an unknown length scale parameter l , a smoothness parameter α (which we fix *a priori*), and a standard deviation parameter ω . The construction ensures a valid (positive-definite) precision matrix, with covariances between weights a function of distances along the manifold. Note that under this construction, unlike a ‘standard’ Gaussian process model, we do not have a closed form expression for the *covariance* between LAT at two different locations. By representing the unknown LAT field using (7), with the precision matrix Q , we can interpret this procedure as using a finite element method to obtain an approximate solution to a stochastic partial differential equation, whose exact solution is a Gaussian process with a Matérn covariance function with length scale l , smoothness parameter $\nu = \alpha - d/2$ (where d is dimension i.e. 2), and variance parameter ω on the atrial domain [21].

Our model linking the observations of LAT to the LAT field on the manifold (LAT_{mesh}) is defined by equations (1) and (4). Combining the errors from the electrograms and projection (assuming independence) gives:

$$\begin{aligned} t_{50} &= LAT_{mesh} + \epsilon \\ \epsilon &\sim \mathcal{N}(0, \sigma^2), \quad \sigma^2 = \sigma_{pos}^2 + \sigma_{EGM}^2 \end{aligned} \quad (8)$$

If we define y_1, \dots, y_N to be the t_{50} values with associated standard deviations $\sigma_1, \dots, \sigma_N$, then in summary, we use the following hierarchical model:

$$y_i | \mu_i, \sigma_i, \tau \sim \mathcal{N}(y_i | \mu_i, \sigma_i^2 / \tau), \quad i = 1, \dots, N \quad (9)$$

$$\mu_i = \beta + g(x_i) \quad (10)$$

$$g(\cdot) | l, \alpha, \omega \sim GF(0, Q^{-1}) \quad (11)$$

The linear predictor μ is the true LAT on the atrial manifold (i.e. $LAT \equiv LAT_{mesh}$) that we want to calculate a distribution for. The field $g(\cdot)$ has a Gaussian field distribution, with precision matrix Q constructed given the mesh vertex

locations and parameters l, α, ω . We fix $\alpha = 2$, corresponding to *Markovian* Matérn fields with smoothness $\nu = 1$ and length-scale parameter l . The parameter β is introduced as an intercept, an average for the entire LAT field, so that $g(\cdot)$ can be assumed to be zero mean). The parameter τ is a global precision which acts to scale the observation precision σ_i^{-2} . Although τ could be fixed to 1, we estimate its value during model fitting, so that our uncertainty assignments can be scaled to achieve a better model fit and for generality.

We use two modeling regimes: (i) LAT spatial interpolation for each S1-S2 interval independently, described in equations (9)-(11) and denoted *SpatialOnly*; and (ii) LAT interpolation over space and S1-S2 interval, denoted *SpatialPacing*. Regime (ii) involves interpolation across different pacing protocols so that the underlying random field becomes a function of both position and S1-S2 interval. We extend the model used in regime (i) to a separable space-time model, the autoregressive process of order 1, AR(1), as described in [25], where S1-S2 takes the place of time in order that the spatial fields at different S1-S2 values are correlated. We can view this as replacing equation (11) by

$$g(\cdot, \cdot) | l, \alpha, \omega, \rho \sim GF(0, Q_T^{-1} \otimes Q^{-1}) \quad (12)$$

where Q_T is the precision matrix of an AR(1) process with correlation parameter ρ ; see [23], [25] for more detail. It is important to note that this method is not simply interpolating between *observation* data from different S1-S2 values, but is correlating the underlying LAT field obtained for different S1-S2 intervals. This is, as far as we are aware, the first instance of interpolating across different pacing protocols.

Fitting the model refers to the training of the parameters, and predicting refers to obtaining statistics for the linear predictor μ , which is synonymous with $LAT \equiv LAT_{mesh}$. For numerical reasons we scaled LAT units into deci-seconds, and unscaled the results afterwards. Distances were measured in millimeters. The uncertain parameters are l, τ, β, ω , and for regime (ii) ρ . Penalized complexity priors [26] were chosen for l, τ, ω and ρ : a 1% probability that l was less than 100; a 1% probability that $\sqrt{\tau}$ is above 2; a 1% probability that ω was above 10; a 80% probability of ρ being above 0.9. A vague uniform prior was chosen for β . To marginalize over the uncertain parameters l, τ, β, ω , we use the integrated nested Laplace approximation [27] implemented in R-INLA. Our procedure is outlined in Figure 4.

E. Cross Validation

We predicted a LAT map for each S1-S2 interval considered. For the *SpatialOnly* model the data y_i were observations from only the considered S1-S2 interval. For the *SpatialPacing* model the data y_i were observations from ‘triplets’ of S1-S2 intervals; the central interval being the considered one.

For model validation, we used k -fold stratified cross-validation implemented in the R package `caret` [28], such that validation folds were optimized for average y_i . For the *SpatialOnly* model, the data (for each pacing) is split into k folds (subsets), and the model is trained on all data except one fold and used to predict the left-out data. This is repeated for

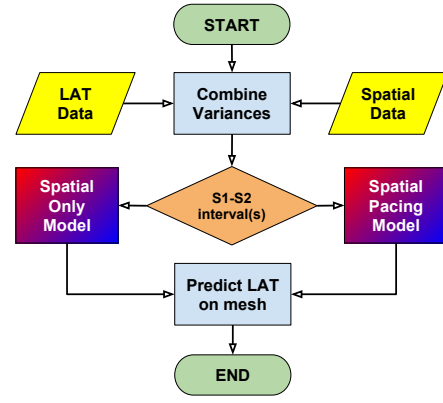


Fig. 4. Schematic of our LAT interpolation methodology. Various methods for assigning LAT to electrograms (LAT data) and assigning measurements to the (smoothed) anatomical mesh (spatial data), with associated uncertainties, could be used in place of our methods. Data from different S1-S2 intervals can be treated separately (*SpatialOnly*) or jointly (*SpatialPacing*).

all folds. We investigated the *SpatialOnly* model for $k = 10$, $k = 4$, and $k = 2$ in order to determine how prediction was affected by the available number of mapping points.

To validate the *SpatialPacing* model, we apply the exact same $k = 10$ folds used for the *SpatialOnly* model to the central interval in the triplet of S1-S2 intervals, such that the validation observations are the same between both modeling regimes e.g. for S1-S2 intervals 329 ms, 336 ms, 343 ms, we perform cross validation for interval 336 ms using the same $k = 10$ folds used for the *SpatialOnly* model for interval 336 ms. This allows a direct comparison of the two models. Also with this cross-validation regime, we (i) vary the spacing between the triplet of S1-S2 intervals; (ii) combine data from triplets of S1-S2 intervals into the *SpatialOnly* model.

We present our cross-validation results as prediction versus observation plots, and summarize them using the *normalized* Root Mean Squared Error (RMSE) for all predictions versus observations, calculated as:

$$\text{RMSE} = \frac{100}{\text{range}} \sqrt{\frac{1}{N} \sum_{i=1}^N (\mathbf{E}[\mu_i] - y_i)^2} \quad (13)$$

where $\text{range} = \max(y) - \min(y)$ for data y_i only for the considered S1-S2 interval. A perfect match of predictions $\mathbf{E}[\mu]$ and noisy observations y_i is not expected (or wanted), and RMSE scores can be highly influenced by data with large uncertainties for which the t_{50} is far from LAT_{mesh} . Therefore we also calculate the Independent Standard Error (ISE):

$$\text{ISE}_i = (\mathbf{E}[\mu_i] - y_i) / \sqrt{\mathbf{V}[\mu_i] + \sigma_i^2} \quad (14)$$

Good predictions will have $\approx 95\%$ of ISEs lie within the interval ± 2 . Both the RMSE and ISE scores, for each S1-S2 interval value, are shown in the legends in Figures 7 - 8.

F. Comparison with alternative interpolation method

Our approach aims to propagate uncertainties in LAT observation to LAT predictions. A direct comparison with other interpolation methods is difficult for two reasons; there is

no ground truth LAT map, and currently used interpolation methods do not incorporate uncertainties. We therefore compared the *SpatialOnly* model against a 3D Gaussian process with heteroscedastic observation variance because this method allowed us to incorporate uncertain LAT observations. Details are provided in Supplementary Material. Although the 3D Gaussian process application itself is somewhat novel, it is comparable with other LAT interpolation in that interpolation is not performed directly on the manifold itself. GP interpolation is very similar to Radial Basis Function interpolation [5] (we used a GP with RBF kernel) and Spline interpolation [6] (splines are special cases of Gaussian processes).

III. RESULTS

In what follows, we summarize the estimated LAT map by the posterior mean and standard deviation of the linear predictor $\mu \equiv LAT$ at each mesh vertex. The GMRF model training and prediction was performed on a HP Elitedesk desktop computer with Intel Core i5-7500 CPU (3.40GHz \times 4) and 8 GB RAM (running Ubuntu 18.04.1 LTS). For the *SpatialOnly* model and *SpatialPacing* model respectively, optimizing with default parameter settings took around 60 s and 600 s, but with well-chosen starting values this reduced to 10 s and 60 s (e.g. parameter values from fitting to all data were used to initialize the optimization during cross validation). To further speed up cross validation we used the empirical Bayes approximation in R-INLA, and made predictions only for mesh points corresponding to the validation data.

We considered S1-S2 intervals between 343 ms and 280 ms, because data coverage became increasingly poor below 280 ms (possibly related to refractory tissue). This also allowed the same parameter priors to be used for all modeling. We present results for three patient cases, for pacing from the coronary sinus (CS). Data summaries are shown in Table I. Patients A and C both had fairly good data coverage overall, but the coverage varied with S1-S2 interval much more for Patient A. Patient B had very low data coverage.

Figures 5 and 6 show $\mathbf{E}[\mu]$ and $\sqrt{\mathbf{V}[\mu]}$ respectively for patient C for S1-S2 interval 336 ms. The LAT map $\mathbf{E}[\mu]$ is typical of our results: the field is smooth, and high precision observations (smaller σ_i) influence and pin the LAT field more than low precision observations (larger σ_i). The heterogeneous observation uncertainty also influences the $\sigma \equiv \sqrt{\mathbf{V}[\mu]}$ map, demonstrating lower prediction certainty in regions of low precision observations. In the lower part of Figure 5, some observations (yellow spheres) appear based on physical considerations of the propagating wave to be poorly assigned to the mesh, and should probably have been identified beforehand and ignored. Nonetheless, the model mostly overlooks these

Patient	Min	Max	Mean	Std	Ignored
A	42	99	71.9	16.3	1
B	6	48	26.6	18.3	2
C	65	72	70.3	2.65	0

TABLE I

SUMMARY OF OBSERVATIONS PER S1-S2 INTERVAL, AND NUMBER OF COORDINATES FOR WHICH ALL OBSERVATIONS WERE IGNORED, FOR THREE PATIENT CASES DISCUSSED BELOW.

problems due to influence from other more precise observations. The GMRF method is also able to smoothly extrapolate beyond the data, although predictions too far beyond the data locations should be treated with caution in general.




Fig. 5. LAT map mean $\mathbf{E}[\mu]$ for patient C for *SpatialOnly* model with observations y_i for S1-S2 interval 336ms. The spheres show LAT observations $y_i \equiv t_{50}$. Sphere sizes increase with observation variance. Higher precision observations pin the LAT map more strongly. There appear to be some observations that were poorly assigned to the mesh (lower figure, yellow spheres), which should ideally be removed followed by refitting.

A. *SpatialOnly* model: single S1-S2 intervals

Figure 7 shows the $k = 10$ cross-validation results for the *SpatialOnly* model, with the RMSE and ISE statistics shown in the legends against the corresponding color for the S1-S2 colorbar. The RMSE statistics for $k = 10, 4, 2$ folds, for both the GMRF model and GP model, are summarized for comparison as averages over S1-S2 intervals in Table II. Bar plots showing the breakdown of these results against S1-S2 interval are given in the Supplementary Material. Figure 7(a) shows good prediction versus observation for all S1-S2 intervals in one patient, whereas Figure 7(c) shows that the prediction appears to get worse for low S1-S2 values in another patient. The RMSE scores in Figure 7(a) are approximately half those in Figure 7(c), and the goodness of fit is further suggested by better ISE coverage in the ± 2 range for the former.

The top three rows of Table II show that when the size of the cross-validation folds increase, the RMSE scores increase but the effect is very minor, demonstrating good robustness of the spatial interpolation method to decreasing amounts of data. However, for Patient B (the least data) both figure 7(b) and Table II demonstrate that prediction is much more difficult with sparse spatial coverage (although the ISE plots in Figure 7(b) suggest that the poorer predictions are compensated for by a




Fig. 6. LAT map standard deviation $\sqrt{\mathbf{V}[\mu]}$ for patient C for *SpatialOnly* model with observation standard deviations σ_i . Sphere sizes increase with observation variance. Uncertainty increases with distance from observations, and more precise observations generally reduce prediction uncertainty.

Model	k -folds	Patient A	Patient B	Patient C
GMRF	10	5.98	18.3	10.9
	4	6.26	18.7	10.7
	2	6.64	19.3	12.0
GP	10	6.42	29.8	11.5
	4	6.84	106.	11.9
	2	20.1	116.	13.4


TABLE II

k -FOLD CROSS-VALIDATION RESULTS FOR THE *SpatialOnly* MODEL (GMRF) AND 3D GAUSSIAN PROCESS MODEL (GP): AVERAGE OF THE RMSE SCORES OVER S1-S2 INTERVALS.


higher uncertainty). Interestingly, predictions with $k = 2$ folds for Patients A and C are still much better than predictions for $k = 10$ folds for Patient B, suggests that spatial coverage overall is more important than the absolute number of data points. This is because withholding validation data where the data is already sparse may result in bad extrapolation.

The bottom three rows of Table II show the cross-validation RMSE summaries for the simpler 3D Gaussian process model, where the same folds from the GMRF spatial model cross-validation were used. In all cases, our model performed better than the 3D GP method, although the GP model performed quite well overall when enough data was available. However, the GP model performance decreased much faster as the number of folds decreased, sometimes due to the inability to fit the data set well. However, as shown in the Supplementary Material, the predicted uncertainty in the GP model is not sensible given the physical system that is being modelled. This is because the interpolation can take no account of the manifold, and low uncertainty can be achieved where


Fig. 7. 10-fold cross-validation results for *SpatialOnly* models. Legends show normalized Root Mean Squared Error and percentage of Independent Standard Errors within ± 2 against corresponding S1-S2 interval color.



(a) Patient A: good data coverage (quite variable with S1-S2), with similar performance for all S1-S2 intervals.



(b) Patient B cross-validation. Very low data coverage.



(c) Patient C cross-validation. Good consistent data coverage, better performance for higher S1-S2 intervals.

the distance to data ‘through the atrium’ is small, even though the connecting distance ‘across the atrium’ is large. Conduction velocity calculations (shown in Supplementary Material) demonstrate that the GMRF method outperforms the GP method substantially; the GP conduction velocity field is riddled with problematic artefacts and unphysical prediction of

Model	Pacing data in triplet	A	B	C
Spatial Pacing	nearest	3.72	8.26	5.99
	next nearest	4.01	10.2	6.03
	next next nearest	3.39	7.62	6.73
	next next next nearest	3.59	9.18	5.38
Spatial Only	nearest	4.94	10.1	7.57
	next nearest	5.61	11.9	8.70
	next next nearest	6.35	12.7	9.60
	next next next nearest	7.97	11.8	10.1

TABLE III

10-FOLD CROSS-VALIDATION RESULTS FOR THE *SpatialPacing* MODEL, AND THE *SpatialOnly* MODEL WITH DATA FOR SEVERAL S1-S2 INTERVALS: AVERAGE OF THE RMSE SCORES OVER S1-S2 INTERVALS.

the electrical wave propagation, whereas the GMRF method yields very sensible results due to the interpolation being performed on the manifold itself.

B. *SpatialPacing* model: multiple S1-S2 intervals

Figure 8 shows the $k = 10$ cross-validation results for the *SpatialPacing* (multiple S1-S2 interval) model, where for each S1-S2 interval the exact same folds for the *SpatialOnly* model were used. The top four rows of Table III show the average RMSE results for the *SpatialPacing* model, where the triplet of S1-S2 intervals used in the model were adjusted across the four rows e.g. for predicting LAT for S1-S2 interval 310 ms, the nearest S1-S2 intervals were 304, 310, and 316 ms, the next nearest 298, 310, and 322 ms, and so on. Both Figure 8 and Table III demonstrate substantially better performance than the *SpatialOnly* model in all cases. Figure 8(c) shows that the decreasing performance with decreasing S1-S2 interval for Patient C all but disappears. The breakdown of these results by S1-S2 interval is given in the Supplementary Material.

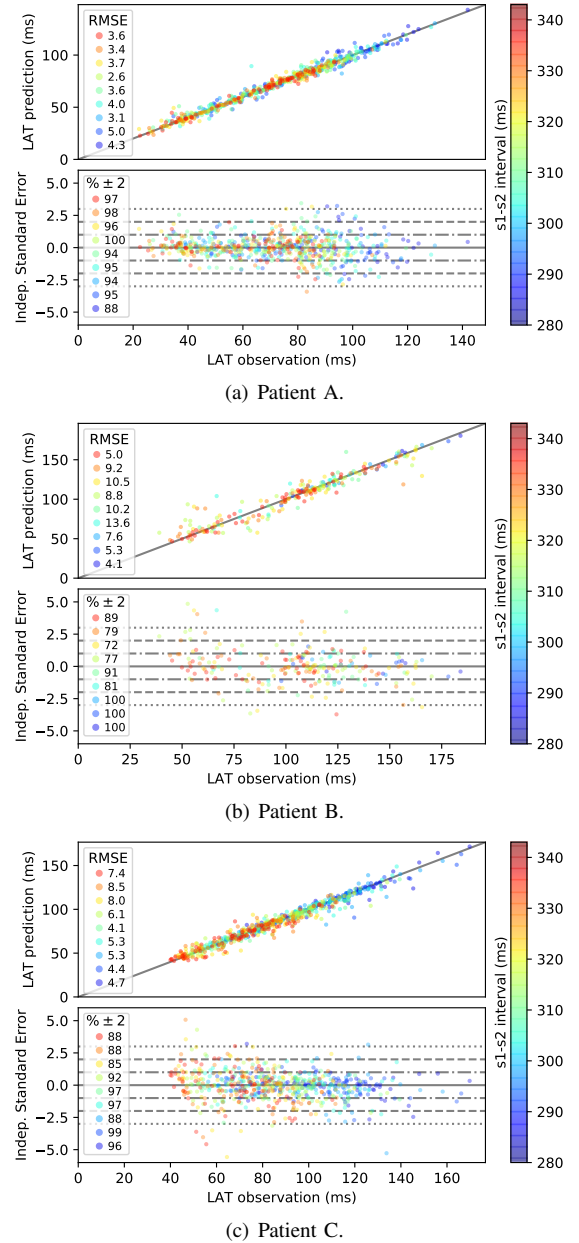
For the *SpatialPacing* model, the predictions seem to be just as good for S1-S2 intervals far apart as for S1-S2 intervals very close together, and certainly we have shown that for S1-S2 intervals of around 25 ms (a range of S1-S2 intervals of 50 ms) the *SpatialPacing* model works well and outperforms the *SpatialOnly* model. This is due to the ability to borrow information about the spatial activation pattern from other pacing protocols, and combine it appropriately in the modelling. This point is emphasized in the bottom four rows of Table III, which show the average RMSE cross-validation results for combining different pacing data into the *SpatialOnly* model. In this case, the RMSE scores are higher and get worse as the S1-S2 interval becomes longer (for Patient B, the result for the next next next nearest neighbours is still significantly better than for modeling using data from only one pacing, due to the low amount of data in this case).

IV. DISCUSSION

A. Uncertainty in LAT observations

Our overall methodology is modular, so other methods for electrogram analysis and data registration to the atrial mesh could be used with our interpolation method. As with this study, uncertainty quantification should be consistent with how Local Activation Time and catheter position were observed. We assumed a normal distribution to LAT in order to combine

Fig. 8. 10-fold cross-validation results for *SpatialPacing* models. Legends show normalized Root Mean Squared Error and percentage of Independent Standard Errors within ± 2 against corresponding S1-S2 interval color.



uncertainties from electrograms and position, and for computational tractability; other choices may be better.

For LAT assignment to electrograms, we extended a ‘Center-of-Mass’ method [13], [14] that does not depend on unreliable fiducial markers, such that uncertainty depended on electrogram morphology consistently with how LAT was assigned. Our uncertainty definition was based on experience with the recordings used in this study, and the LAT distribution embraced a range of plausible features for the ‘true’ LAT. A more systematic calibration would optimize the relative consistency of electrogram uncertainty from one recording to another, and identify suitable timing thresholds.

In clinical data, electrograms can be highly fractionated such that there are multiple activation complexes that need

to be separated; see [29] and [30]. In these cases, there is additional uncertainty associated with bracketing (separating) the complexes in the signal. It will be important in these cases not to overlook this ‘bracketing uncertainty’; LAT assignment will depend heavily on how activation complexes are separated in highly fractionated signals, and it would be better to define large uncertainties than to be overconfident.

Mismatch between the electrode location provided by the electro-anatomic mapping system and the atrial mesh is a further source of uncertainty, which we have also included in our study. Registration of electrode locations with the anatomical mesh is a difficult problem because the anatomical map is produced at a different time to the measurements. Our method takes this uncertainty into account in meaningful way by directly relating the error in the registration method and the corresponding uncertainty in LAT that could likely occur. However, more accurate and robust methods of registration, such as [31] and [32], which deal both with improving the accuracy of registration as well estimating errors, could be used within our framework.

B. GMRF method for interpolation

The GMRF method can include uncertainty in LAT measurements and yield probabilistic LAT maps. In this framework the LAT observations are not taken to be the ground truth, but rather are noisy observations from which the ground truth must be inferred. Notably, interpolation is performed directly on the atrial manifold. For this reason, in comparison with a 3-dimensional Gaussian process interpolation (which is qualitatively similar to previous methods), the GMRF method performed better and made more plausible estimates of uncertainty. GP interpolation did not take account of the manifold, and so resulted in problematic artefacts in the conduction velocity field. Mesh smoothing and adjustment may also affect results in important ways, since the truncation of pulmonary veins (and the left atrial appendage, if this is removed) changes the connectivity of the mesh; it is potentially very important that the actual geometry is properly taken into account as done here.

Accurate assignments of measurements to the atrial mesh is very important and will have large consequences for accuracy; inaccurate/incorrect assignments cannot necessarily be compensated for by low precision. It is difficult to determine a minimum number of mapping points required to use the GMRF method; good predictions can be expected near to training data, but not necessarily in regions of far extrapolation. However, predictive uncertainty increases far away from the data, and so the true underlying LAT may still be captured by the predictive distribution. Our results suggest that the method is robust to decreasing the number of observations, and that an even coverage of accurate LAT recordings is more important than the absolute number of recording locations.

Our method should be applicable to more complex activation patterns such as those which occur during atrial arrhythmias. However, in order to accurately interpolate such patterns, there will need to be sufficient spatial coverage and measurements will need to have high precision, such that

a more complex pattern with lower uncertainty is a better fit to the observations than a smoother pattern with higher uncertainty. It may be necessary to investigate the effect of assigning different hyperparameter priors, and to fix the absolute value of the observation uncertainty in the modeling, in order to fit more complex activation patterns. It should be noted that there is no enforcement of monotonicity in the predictions, so it is possible that our method could produce implausible predictions in regions where data coverage is extremely poor or in regions where the data has been poorly assigned to the anatomical mesh.

The *SpatialPacing* model demonstrates the first case of interpolation across different pacing protocols that we are aware of. Our results demonstrate that information can be shared between different pacing protocols in order to improve LAT map predictions. This potentially allows for more efficient data collection in procedures. For example, it will be possible to use totally different spatial locations for each pacing protocol in the procedure, or to collect data for only a subset of pacing protocols for each placement of the recording catheter. This could radically improve the spatial coverage of data obtained during procedures with multiple protocols.

LAT maps are often patchy, but our results suggest that non-smooth patterns resulting from overlooking uncertainties should be treated with caution. The smoothness of the LAT maps produced by our method is somewhat a result of the observation uncertainties. From a physical point of view, the activation pattern should be assumed to be smooth, unless the data (with observation uncertainty included) strongly suggest otherwise. For this reason, our method of interpolation could be thought of as assuming a smooth prior for the activation pattern, such that more complex activation patterns must be enforced by the data. Probabilistic interpolation of uncertain Local Activation Times should therefore be an important new tool for constructing LAT maps for clinical purposes.

V. CONCLUSION

In this paper we have described and evaluated a novel method to interpolate uncertain Local Activation Times directly on the left atrial manifold and between pacing protocols. Our approach is to first quantify uncertainty in assigning Local Activation Times (LAT) to bipolar electrograms, as well as uncertainty in registration of catheter positions to patient specific anatomical meshes, and then to explicitly incorporate this uncertainty in LAT into a probabilistic interpolation scheme using Gaussian Markov Random Fields (GMRFs) in a Bayesian hierarchical framework. Our methods yield probabilistic LAT maps with predictive mean and standard deviation which account for heterogeneous observation uncertainty and interpolation uncertainty. We used cross-validation to demonstrate performance, showing that our method outperforms a 3D interpolation method which is comparable to existing methodologies, and demonstrated that our method can sensibly combine data from different pacing protocols to improve predictive performance. This approach enables an approach for building maps of LAT in individual patients that is robust in the face of measurement uncertainties.

REFERENCES

- [1] S. E. Williams *et al.*, “Local activation time sampling density for atrial tachycardia contact mapping: How much is enough?” *Europace*, vol. 20, pp. e11–e20, 2018.
- [2] N. M. S. De Groot *et al.*, “Voltage and Activation Mapping: How the Recording Technique Affects the Outcome of Catheter Ablation Procedures in Patients With Congenital Heart Disease,” *Circulation*, vol. 108, pp. 2099–2106, 2003.
- [3] C. H. Roney *et al.*, “A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction,” *Computers in Biology and Medicine*, vol. 104, pp. 278–290, 2019.
- [4] S. E. Williams *et al.*, “The effect of activation rate on left atrial bipolar voltage in patients with paroxysmal atrial fibrillation,” *Journal of Cardiovascular Electrophysiology*, vol. 28, pp. 1028–1036, 2017.
- [5] M. Masè and F. Ravelli, “Automatic reconstruction of activation and velocity maps from electro-anatomic data by radial basis functions,” *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC’10*, pp. 2608–2611, 2010.
- [6] B. Yilmaz, U. Cünedioğlu, and E. Baysoy, “Usage of spline interpolation in catheter-based cardiac mapping,” *Turkish Journal of Electrical Engineering and Computer Sciences*, vol. 18, pp. 989–1002, 2010.
- [7] C. D. Cantwell *et al.*, “Techniques for automated local activation time annotation and conduction velocity estimation in cardiac mapping,” *Computers in Biology and Medicine*, vol. 65, pp. 229–242, 2015.
- [8] K. P. Anderson *et al.*, “Determination of Local Myocardial Electrical Activation for Activation Sequence Mapping A Statistical Approach,” *Circulation Research*, vol. 69, pp. 898–917, 1991.
- [9] S. Liverani and C. D. Cantwell, “Elicitation and visualisation of uncertainty in electrograms for activation time maps,” *Proceedings of the 47th Scientific Meeting of the Italian Statistical Society, Cagliari, 2014*, pp. 1–6, 2014.
- [10] E. Konukoglu *et al.*, “Efficient probabilistic model personalization integrating uncertainty on data and parameters: Application to Eikonal-Diffusion models in cardiac electrophysiology,” *Progress in Biophysics and Molecular Biology*, vol. 107, pp. 134–146, 2011.
- [11] F. Murgatroyd, A. Krahn, and G. J. Klein, *Handbook of Cardiac Electrophysiology: A Practical Guide to Invasive EP Studies and Catheter Ablation*. REMEDICA, 2002.
- [12] K. P. Anderson *et al.*, “Criteria for local myocardial electrical activation: effects of electrogram characteristics,” *IEEE Transactions on Biomedical Engineering*, vol. 40, pp. 169–181, 1993.
- [13] M. El Haddad *et al.*, “Algorithmic detection of the beginning and end of bipolar electrograms: Implications for novel methods to assess local activation time during atrial tachycardia,” *Biomedical Signal Processing and Control*, vol. 8, pp. 981–991, 2013.
- [14] M. El Haddad *et al.*, “Novel algorithmic methods in mapping of atrial and ventricular tachycardia,” *Circulation: Arrhythmia and Electrophysiology*, vol. 7, pp. 463–472, 2014.
- [15] GPy, “GPpy: A gaussian process framework in python,” <http://github.com/SheffieldML/GPy>, since 2012.
- [16] P. Cignoni, G. Ranzuglia, and Visual Computing Lab ISTI-CNR, “Meshlab,” <http://meshlab.sourceforge.net/>.
- [17] L. Antiga *et al.*, “An image-based modeling framework for patient-specific computational hemodynamics,” *Medical & Biological Engineering & Computing*, vol. 46, pp. 1097–1112, 2008.
- [18] D. Frisch, “point2trimesh() — distance between point and triangulated surface,” <https://uk.mathworks.com/matlabcentral/fileexchange/52882-point2trimesh-distance-between-point-and-triangulated-surface>, 2016.
- [19] N. Cressie, *Statistics for spatial data*, ser. Wiley series in probability and mathematical statistics: Applied probability and statistics. J. Wiley, 1993.
- [20] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [21] F. Lindgren, H. Rue, and J. Lindström, “An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach,” *Journal of the Royal Statistical Society, Series B*, vol. 73, pp. 423–498, 2011.
- [22] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2018. [Online]. Available: <https://www.R-project.org/>
- [23] E. Krainski *et al.*, *Advanced Spatial Modeling with Stochastic Partial Differential Equations Using R and INLA*. CRC Press/Taylor and Francis Group, 2019. [Online]. Available: <https://becarioprecario.bitbucket.io/spde-gitbook/>
- [24] “The r-inla project,” <http://www.r-inla.org/>, 2018.
- [25] M. Cameletti *et al.*, “Spatio-temporal modeling of particulate matter concentration through the SPDE approach,” *AStA Advances in Statistical Analysis*, vol. 97, pp. 109–131, 2013.
- [26] D. Simpson *et al.*, “Penalising model component complexity: A principled, practical approach to constructing priors,” *Statistical Science*, vol. 32, pp. 1–28, 02 2017.
- [27] H. Rue, S. Martino, and N. Chopin, “Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 71, pp. 319–392, 2009.
- [28] Kuhn, M., “Building Predictive Models in R Using the caret Package,” *Journal of Statistical Software*, vol. 28, pp. 1–26, 2008.
- [29] F. Ravelli and M. Masè, “Computational mapping in atrial fibrillation : how the integration of signal-derived maps may guide the localization of critical sources,” *Europace*, vol. 16, pp. 714–723, 2014.
- [30] M. Baumert, P. Sanders, and A. Ganesan, “Quantitative-Electrogram-Based Methods for Guiding Catheter Ablation in Atrial Fibrillation,” *Proceedings of the IEEE*, vol. 104, pp. 416–431, 2016.
- [31] M. E. Rettmann *et al.*, “Quantitative modeling of the accuracy in registering preoperative patient-specific anatomic models into left atrial cardiac ablation procedures,” *Medical Physics*, vol. 41, pp. 1–11, 2014.
- [32] A. Cristoforetti *et al.*, “A stochastic approach for automatic registration and fusion of left atrial electroanatomic maps with 3D CT anatomical images,” *Physics in Medicine and Biology*, vol. 52, pp. 6323–6337, 2007.