

Research



Cite this article: Coveney *S et al.* 2020
Gaussian process manifold interpolation for
probabilistic atrial activation maps and
uncertain conduction velocity. *Phil. Trans. R.
Soc. A* **378**: 20190345.
<http://dx.doi.org/10.1098/rsta.2019.0345>

Accepted: 21 April 2020

One contribution of 16 to a theme issue
'Uncertainty quantification in cardiac and
cardiovascular modelling and simulation.'

Subject Areas:

cardiac electrophysiology, statistical methods,
uncertainty quantification

Keywords:

cardiac conduction velocity, atrial fibrillation,
Gaussian process, probabilistic interpolation,
local activation time, manifold

Author for correspondence:

Sam Coveney
e-mail: s.coveney@sheffield.ac.uk


Electronic supplementary material is available
online at <https://doi.org/10.6084/m9.figshare.c.4982324>.

Gaussian process manifold interpolation for probabilistic atrial activation maps and uncertain conduction velocity

Sam Coveney¹, Cesare Corrado³, Caroline H. Roney³,
Daniel O'Hare³, Steven E. Williams³, Mark D. O'Neill³,
Steven A. Niederer³, Richard H. Clayton¹,
Jeremy E. Oakley² and Richard D. Wilkinson²

¹Insigneo Institute for in-silico medicine and Department of
Computer Science, and ²School of Mathematics and Statistics,
University of Sheffield, Sheffield, UK

³Division of Imaging Sciences and Biomedical Engineering, King's
College London, London, UK

 SC, 0000-0002-7134-3196; CC, 0000-0002-8914-8735;
CHR, 0000-0001-6809-0928; SAN, 0000-0002-4612-6982;
RHC, 0000-0002-8438-7518; JEO, 0000-0002-9860-4093;
RDW, 0000-0001-7729-7023

In patients with atrial fibrillation, local activation time (LAT) maps are routinely used for characterizing patient pathophysiology. The gradient of LAT maps can be used to calculate conduction velocity (CV), which directly relates to material conductivity and may provide an important measure of atrial substrate properties. Including uncertainty in CV calculations would help with interpreting the reliability of these measurements. Here, we build upon a recent insight into reduced-rank Gaussian processes (GPs) to perform probabilistic interpolation of uncertain LAT directly on human atrial manifolds. Our Gaussian process manifold interpolation (GPMI) method accounts for the topology of the atrium, and allows for calculation of statistics for predicted CV. We demonstrate our method on two clinical cases, and perform validation against a simulated ground truth. CV uncertainty depends on data density, wave propagation direction

and CV magnitude. GPMI is suitable for probabilistic interpolation of other uncertain quantities on non-Euclidean manifolds.

This article is part of the theme issue ‘Uncertainty quantification in cardiac and cardiovascular modelling and simulation’.

1. Background

Electrical activation of cardiac tissue acts to initiate and synchronize mechanical contraction. The spread of the activation wave is characterized by local activation times (LATs) and conduction velocity (CV), where CV describes the local speed and direction of the wavefront [1]. CV is a clinically important material property of cardiac tissue [2,3], with decreases in CV often associated with abnormalities such as tissue fibrosis. CV can also be used to calibrate patient-specific models of electrophysiology [4], which have the potential to be used for planning clinical interventions [5]. CV can be derived from measurements of LAT at a set of locations. In the clinical setting, LAT can be obtained from electrograms recorded using catheters placed inside the heart. However, reliable CV estimation based on uncertain estimates of LAT from noisy electrograms is difficult [1,6]. Wave collision and lines of block also present challenges for CV estimation [7,8].

In a previous study of electrical activation in the human left atrium, we identified uncertainties associated with identifying LAT from individual electrograms, as well as uncertainty associated with the registration of the recording electrode to an anatomical mesh [9]. We demonstrated that it is possible to create probabilistic LAT maps over the left atrium, which is modelled as a two-dimensional (2D) manifold in three-dimensional (3D) space represented as a triangle mesh. These LAT maps take account of uncertainties in both LAT measurements and the interpolation procedure itself. Importantly, the interpolation was performed directly on the manifold geometry, taking account of the physical distances travelled by the wavefront and the connectivity between points on the mesh. We used a stochastic partial differential equation (SPDE) approach linking Gaussian Markov random fields (GMRFs) to a class of Gaussian processes (GPs) with Matérn covariance functions [10]. This enabled us to construct random fields on non-Euclidean domains such as the manifold describing the left atrium.

Given a probabilistic LAT map on an atrial mesh, a CV map can easily be calculated from the LAT map posterior mean. A natural approach to compute how our uncertainty about LAT affects our estimate of CV is to use stochastic sampling. We can simulate samples from our posterior distribution for LAT, calculating CV for each of them, thus estimating a distribution for CV at every mesh location. However, doing this with the GMRF model used in [9] shows that this model is not suitable for estimating uncertainty about CV, because random samples of LAT are insufficiently smooth. In practical terms, this leads to samples of CV that are not consistent with observations of LAT. In technical terms, the Matérn class of covariance functions depends upon a smoothness parameter ν , such that process samples are k -times mean-squared differentiable if and only if $\nu > k$ [11]. For the SPDE-GMRF, ν is constrained to be one, so that posterior LAT samples are not differentiable. Consequently, they are too rough to represent a spreading electrical wavefront, meaning the approach cannot be used to calculate the CV distribution. In this paper, we describe a new GP approach that does allow us to estimate CV distributions from a probabilistic interpolation of noisy LAT measurements.

The reason for using the SPDE-GMRF approach was that an ‘ordinary GP’, for which the covariance matrix is a function of the Euclidean distance matrix between 3D spatial locations, would not account for the geometry or topology of the atrial manifold. It is not possible to simply replace Euclidean distances with for example geodesic distances because the resulting covariance matrix is not positive semi-definite. However, the meaningful distance between locations for this physical system is geodesic distance on the manifold, not distance through the manifold that ignores the topology. A recent advance in reduced-rank GPs showed that it is possible to expand a stationary covariance kernel in terms of eigenfunctions of the Laplacian operator on the spatial

domain [12]. The contribution of this paper is to address the problem of obtaining distributions of CV from probabilistic interpolation of LAT. We apply this recent advance to a 2D manifold in 3D space, so we are able to form a GP on the manifold using any stationary covariance function, including those which lead to differentiable samples. This allows us to calculate CV distributions by probabilistic interpolation of LAT.

This paper is organized as follows. We first explain how to interpolate LAT using ordinary GPs. We then review how to expand the covariance function as in [12], but applied to a 2D triangulated manifold representing the atrium, such that we can simply replace the covariance matrix in ordinary GP regression to obtain a GP confined to the manifold. Using clinical data, we perform LAT interpolation and obtain corresponding probabilistic CV maps. We also validate our methodology against a simulated ground truth. We find that CV uncertainty depends on LAT observation uncertainty, data density, alignment of observations with respect to wave propagation direction, and on CV magnitude. We then discuss future work made possible by solving the geometrical aspect of probabilistic interpolation on atrial manifolds.

2. Methods

(a) Data collection

As proof of concept, we present two clinical cases in this paper. Additional details on data collection are given in the electronic supplementary material. For the first case, Patient 1, (identical to Patient A in [9]) observations were collected using an S1–S2 pacing interval: a 2-beat drive train with a cycle length of S1 470 ms followed by a single premature extra stimulus S2, where $S2 < S1$. For the second case, Patient 2, a constant pacing at cycle length 500 ms was used, i.e. S1 only. For Patient 1, electrograms were processed to obtain uncertain LAT as described in [9]. For Patient 2, visual inspection of each point was undertaken by one of the authors to assign LAT to the earliest sharp deflection, without an assigned uncertainty based on the electrogram analysis.

In both cases, a triangulated mesh of the left atrium, obtained using the respective mapping systems, was processed to clip away the pulmonary veins and mitral valves to leave holes, and the triangulation was smoothed (using mainly surface-preserving Laplacian smoothing, element subdivision and short-edge collapsing) to be reasonably regular, with approximately 10^4 non-intersecting mesh faces. This yields an interpolation resolution mesh, in which LAT observations were assigned to the nearest vertex, and given an additional observation uncertainty arising from this mesh assignment as in [9]. Note that our method does not require observation-specific uncertainty for LAT, or specific methods for determining LAT and assigning observations to a mesh.

(b) Gaussian process regression

A function f , for which only noisy observations y_i are available at spatial locations \mathbf{x}_i , can be modelled by a GP as follows:

$$\left. \begin{aligned} f(\mathbf{x}) &\sim \mathcal{GP}(0, \tau^2 k(\mathbf{x}, \mathbf{x}')), \\ \text{and} \quad y_i &= f(\mathbf{x}_i) + \epsilon_i, \quad \text{where } \epsilon_i \sim \mathcal{N}(0, \sigma_i^2), \end{aligned} \right\} \quad (2.1)$$

where ϵ_i is the (heteroscedastic) noise for each observation. The covariance function or 'kernel' k specifies the covariance $\text{cov}(f(\mathbf{x}), f(\mathbf{x}')) := \tau^2 k(\mathbf{x}, \mathbf{x}')$ as a function of location. Here, the mean of the GP is set to zero. We can generalize by using a non-zero parametric mean function, but by centring and scaling the data y to be mean zero, variance one, we find that a zero mean is sufficient and simplifies the exposition. A GP is 'fit' to data by conditioning upon the observations $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^n$ and optimizing τ^2 and any kernel hyperparameters to maximize the log-likelihood. Here, σ_i is obtained by pre-processing LAT and mesh assignment rather than by optimization, although such optimization would be possible.

Posterior predictions (i.e. using the GP conditioned upon the data) are multivariate Gaussian distributions specified by the posterior mean and variance. At input (e.g. location) \mathbf{x}^* , the posterior mean and variance is given by

$$\left. \begin{aligned} \mathbb{E}[f(\mathbf{x}^*) | \mathcal{D}] &= \mathbf{k}_*^T (\mathbf{K} + \mathbf{\Sigma})^{-1} \mathbf{y} \\ \text{and} \quad \mathbb{V}[f(\mathbf{x}^*) | \mathcal{D}] &= \tau^2 k(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{k}_*^T (\mathbf{K} + \mathbf{\Sigma})^{-1} \mathbf{k}_* \end{aligned} \right\} \quad (2.2)$$

where \mathbf{K} is an $n \times n$ matrix containing the covariance between all pairs of observations $\tau^2 k(\mathbf{x}_i, \mathbf{x}_j)$; \mathbf{y} is the vector of observations y_i ; and $\mathbf{\Sigma} \equiv \text{diag}(\sigma_1^2 + \eta, \dots, \sigma_n^2 + \eta)$ contains the observation error variances σ_i^2 (for homoscedastic noise, all σ_i are equal) plus a ‘nugget’ η representing additional observation uncertainty common to all observations (η is optimized with other hyperparameters). The cross-covariance between observations and predictions is given by the vector \mathbf{k}_* , and $\tau^2 k(\mathbf{x}^*, \mathbf{x}^*)$ is the covariance between prediction locations (both are matrices when \mathbf{x}^* contains multiple locations). By predicting at multiple locations simultaneously, e.g. all points on the atrial mesh, we obtain expressions for the posterior mean and variance at all locations.

(c) Local activation time interpolation with a Gaussian process

We assume the underlying LAT field is a smoothly varying function of location and model it as a GP: $\text{LAT} \equiv f(\mathbf{x}) \sim \mathcal{GP}(0, \tau^2 k(\mathbf{x}, \mathbf{x}'))$. We model the relationship between the noisy LAT observations \mathbf{y} and the ‘true’ LAT as $y_i = f(\mathbf{x}_i) + \epsilon_i$ where $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$, as in equation (2.1). LAT observation y_i is assigned to a coordinate on the atrial mesh \mathbf{x}_i .

For the GP kernel, we choose the Matérn 3/2 kernel ($\nu = 3/2$), which gives a once (mean squared) differentiable process. This choice was made because CV is obtained from the derivative of LAT, but we do not want to make assumptions about the smoothness of CV. The Matérn 3/2 kernel is a stationary kernel that, when working in Euclidean space, depends only on the Euclidean distance $r_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|_2$ between observations i and j , i.e. $k(\mathbf{x}_i, \mathbf{x}_j) \equiv k(r_{ij})$, and can be written, with length-scale hyperparameter l , as

$$k_{3/2}(r) = \left(1 + \frac{\sqrt{3}r}{l} \right) \exp \left(-\frac{\sqrt{3}r}{l} \right). \quad (2.3)$$

In [9], we jointly interpolated LAT across both space and S1-S2 pacing using an order one autoregressive model (i.e. AR(1)) for the dependence on S2. We do not interpolate over S2 in this paper, but could do so using a separable covariance structure such as $k_{\text{spatial-pacing}}((x, S2), (x', S2')) = k_{3/2}(x, x') \times k_{\text{pacing}}(S2, S2')$, where $k_{\text{pacing}}(S2, S2')$ is a suitable standard kernel (e.g. squared exponential), simplifying the approach from [9].

(d) Gaussian processes on a manifold

The standard GP approach described above assumes, due to the kernel, that correlations are a function of Euclidean distance. Since electrical activation advances over the atrial manifold, instead of working with Euclidean distances $\|\mathbf{x} - \mathbf{x}'\|_2$ in the domain \mathbb{R}^3 , we want to limit the domain to locations on the atria, which we denote by Ω , and assume correlations are a function of geodesic path length on that domain. Unfortunately, it is not possible to simply replace the Euclidean distance in a kernel such as equation (2.3) with the geodesic distance and still form a valid covariance matrix, as this leads to non-positive definite kernels [13].

Here, we take advantage of a recently developed [12,14] reduced-rank method for expressing a stationary kernel in terms of the kernel spectral density and eigensolutions of the Laplacian operator on Ω , that allows us to bypass the need to explicitly define a covariance function. We choose a spectral density, and then define a stochastic process with that spectral density, which then implicitly defines a covariance function on Ω (by the Wiener–Khinchin theorem).

We use a finite-dimensional model of the form

$$\tilde{f}(\mathbf{x}) = \sum_{k=1}^M f_k \phi_k(\mathbf{x}), \quad \text{where } f_k \sim N(0, S(\sqrt{\lambda_k})), \quad (2.4)$$

where the f_k are parameters to be learnt. $\tilde{f}(x)$ is a (finite dimensional) GP with covariance function given by the basis function expansion

$$\text{Cov}(\tilde{f}(\mathbf{x}_i), \tilde{f}(\mathbf{x}_j)) = \tilde{k}(\mathbf{x}_i, \mathbf{x}_j) = \sum_{k=1}^M S(\sqrt{\lambda_k}) \phi_k(\mathbf{x}_i) \phi_k(\mathbf{x}_j). \quad (2.5)$$

The basis vectors $\phi_k(\mathbf{x})$ are chosen to be eigenfunctions corresponding to eigenvalues λ_k of the negative Laplacian operator on the domain Ω , i.e. $\{\lambda_k, \phi_k(\cdot)\}$ are solutions of the system

$$\text{and} \quad \left. \begin{aligned} -\nabla^2 \phi_k(\mathbf{x}) &= \lambda_k \phi_k(\mathbf{x}), & \mathbf{x} \in \Omega \\ \nabla \phi_k(\mathbf{x}) \cdot \mathbf{n}_{\partial\Omega} &= 0, & \mathbf{x} \in \partial\Omega \end{aligned} \right\} \quad (2.6)$$

with $0 \leq \lambda_1 < \lambda_2 < \dots$. Note that this basis is independent of our choice of covariance structure/spectral density, allowing us to try different structures without needing to recompute the basis vectors. We use Neumann boundary conditions specifying zero derivative normal to the mesh edges (pulmonary veins and mitral valve), as these do not specify LAT in advance of the interpolation (by contrast [12,14] used Dirichlet boundary conditions).

$S(w)$ is a spectral density [11] which we are free to choose; it is S that determines the nature of the covariance structure of our model. In this paper, we use the spectral density corresponding to the Matérn family of covariance functions, which is given by

$$S(w) := \frac{2^D \pi^{D/2} \Gamma(v + D/2) (2v)^v}{\Gamma(v) l^{2v}} \left(\frac{2v}{l^2} + 4\pi^2 \omega^2 \right)^{-(v+D/2)}, \quad (2.7)$$

where v is smoothness, l is length-scale, D is dimensionality ($D = 2$ for a 2D manifold in 3D space) and Γ is the Gamma function. The spectral density of our model \tilde{f} is only approximately $S(w)$, but converges to $S(w)$ as M and Ω grow [12].

The motivation in [12] for using expansion (2.4) is to reduce the computational complexity of hyperparameter estimation, as the basis is independent of the covariance function. By contrast, our motivation for using this approach is that the solution of equation (2.6) can be computed numerically on a discrete mesh, even for non-Euclidean domains, allowing us to define a GP model directly on the manifold defined by Ω . In particular, we have chosen to model LAT on an atrial manifold using a stochastic process with Matérn spectral density.

(e) Eigenfunction calculations

To apply this kernel expansion (2.5) to a 2D manifold in 3D space, we represent the manifold as a discrete triangle mesh with N vertices and represent the Laplacian operator with the ‘mesh/cotan Laplacian’ (see [15]) given by

$$(\nabla^2 s)_i = \frac{1}{2A_i} \sum_{j \in N(i)} (\cot a_{ij} + \cot b_{ij})(s_i - s_j), \quad (2.8)$$

where s represents an arbitrary scalar field defined at vertices on the mesh. Figure 1a represents the terms in equation (2.8). Applied to a triangular mesh, the $N \times N$ operator matrix \mathcal{L} represents the Laplacian with Neumann boundary conditions at the mesh edges (mitral valve and pulmonary veins). We calculate the Laplacian matrix \mathcal{L} neglecting the areas A_i , and solve the eigenvalue problem $\mathcal{L}\phi = \lambda A\phi$, where A is a diagonal matrix of areas A_i (see [16]), solving for the M smallest eigenvalues and corresponding eigenvectors using the *eigsh* function from the Python package *Scipy* [17].




Figure 1. (a) Representation of the terms in equation (2.8), where the shaded area represents the Voronoi cell with area A_i and the angles opposing the edge between vertex i and j are a_{ij} and b_{ij} . (b) Subdivision of a mesh triangle face into nine triangles by addition of vertices, including a vertex at the mesh face centroid. (Online version in colour.)

We obtain M eigenvectors such that $\phi_k(\mathbf{x}_i) \equiv \phi_k^i$, where $k = 1 \dots M$ index the eigenfunction and $i = 1 \dots N$ index the mesh vertices. These eigenfunctions belong to the spatial domain represented by the mesh, and as such the topology (connectivity) and geometry (distances) of the atrial manifold are properly accounted for. Note that the eigenproblem only needs to be solved once for a particular mesh. We can then construct all of the covariance matrices needed for GP regression. We refer to this method as Gaussian process manifold interpolation (GPMI). Our code for applying GPMI to atrial manifolds can be found at [18].

To represent the atrial mesh smoothly, we use an interpolation resolution mesh of approximately 10^4 vertices which corresponds to an edge length approximately 1 mm. In order to ensure a fine discretization of the Laplacian, and also to allow the calculation of eigenfunction gradients (required for the posterior of the GP gradient; see §2f), we subdivide the interpolation mesh: each triangular element (three vertices) is subdivided into nine elements by creating two additional vertices along each edge and one additional vertex at the element centroid, as shown in figure 1b. The Laplacian and the corresponding eigenfunctions are calculated for this finer mesh.

For each original mesh face, there are now 10 vertices which lie in the plane of this face. To calculate the gradients of the eigenfunctions at a face centroid, we first rotate these (finer mesh) vertex coordinates into a 2D frame, and then fit a cubic polynomial to the vertex eigenfunction values and 2D coordinates using least-squares regression. We then evaluate the gradient of the eigenfunctions at the centroid using the derivative of the polynomial expression, and rotate this (vector) gradient back into the original 3D coordinate frame. This allows us to calculate the eigenfunction gradients at the centroids, even when the eigenfunctions vary nonlinearly across a face. For simplicity, we retain the eigenfunction values at the mesh vertices and face centroids of the original mesh only (i.e. we discard the vertices from subdividing the face edges), which allows us to assign LAT observations and make LAT predictions at both vertices and centroids.

Boundary conditions at the edges of the mesh (holes representing the mitral valve and pulmonary veins) will leave unwanted artefacts in the interpolation. To deal with this problem, after forming the interpolation resolution mesh and prior to subdivision as described above, we append tubes of triangular elements to the atrial holes, consisting of 10–20 layers of elements (depending on the mesh resolution) and extending away from the mesh centre of mass, using a custom algorithm described in the electronic supplementary material. This takes the boundaries away from the ‘real atrium’ and removes the unwanted boundary effects on the interpolation. Note that the topology is unchanged by this procedure. After subdivision as described above, we retain only the eigenvector entries corresponding to the original mesh vertices and centroids before extension. A figure of the mesh extension result for Patient 2 is given in the electronic supplementary material.

(f) Conduction velocity distribution

CV is the inverse of the LAT gradient. The distribution for either the LAT gradient or CV could be calculated by first generating samples from the posterior distribution for LAT, and

then calculating the desired quantities for each sample. However, since applying a linear operator (such as the gradient) to a GP results in another GP, the gradient of LAT is also a GP. Therefore, we can instead sample directly from this gradient distribution. The posterior distribution for the gradient of LAT is given by similar expressions to the posterior distribution for LAT (equation (2.2)), except that the covariance entries involving prediction locations require derivatives of the covariance function with respect to these prediction locations

$$\mathbb{E} \left[\frac{\partial f(\mathbf{x}^*)}{\partial \mathbf{x}^*} \mid \mathcal{D} \right] = \frac{\partial \mathbf{k}_*^T}{\partial \mathbf{x}^*} (\mathbf{K} + \boldsymbol{\Sigma})^{-1} \mathbf{y}$$

and

$$\mathbb{V} \left[\frac{\partial f(\mathbf{x}^*)}{\partial \mathbf{x}^*} \mid \mathcal{D} \right] = \tau^2 \left. \frac{\partial^2 k(\mathbf{x}_a, \mathbf{x}_b)}{\partial \mathbf{x}_a \partial \mathbf{x}_b} \right|_{\mathbf{x}_a = \mathbf{x}_b = \mathbf{x}^*} - \frac{\partial \mathbf{k}_*^T}{\partial \mathbf{x}^*} (\mathbf{K} + \boldsymbol{\Sigma})^{-1} \frac{\partial \mathbf{k}_*}{\partial \mathbf{x}^*} \quad (2.9)$$

When the kernel is expressed as in equation (2.5), the necessary kernel derivatives are given as

$$\left. \begin{aligned} \frac{\partial k(\mathbf{x}_i, \mathbf{x}_j)}{\partial \mathbf{x}_i} &= \sum_{k=1}^M S(\sqrt{\lambda_k}) \frac{\partial \phi_k(\mathbf{x}_i)}{\partial \mathbf{x}_i} \phi_k(\mathbf{x}_j) \\ \frac{\partial^2 k(\mathbf{x}_i, \mathbf{x}_j)}{\partial \mathbf{x}_i \partial \mathbf{x}_j} &= \sum_{k=1}^M S(\sqrt{\lambda_k}) \frac{\partial \phi_k(\mathbf{x}_i)}{\partial \mathbf{x}_i} \frac{\partial \phi_k(\mathbf{x}_j)}{\partial \mathbf{x}_j} \end{aligned} \right\} \quad (2.10)$$

and

The derivatives of the eigenfunctions can all be calculated in advance of the interpolation, after calculating the eigenfunctions themselves (see §2e). Note that it is therefore possible to include uncertain *observations* of the gradient of LAT into the modelling.

Given the gradients of the eigenfunctions at the face centroids (see §2e), the posterior distribution of the gradient of LAT can be calculated at all centroids. However, since neither magnitude or inverse are linear operations, we calculate the distributions for the magnitude of the gradient of LAT and CV by sampling from the posterior LAT gradient and applying the operations to the samples. At each centroid, we generate 2000 samples from the multivariate normal distribution for the components of the gradient of LAT given by equation (2.9). We calculate the standard deviation, 9th, 25th, 50th, 75th, 91st percentiles for the magnitude of the posterior samples of the gradient of LAT (the percentiles can be inverted to obtain CV percentiles). We use the posterior mean of (2.9) to represent prediction of the magnitude, and statistics from sampling for prediction uncertainty.

(g) Simulations

To quantitatively evaluate our methodology, we compare predictions of LAT and CV against a simulation ‘ground truth’. A computational mesh with a regular triangulation with edge length of 215 μm was obtained for Patient 1 (1047446 vertices). Simulations were performed with the Cardiac Arrhythmia Research Package (CARP), an electrophysiology solver suitable for high-performance computing [19]. We simulate S1 pacing from the coronary sinus for S1 750 ms, using the modified Mitchell–Schaeffer ionic model with heterogeneous parameters to generate a spatially heterogeneous CV field (details in electronic supplementary material). To create a ‘virtual patient’ from these simulations, we downsample the simulation mesh to an interpolation resolution mesh, and match the interpolation mesh vertices to the nearest simulation mesh vertices. For each interpolation mesh vertex, we record LAT from the matching simulation vertex.

We calculate CV for the simulation in two ways. Firstly, we calculate CV element-wise at each interpolation centroid from LAT at the face vertices (‘element’ method): given a face defined by anti-clockwise vertices $\{i, j, k\}$, face area A , face normal \mathbf{N} , edge vectors \mathbf{e}_{ab} , with LAT at vertex l given by t_l , the centroid LAT gradient and CV are given by (\times denotes the vector product)

$$\text{CV} = v / |v|^2, \quad \text{where } v \equiv \nabla \text{LAT} = -\frac{1}{2A} \mathbf{N} \times (t_l \mathbf{e}_{jk} + t_j \mathbf{e}_{ki} + t_k \mathbf{e}_{ij}). \quad (2.11)$$

Table 1. Number of interpolation mesh vertices and centroids, and number of observations, for each case.

case	vertices	centroids	observations
Patient 1	9460	18 680	61
Patient 2	17 235	34 098	1228
virtual	45 180	89 744	various

Secondly, since the CV field is heterogeneous and CV often varies on a length-scale smaller than the distance between observations, we also calculate CV taking account of a larger area ('wave' method). For the 20 vertices nearest to where CV is to be calculated, the vertex coordinates are projected into 2D such that the geodesic distances are optimally conserved. A least-squares linear fit to LAT is used to calculate a gradient corresponding to plane wave propagation [8]. We do this for 1000 random centroids locations, optimized to be well spaced over the mesh using a maximin distance criterion over 10 000 designs. Note that 20 vertices corresponds to at least two 'rings' of vertices around a centroid.

To validate our approach for predicting LAT and CV, we calculate the *normalized* root mean square error (nRMSE), reported as a percentage, and independent standard error (ISE) for all predictions $E[z_i]$ (with variance $V[z_i]$) versus ground truths z_i , where nRMSE (%) and ISE are given by

$$\left. \begin{aligned} \text{nRMSE} &= \frac{100}{\text{range}} \sqrt{\frac{1}{N} \sum_{i=1}^N (E[z_i] - z_i)^2} \\ \text{and} \quad \text{ISE}_i &= (E[z_i] - z_i) / \sqrt{V[z_i]} \end{aligned} \right\} \quad (2.12)$$

where $\text{range} = \max(z_i) - \min(z_i)$.

For LAT, about 95% of ISE scores should lie within the interval ± 2 . We emphasize that LAT, and therefore (implicitly) the gradient of LAT, is what is directly modelled in our framework. Therefore, for comparison with the ground truth, we predict and compare against the magnitude of the gradient of LAT, rather than against CV, using the standard deviation of the magnitudes of LAT gradient samples. This prevents the nRMSE and ISE scores from being heavily distorted by the inversion involved in calculating CV from the LAT gradient (see §4). We calculate these scores for the LAT gradient for comparison with the two methods of calculating CV outlined above. The meaning of the terms in equation (2.12) for the magnitude of the LAT gradient must be explained: we use the magnitude of the posterior mean of the LAT gradient as $E[z_i]$, and the standard deviation of the magnitude of the posterior LAT gradient samples as $\sqrt{V[z_i]}$.

3. Results

The number of interpolation mesh vertices and centroids, and the number of observations, are given in table 1. For the virtual patient, the number of observations will be given in the text for each case. We use $M = 256$ basis functions, which we determined to be sufficient for modelling LAT (see electronic supplementary material, S3).

Validation results for the virtual patient are given in table 2, for prediction of LAT and the magnitude of LAT gradient. For each table row, LAT observations were chosen by picking n random vertices and corresponding LAT values, and corrupting these values with additive random noise $\epsilon_i \sim \mathcal{N}(0, 1)$. The nRMSE and ISE scores were calculated and averaged across 10 different observation designs for each n .

(a) Local activation time interpolation

Figure 2 shows LAT interpolation for Patient 1, overlaid with CV vectors coloured by the LAT posterior mean (a) and standard deviation (b) at the face centroids. The LAT observations y_i and




Figure 2. Posterior LAT distribution for Patient 1, shown as CV vectors at face centroids coloured by LAT posterior mean (*a*) and standard deviation (*b*). Spheres represent LAT observations y_i (*a*) and observation noise σ_i (*b*). (Online version in colour.)

Table 2. Validation results for prediction against ground truth (simulation), averaged over 10 predictions using different random observation data. ‘Wave’ and ‘element’ refer to methods for calculating CV described in S2f, but the corresponding validation scores are for magnitude of LAT gradient. Both nRMSE and ISE coverage are in %.

number of observations	LAT		wave $ \nabla\text{LAT} $		element $ \nabla\text{LAT} $	
	nRMSE	$ \text{ISE} \leq 2$	nRMSE	$ \text{ISE} \leq 2$	nRMSE	$ \text{ISE} \leq 2$
50	2.34	95.7	15.0	91.8	14.0	91.5
100	1.51	94.3	11.7	93.7	11.1	93.4
250	0.91	96.0	9.46	94.5	9.16	93.9
500	0.66	96.1	8.38	93.1	8.22	92.2
750	0.56	96.1	7.62	93.3	7.56	92.2
1000	0.50	95.6	7.22	92.6	7.22	91.6

noise standard deviations σ_i are shown as coloured spheres. LAT prediction uncertainty depends on the distance from the observations and on the observation uncertainty. These interpolation results are comparable with those in [9]. Figure 3 shows the LAT interpolation for Patient 2. In this case, the observation noise only arises from assigning observations to mesh vertices and not from LAT assignment to electrograms, resulting in a very small observation noise. Prediction uncertainty is significantly lower than for Patient 1 due to the increased spatial data coverage.




Figure 3. Posterior LAT distribution for Patient 2, shown as CV vectors at face centroids coloured by LAT posterior mean (a) and standard deviation (b). Spheres represent LAT observations y_i (a) and observation noise σ_i (b). (Online version in colour.)

The validation scores for LAT prediction for the virtual patient show that ISE coverage ($|ISE| \leq 2$) is consistently approximately 95%, and the nRMSE is already less than 1.0% for 250 observations. Note that these scores were comparing all LAT predictions against all ground truth LAT values. These scores demonstrate excellent prediction of LAT using the GPMI method.

(b) Conduction velocity

Figure 4 shows CV for the virtual patient. In the top row, CV is calculated at centroids using the ‘element’ method equation (2.11) with known (ground truth) LAT at the mesh vertices. The second and third rows show CV prediction using 1000 random observations of LAT at mesh vertices, with random noise $\epsilon \sim \mathcal{N}(0,1)$ corrupting these observations for the third row. The predictions capture the heterogeneous CV field produced in the simulation, though the features are more smoothed in the prediction, especially where the true CV varied on a lengthscale shorter than the resolution given by the distances between observations. The predicted CV field is slightly smoother in the case with noisy LAT observations. Overall, the predictions seem to perform very well, capturing even the locations of wave collision and stimulus location.




Figure 4. CV vectors at face centroids coloured by CV magnitude. CV from simulation (a), and CV predictions using 1000 observations randomly selected from the simulation: noiseless observations (b) and noisy observations (c). (Online version in colour.)

Figure 5 shows the prediction uncertainty (interquartile range) for CV magnitude, corresponding to the predictions in figure 4. There is significantly more uncertainty in CV for prediction using noisy LAT observations (the colourbar scales are different). Higher CV correlates with higher prediction uncertainty (as is also the trend for Patient 2, discussed below).

Table 2 gives the validation scores for the magnitude of the LAT gradient for the virtual patient. The scores are very similar for both methods of calculating the ground truth. Notably, the nRMSE scores are an order of magnitude larger for $|\nabla\text{LAT}|$ prediction than for LAT prediction, although the ISE scores suggest that the truth is captured well by the uncertainty. For 250 observations, the nRMSE score for $|\nabla\text{LAT}|$ is less than 10% (and the nRMSE score for LAT is less than 1%). We propose that 250 observations (spaced over the atria) could be considered a lower bound for having confidence in CV predictions calculated with our methodology.

Figure 6 shows prediction versus observation plots (left) and ISE plots (right) for the virtual patient for 1000 noisy observations, where the ground truth is for $|\nabla\text{LAT}|$ calculated by the ‘wave’ method. Points are ordered so that larger standard deviations are plotted in the foreground. The error bars in the prediction versus observation plots correspond to the 9th and 91st percentiles. The uncertainty mostly captures the ground truth. Notably, the uncertainty is highest where $|\nabla\text{LAT}|$ is overestimated (i.e. CV is underestimated) relative to the ground truth. This is likely because regions of fast CV cannot be captured when uncertainty is high.




Figure 5. CV vectors at face centroids coloured by CV magnitude interquartile range, for the predictions in figure 4. Note that the colour scales are different for each case, due to differences in prediction uncertainty for CV magnitude. (Online version in colour.)




Figure 6. Validation plots for the magnitude of the gradient of LAT. Simulation (ground truth) values were obtained with the ‘wave’ method for 1000 well-spaced centroid locations. For prediction, 1000 noisy LAT observations were used. The error bars on the prediction versus truth plot (left) represent the 9th and 91st percentiles. (Online version in colour.)

For Patient 1, there are only 61 LAT observations, much less than the 250 minimal number of observations indicated from the simulation studies. These observations are also highly clustered such that large sections of the mesh have no observations at all (figure 2). Therefore, we do not present plots of CV magnitude for Patient 1. Figure 7 shows CV calculated from the posterior of the LAT gradient for Patient 2, coloured by the CV magnitude (figure 7a) and the interquartile




Figure 7. CV predictions for Patient 1, shown as CV vectors at face centroids coloured by CV magnitude (*a*) and interquartile range (*b*). Spheres show these quantities at observation vertices (interpolated from neighbouring faces). (Online version in colour.)

range (figure 7*b*). The coloured spheres represent these values (interpolated from neighbouring faces) at observation vertices. The predicted CV magnitude is sometimes very high far from the observations, where the GP flattens out beyond the data (shallow LAT gradients correspond to high CV), and in regions of wave collision where LAT is interpolated across two different wavefronts (the calculated value of CV does not represent wave propagation).

CV uncertainty depends highly on distance from observations, but also on the direction of the (predicted) CV field, so that physiological and more certain CV predictions are made where there are several nearby observations aligned along the direction of (predicted) wave propagation. This makes physical sense as observations aligned along the direction of changing LAT are required to inform the gradient of LAT, and is especially useful since we do not know the direction of propagation in advance of the interpolation.

4. Discussion

In this study, we used GPs to interpolate LAT directly on atrial manifolds in order to calculate CV distributions. The results show that point-wise uncertainty in LAT interpolation is not sufficient

for inferring uncertainty in CV. Rather, CV uncertainty depends on LAT observation uncertainty, data density, and the presence of sufficient observations along the direction of wave propagation. Also, larger CV will generally have higher associated uncertainties due to the inverse relationship between the gradient of LAT and CV, such that the same uncertainty in LAT will have much larger effects on CV for shallow LAT gradients (for a given spacing of percentiles for the magnitude of the gradient of LAT, the percentiles on the distribution for CV will become further apart as the gradient of LAT is decreased). Precisely measuring fast CV is therefore difficult. Uncertainty in CV restitution curves could be important for calibrating atrial models, e.g. in [4,20] where the CV uncertainty could weight the comparison between simulated and observed restitution curves.

We have focused here on solving the geometry aspect of the problem of probabilistic interpolation on atrial manifolds, which had only been possible previously for the special case of GMRF (as in our previous publication [9]). The method presented here, GPML, allows us to solve the problem of calculating CV statistics from a probabilistic LAT interpolation. It would also be suitable for probabilistic interpolation of other quantities such as voltage and tissue parameters. The flexible modelling framework presented here means that *observations* of LAT gradient could also be included into the modelling. This could allow both LAT and LAT gradient observations (for example, from multi-polar catheters where consideration of wave propagation is used to obtain precise CV values [6,8]), to be used in a unified and consistent modelling framework. It may also be possible to include linear constraints on interpolation via the covariance function [21], and to directly include physics using latent force models as in [22].

In theory, it should be possible to apply this modelling framework to a 3D domain such as the ventricles. Besides the increase in the number of basis functions that may be required to accurately expand the kernel [12], the main difficulty may be in the boundary conditions at the *surface* of the domain. Some form of mesh extension analogous to that done here may be possible, or perhaps different boundary conditions might be applied. One possibility may be to use a separable kernel structure, similar to the ‘spatial-pacing’ model mentioned in §2c, whereby the ‘thickness’ direction is accounted for with a separate kernel that multiplies the manifold kernel.

5. Limitations

LAT interpolation is difficult in part because of genuine discontinuities where wavefronts collide and where lines of block may exist, and also because of the quasi-monotonic nature of the underlying activation field. These difficulties are not solved with a GP, which may further introduce artefacts such as reverting to the mean in regions of extrapolation and oscillations in larger gaps between observations. Measurements of CV made by fitting plane waves to LAT observations can explicitly model a propagating wavefront, but the LAT interpolation here does not. Furthermore, CV calculations using LATs measured from individual placements of a multipolar catheter have the advantage that the distance between recording sites is fixed and known, whereas LAT interpolation and CV prediction from combining observations for many separate placements of a catheter (as done here) is potentially subject to large inaccuracies, i.e. some non-physiological CV predictions simply reflect the data, and it may be hard to correct for this without making strong prior assumptions about CV.

Accurately predicting the gradient of LAT is more challenging than predicting LAT, as noisy observations of LAT provide limited gradient information, particularly when sparsely spaced. Furthermore, the *magnitude* of the LAT gradient is not a GP. Nonetheless, we find that the calibration of the predictions (the coverage achieved by credible intervals averaged over all predictions) is reasonable for the posterior distribution of the LAT gradient magnitude.

6. Conclusion

We have introduced the GPML method which allows us to perform probabilistic interpolation directly on atrial manifolds using a reduced-rank GP that expresses the covariance function in terms of eigenfunctions of the Laplacian operator on a mesh triangulation of the manifold.

This allows us to interpolate uncertain LAT with a differentiable model, which then allows us to calculate probabilistic estimates of conduction velocity maps. In general, this method allows GPs to be used to model data on irregular spatial domains such as the left atrium, while directly accounting for the topology of the domain and the physically meaningful distances that correlations ought to depend upon.

Ethics. Clinical data were recorded as part of routine care for first-time ablation for atrial fibrillation. Ethical approval was granted by the National Research Ethics Service (10/H0802/77) for Patient 1, and (REC number 18/HRA/0083) for Patient 2.

Authors' contributions. S.C., R.D.W., J.E.O. and R.H.C. conceived the study. S.C. designed and implemented the main methods. C.C., C.H.R. and D.O.H. processed the raw clinical data. S.E.W. and M.D.N. performed the clinical data collection. All authors read and approved the manuscript.

Competing interests. The authors declare that they have no competing interests.

Funding. This work was funded by grants from the UK Engineering and Physical Sciences Research Council (EP/M012492/1, NS/A000049/1, EP/P01268X/1, EP/K037145/1), the British Heart Foundation (PG/15/91/31812, PG/13/37/30280), Kings Health Partners London National Institute for Health Research (NIHR) Biomedical Research Centre, the Wellcome/EPSRC Centre for Medical Engineering [WT 203148/Z/16/Z], Medical Research Council Skills Development Fellowship (MR/S015086/1), and funding for research from Abbott, EBR systems, Pfizer and Siemens.

Acknowledgements. We are grateful to Dr Mauricio Alvarez, Dr Arno Solin and Prof. Finn Lindgren for helpful discussions.

References

1. Cantwell CD, Roney CH, Ng FS, Siggers JH, Sherwin SJ, Peters NS. 2015 Techniques for automated local activation time annotation and conduction velocity estimation in cardiac mapping. *Comput. Biol. Med.* **65**, 229–242. (doi:10.1016/j.compbimed.2015.04.027)
2. Honarbakhsh S *et al.* 2019 Left atrial scarring and conduction velocity dynamics: rate dependent conduction slowing predicts sites of localized reentrant atrial tachycardias. *Int. J. Cardiol.* **278**, 114–119. (doi:10.1016/j.ijcard.2018.10.072)
3. Bellmann B *et al.* 2019 Velocity characteristics of atrial fibrillation sources determined by electrographic flow mapping before and after catheter ablation. *Int. J. Cardiol.* **286**, 56–60. (doi:10.1016/j.ijcard.2019.02.006)
4. Corrado C, Williams S, Karim R, Plank G, O'Neill M, Niederer S. 2018 A work flow to build and validate patient specific left atrium electrophysiology models from catheter measurements. *Med. Image Anal.* **47**, 153–163. (doi:10.1016/j.media.2018.04.005)
5. Boyle PM *et al.* 2019 Computationally guided personalized targeted ablation of persistent atrial fibrillation. *Nat. Biomed. Eng.* **3**, 870–879. (doi:10.1038/s41551-019-0437-9)
6. Verma B, Oesterlein T, Loewe A, Luik A, Schmitt C, Dössel O. 2018 Regional conduction velocity calculation from clinical multichannel electrograms in human atria. *Comput. Biol. Med.* **92**, 188–196. (doi:10.1016/j.compbimed.2017.11.017)
7. Loewe A, Poremba E, Oesterlein T, Luik A, Schmitt C, Seemann G, Dössel O. 2019 Patient-specific identification of atrial flutter vulnerability—a computational approach to reveal latent reentry pathways. *Front. Physiol.* **9**, 1910. (doi:10.3389/fphys.2018.01910)
8. Roney CH *et al.* 2019 A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction. *Comput. Biol. Med.* **104**, 278–290. (doi:10.1016/j.compbimed.2018.10.019)
9. Coveney S *et al.* 2019 Probabilistic interpolation of uncertain local activation times on humana atrial manifolds. *IEEE Trans. Biomed. Eng.* **67**, 99–109. (doi:10.1109/TBME.2019.2908486)
10. Lindgren F, Rue H, Lindström J. 2011 An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *J. R. Stat. Soc. B* **73**, 423–498. (doi:10.1111/j.1467-9868.2011.00777.x)
11. Rasmussen CE, Williams CKI, Sutton RS, Sutton AG, Spirtes P, Glymour C, Scheines R, Schölkopf B, Smola AJ. 2006 *Gaussian processes for machine learning*. Cambridge, MA: MIT Press.
12. Solin A, Särkkä S. 2019 Hilbert space methods for reduced-rank Gaussian process regression. *Stat. Comput.* **30**, 419–446. (doi:10.1007/s11222-019-09886-w)

13. Wood AT. 1995 When is a truncated covariance function on the line a covariance function on the circle? *Stat. Probab. Lett.* **24**, 157–164. (doi:10.1016/0167-7152(94)00162-2)
14. Solin A, Kok M. 2019 Know Your Boundaries: constraining Gaussian Processes by Variational Harmonic Features. *arXiv*, p. 1904.05207v1.
15. Sorkine O. 2005 Laplacian Mesh Processing, *Eurographics - State of the Art Reports*, no. Section 4, pp. 53–70.
16. Reuter M, Biasotti S, Giorgi D, Patanè G, Spagnuolo M. 2009 Discrete Laplace-Beltrami operators for shape analysis and segmentation. *Comput. Graph. (Pergamon)* **33**, 381–390. (doi:10.1016/j.cag.2009.03.005)
17. Virtanen P *et al.* 2019 SciPy 1.0—Fundamental Algorithms for Scientific Computing in Python. *arXiv e-prints*, p. arXiv:1907.10121, July
18. Coveney S. 2020 quLATi. <https://zenodo.org/record/3758043#Xp2s5nVKjeQ>. (doi:10.5281/zenodo.3758043)
19. Vigmond EJ, Hughes M, Plank G, Leon LJ. 2003 Computational tools for modeling electrical activity in cardiac tissue. *J. Electrocardiol.* **36** Suppl, 69–74. (doi:10.1016/j.jelectrocard.2003.09.017)
20. Corrado C *et al.* 2017 Personalized models of human atrial electrophysiology derived from endocardial electrograms. *IEEE Trans. Biomed. Eng.* **64**, 735–742. (doi:10.1109/TBME.2016.2574619)
21. Jidling C, Wahlström N, Wills A, Schön TB. 2017 Linearly constrained Gaussian processes, no. Nips.
22. Alvarado PA, Alvarez MA, Daza-Santacoloma G, Orozco A, Castellanos-Dominguez G. 2014 A latent force model for describing electric propagation in deep brain stimulation: a simulation study. *2014 36th Annual Int. Conf. of the IEEE Engineering in Medicine and Biology Society, EMBC 2014, Chicago, IL, 26–30 August*, pp. 2617–2620. (doi:10.1109/EMBC.2014.6944159).