

Physics-informed Gaussian Processes for Nonlinear Partial Differential Equations in a Fluid-Dynamics Application

L. Mihaela Paun¹, Mitchel J. Colebank², and Dirk Husmeier³

¹School of Mathematical Sciences, University of Southampton, Southampton, UK

²Department of Mathematics, University of South Carolina, Columbia, SC, USA

³School of Mathematics and Statistics, University of Glasgow, Glasgow, UK

Emails: L-M.Paun@soton.ac.uk; mjcolebank@sc.edu; Dirk.Husmeier@glasgow.ac.uk;

Abstract We investigate the novel application of a state-of-the-art physics-informed Gaussian Process method to a computational haemodynamics model described by nonlinear partial differential equations (PDEs), commonly used to describe blood flow in a network of blood vessels. We estimate and quantify the uncertainty of the PDE solution and the unknown PDE parameters in a Bayesian framework without resorting to computationally expensive evaluations of the PDEs. Our results indicate that with low noise, both the PDE parameter and solutions are well inferred.

Keywords: physics-informed machine learning, Gaussian processes, uncertainty quantification, partial differential equations, fluid-dynamics

1. Introduction

Computer simulations based on partial differential equations (PDEs) are powerful tools that enable understanding of many physical and engineering phenomena. Critical to these PDEs is the estimation of unknown parameters from measurements. Typically, methods for PDE parameter inference require repeated numerical evaluations of the PDEs a large number of times, which is computationally expensive, and thus often deemed impractical in real-world applications. To overcome this issue, a physics-informed machine learning approach can be taken, in which the PDE numerical evaluation is bypassed altogether. In the current study, we employ the PDE-informed Gaussian Process (PIGP) approach developed by Li et al. [1] to a fluid-dynamics application of the cardiovascular circulation that is based on the Navier-Stokes equations [2].

The idea behind PIGP is that the PDE solution is modelled as a Gaussian Process (GP). The PIGP method can handle nonlinear PDEs by the use of a PDE augmentation, whereby the nonlinear PDEs are reformulated as a system of equivalent *linear* PDEs in all derivatives. Besides estimating the unknown PDE parameters and solutions, the PIGP method provides uncertainty quantification (UQ) by constructing the joint posterior distribution of the PDE parameters, solution and measurement noise variance.

Here we are interested in the novel application of PIGP to an idealised model of the pulmonary blood circulation described by nonlinear PDEs. The model is able to predict features of the pulmonary haemodynamics, such as pulmonary blood flow, pressure or cross-sectional area under normal physiological and pathological conditions. This is of medical relevance as it allows pulmonary disease monitoring and treatment [3, 4].

We find that with low noise, both the PDE parameter and solutions are well estimated, and the uncertainty is quantified.

2. PIGP methodology

Literature studies have mainly employed GP models for PDE parameter inference in the context of emulation [5, 6]. More specifically, a GP emulator is employed to approximate the PDE solution based on a training data set obtained by numerically evaluating the PDEs a large number of times on a design. The disadvantage of this approach is that to build the training data set, the PDEs are numerically evaluated a large number of times, which is computationally burdensome for time-consuming PDEs. A second approach exists, where GPs are used as a prior on the PDE solution while the physics are naturally incorporated without relying on repeated expensive forward evaluations [1, 7]. The latter approach is adopted by the PIGP method. Most studies investigating the latter approach only consider linear PDEs [8, 7], and make use of the fact that the GP is closed under differentiation (a linear transformation of a GP is also a GP). However, for nonlinear PDEs, this property does not hold, and hence, a linearisation step is required. For instance, the study in [7] linearises the nonlinear PDE operator by a discretisation

of the PDE in the time domain before applying GP regression.

Below we proceed to describe the PIGP methodology [1] employed in this work.

Consider a PDE operator $\mathbb{A}(\mathbf{x}, u(\mathbf{x}), \boldsymbol{\theta})$ parameterized by the unknown PDE parameter $\boldsymbol{\theta}$, and PDE solution $u(\mathbf{x})$ for $\mathbf{x} \in \Omega$. We further consider a function $f(\mathbf{x}, u(\mathbf{x}), \boldsymbol{\theta})$, which is the zeroth-order term of the PDE, such that

$$\mathcal{A}(\mathbf{x}, u(\mathbf{x}), \boldsymbol{\theta}) = f(\mathbf{x}, u(\mathbf{x}), \boldsymbol{\theta}). \quad (1)$$

We assume that the PDE solution is observed with noise, hence $y_i = u(\mathbf{x}_i) + \epsilon_i$, where the errors ϵ_i , $i = 1, \dots, n$ are assumed Normally distributed, independent and identically distributed (iid), $\epsilon_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$. Furthermore, we let $\boldsymbol{\tau} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ denote the set of observation locations and/or times. PIGP aims to estimate and quantify the uncertainty of the PDE parameter $\boldsymbol{\theta}$, the PDE solution $u(\mathbf{x})$ and the noise variance σ^2 from data $\mathbf{y}(\boldsymbol{\tau})$ in a Bayesian framework.

As part of PIGP, the PDE solution $u(\mathbf{x})$ is placed a GP prior on, hence the random variable $U(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, where $\mu(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}')$ are the mean and covariance function of the GP.

Linear PDE operator: For a *linear* PDE operator denoted $\mathcal{L}_{\mathbf{x}}^{\boldsymbol{\theta}}$, $\mathcal{L}_{\mathbf{x}}^{\boldsymbol{\theta}}U(\mathbf{x})$ exists for a smooth enough mean and covariance function, and $\mathcal{L}_{\mathbf{x}}^{\boldsymbol{\theta}}U(\mathbf{x})$ is also a GP. Further, a new variable W is introduced,

$$W = \sup_{\mathbf{x} \in \Omega} \|\mathcal{L}_{\mathbf{x}}^{\boldsymbol{\theta}}U(\mathbf{x}) - f(\mathbf{x}, U(\mathbf{x}), \boldsymbol{\theta})\|_{\infty}, \quad (2)$$

which quantifies the difference between the prediction of the PDE solution and the PDE term. Hence, for a perfect prediction, $W \equiv 0$. However, in practice W cannot be computed exactly, and hence it is approximated by a finite discretisation set, i.e. $\mathbf{I} = \{\mathbf{x}_1, \dots, \mathbf{x}_{n_I}\}$ chosen such that $\boldsymbol{\tau} \subset \mathbf{I} \subset \Omega$, i.e. the discretisation set includes the set of observations. Hence, we define

$$W_I = \max_{\mathbf{x} \in \mathbf{I}} \|\mathcal{L}_{\mathbf{x}}^{\boldsymbol{\theta}}U(\mathbf{x}) - f(\mathbf{x}, U(\mathbf{x}), \boldsymbol{\theta})\|_{\infty}, \quad (3)$$

and for a sufficiently dense \mathbf{I} , $W_I \rightarrow W$. By denoting $\mathbf{u}(\mathbf{I}) = (u(\mathbf{x}_1), \dots, u(\mathbf{x}_{n_I}))$, we construct the joint posterior distribution

$$\begin{aligned} p(\mathbf{u}(\mathbf{I}), \boldsymbol{\theta}, \sigma^2 | W_I = 0, \mathbf{y}(\boldsymbol{\tau})) \\ \propto p(\mathbf{u}(\mathbf{I}) | \boldsymbol{\theta}) \times p(\boldsymbol{\theta}) \times p(\sigma^2) \times p(\mathbf{y}(\boldsymbol{\tau}) | \mathbf{u}(\mathbf{I}), \boldsymbol{\theta}, \sigma^2) \times p(W_I = 0 | \mathbf{u}(\mathbf{I}), \boldsymbol{\theta}, \sigma^2) \\ \propto p(\boldsymbol{\theta}) \times p(\sigma^2) \times \exp\{-0.5[n_I \log(2\pi) + \log(\det(\mathbf{C})) + \|\mathbf{u}(\mathbf{I}) - \boldsymbol{\mu}(\mathbf{I})\|_{\mathbf{C}^{-1}}]\} \\ \times \exp\{-0.5[n \log(2\pi) + n \log(\sigma^2) + \|\mathbf{u}(\boldsymbol{\tau}) - \mathbf{y}(\boldsymbol{\tau})\|_{\sigma^{-2}}]\} \\ \times \exp\{-0.5[n_I \log(2\pi) + \log(\det(\mathbf{K})) + \|f(\mathbf{I}, \mathbf{u}(\mathbf{I}), \boldsymbol{\theta}) - \mathcal{L}_{\mathbf{x}}^{\boldsymbol{\theta}}\boldsymbol{\mu}(\mathbf{I}) - \mathbf{m}\{\mathbf{u}(\mathbf{I}) - \boldsymbol{\mu}(\mathbf{I})\}\|_{\mathcal{K}^{-1}}]\}, \end{aligned} \quad (4)$$

where $\mathbf{C} = \mathbf{K}(\mathbf{I}, \mathbf{I})$, $\mathbf{m} = \mathcal{L}\mathbf{K}(\mathbf{I}, \mathbf{I})\mathbf{K}(\mathbf{I}, \mathbf{I})^{-1}$, $\mathcal{K} = \mathcal{L}\mathbf{K}\mathcal{L}(\mathbf{I}, \mathbf{I})\mathbf{K}(\mathbf{I}, \mathbf{I})^{-1}\mathbf{K}\mathcal{L}(\mathbf{I}, \mathbf{I})$, and $\mathbf{K}(\mathbf{I}, \mathbf{I})$ is an (n_I, n_I) covariance matrix with (i,j) element $k(\mathbf{x}_i, \mathbf{x}_j)$, $\mathcal{L}\mathbf{K}(\mathbf{I}, \mathbf{I})$ is a matrix with elements of the form $\mathcal{L}_{\mathbf{x}}^{\boldsymbol{\theta}}(k(\mathbf{x}, \mathbf{x}'))$, $\mathbf{K}\mathcal{L}(\mathbf{I}, \mathbf{I})$ is a matrix with elements of the form $\mathcal{L}_{\mathbf{x}'}^{\boldsymbol{\theta}}(k(\mathbf{x}, \mathbf{x}'))$, and $\mathcal{L}\mathbf{K}\mathcal{L}(\mathbf{I}, \mathbf{I})$ is a matrix of the form $\mathcal{L}_{\mathbf{x}}^{\boldsymbol{\theta}}(\mathcal{L}_{\mathbf{x}'}^{\boldsymbol{\theta}}(k(\mathbf{x}, \mathbf{x}')))$. The reader is referred to [1] for details.

Nonlinear PDE operator: For a *nonlinear* PDE operator, $\mathcal{A}(\mathbf{x}, u(\mathbf{x}), \boldsymbol{\theta})$ in eq (1) may not be Gaussian, and hence we cannot directly write down the joint posterior distribution in eq (4). To overcome this issue, the nonlinear PDE can be turned into a system of augmented linear PDEs, which are equivalent to the original nonlinear PDE. Thus, we write the nonlinear PDE in eq (1) as [1]:

$$\nabla^{\alpha_1} u = \mathcal{A}_1(\boldsymbol{\theta}, u, \nabla^{\alpha_2} u, \dots, \nabla^{\alpha_l} u) + f(\mathbf{x}, u(\mathbf{x}), \boldsymbol{\theta}), \quad (5)$$

where ∇^{α_i} is a partial derivative operator, \mathcal{A}_1 is nonlinear in its arguments and may contain parameter-dependent terms. The constructed augmented PDE has two key advantages which enable the PIGP methodology described above to be applied: (i) it is linear in all non-zeroth order derivatives, and (ii) the corresponding PDE operator is independent of the parameter $\boldsymbol{\theta}$. The latter is important for computational reasons, since the dependence of the PDE operator on $\boldsymbol{\theta}$ would imply the updating of \mathbf{K} , $\mathcal{L}\mathbf{K}$, $\mathbf{K}\mathcal{L}$, and $\mathcal{L}\mathbf{K}\mathcal{L}$ (eq (4)) for every $\boldsymbol{\theta}$ update within a Bayesian inference scheme, which would increase computational costs significantly.

By denoting $u_1 = u$, eq (5) is rewritten such that: $\nabla^{\alpha_2} u_1(\mathbf{x}) = u_2(\mathbf{x}), \dots, \nabla^{\alpha_l} u_1(\mathbf{x}) = u_l(\mathbf{x})$, and $\nabla^{\alpha_1} u_1(\mathbf{x}) = f(\mathbf{x}, u_1(\mathbf{x}), \theta) + \mathcal{A}_1(\mathbf{x}, \theta, u_1(\mathbf{x}), u_2(\mathbf{x}), \dots, u_l(\mathbf{x}))$. Further, by placing independent GPs on U_1, \dots, U_l , the joint posterior distribution of $u_1, \dots, u_l, \theta, \sigma^2$ can be derived, and takes a similar form to that in eq (4). The reader is referred to [1] for further details.

3. Fluid-dynamics application

We apply the PIGP methodology to a computational haemodynamics model commonly used to describe blood flow in a network of blood vessels [2]. The computational framework is a reduced form of the full Navier-Stokes equation in cylindrical coordinates, only accounting for axial dynamics down a blood vessel and hence it is one-dimensional (1D) in space. While the model is considered low-fidelity in its description of spatial fluid flow, it is computationally efficient and enables the use of more physiological stress-strain dynamics, where vessel walls can deform under load [2].

3.1. Nonlinear PDEs

The fluid-dynamics model is determined by a set of coupled nonlinear PDEs. The equations simulate blood flow, $q(x, t)$, blood pressure, $p(x, t)$ and cross-sectional area $A(x, t) = \pi(R(x, t))^2$ at location x and time t by

$$\frac{\partial A}{\partial t} + \frac{\partial q}{\partial x} = 0, \quad \frac{\partial q}{\partial t} + \frac{\partial}{\partial x} \left(\frac{q^2}{A} \right) + \frac{A}{\rho} \frac{\partial p}{\partial x} = -\frac{2\pi\nu R}{\delta} \frac{q}{A}. \quad (6)$$

For notational simplicity, we have dropped the indices (x, t) . Here $0 \leq x \leq L$ (cm) is the axial position along the vessel, μ (g/cm-s) is the viscosity, ρ (g/cm³) is blood density, $\nu = \mu/\rho$ (cm²/s) is the kinematic viscosity, $R(x, t)$ (cm) is the radius, and t (s) is time. The axial blood velocity, $u_x = q/A$, has a shape profile given by the Stokes boundary layer

$$u_x(r, x, t) = \begin{cases} \bar{u}_x, & r < R - \delta, \\ \bar{u}_x \frac{(R-r)}{\delta}, & R - \delta < r \leq R. \end{cases} \quad (7)$$

Here $\delta = \sqrt{\nu T/2\pi}$ (cm) is the boundary layer thickness, T (s) is the duration of the cardiac cycle, and \bar{u}_x is the axial velocity outside of the boundary layer towards the center of the vessel. Finally, we define a pressure-area relationship of the form

$$p(x, t) = p_0 + \frac{4}{3} \frac{Eh}{r_0} \left(1 - \sqrt{\frac{A_0}{A}} \right), \quad \frac{Eh}{r_0} = k_1 \exp(-k_2 r_0) + k_3, \quad (8)$$

where E (g/cm-s²) is Young's modulus, h (cm) is the vessel wall thickness, p_0 (g/cm-s²) is the reference pressure corresponding to the reference area $A_0 = \pi r_0^2$ (cm²). For our analysis, we consider a single blood vessel, and impose a dynamic inflow at the start of the vessel $q(0, t)$ and a constant pressure at the end point of the vessel $p(L, t)$ as boundary conditions. We use this as a starting point for method development, with further applications for models that include a network of blood vessels [2, 3].

3.2. Augmented linear PDEs

Equations (6) and (8) contain three variables, area $A(x, t)$, flow $q(x, t)$ and pressure $p(x, t)$. We reduce the system by expressing $p(x, t)$ in terms of $A(x, t)$. Further, by denoting $A = A_1$ and $q = q_1$, and after dropping the indices (x, t) for

notational simplicity, we reformulate the system of nonlinear PDEs in eqs (6) and (8) by the equivalent linear PDEs:

$$\begin{aligned}
\frac{\partial A_1}{\partial x} &= A_2, \\
\frac{\partial q_1}{\partial x} &= q_2, \\
\frac{\partial A_1}{\partial t} &= -q_2, \\
\frac{\partial q_1}{\partial t} &= -2\frac{q_1}{A_1}q_2 + \left(\frac{q_1}{A_1}\right)^2 A_2 - \frac{2}{3\rho} \frac{Eh}{r_0} \sqrt{\frac{A_0}{A_1}} A_2 - \frac{2\pi\nu R}{\delta} \frac{q_1}{A_1},
\end{aligned} \tag{9}$$

where $R = \sqrt{\frac{A_1}{\pi}}$. The PIGP method will therefore infer the PDE solutions $A(x, t)$ and $q(x, t)$; $p(x, t)$ can then be obtained from $A(x, t)$ using the first equation in (8). We also learn the parameter Eh/r_0 by assuming $k_1 = 0$ and $k_2 = 0$, and only infer k_3 from the second equation in (8). The inference data consist of observations for both variables $A(x, t)$ and $q(x, t)$ at various location points x and time points t .

4. Simulation study

We consider 225 data points, 15 in each direction, x and t , hence $\tau = \{(x_i, t_i), i = 1, \dots, n\}$, where $n = 225$. To mimic the real data collection, we take t to be evenly spaced between $[0, T]$, where T is the time interval (0.9 s). In addition, we take x evenly spaced between $[0, L]$, where $L = 50$ (mm) is the vessel length. We add noise to the data, $\epsilon_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$, where $\sigma^2 = 10^{-6}$. These data act as training data for GP fitting. We use a product Matérn kernel with the degree of freedom set to 2.1, see [1] for details.

For the purpose of this simulation study, we assume no knowledge of the boundary conditions, which is a scenario often encountered in practice, in order to test the method's ability to infer the parameter and PDE solution without access to the boundary conditions. In addition, we take the number of discretisation points n_I to be equal to the number of observations n .

Lastly, we obtain 20,000 Hamiltonian Monte Carlo (HMC) samples from the generalisation of the posterior distribution in eq (4) for multiple GPs, see [1] for details. As a starting point we use the maximum a posteriori point (MAP) found using the LBFGS algorithm to speed up computations, and using the default step size of 10^{-5} and the number of leapfrog steps of 200 [1]. For the PDE parameters we use a uniform prior within the physiologically chosen range $[7.5, 8.5] \times 10^4$, and for the noise variance we use Jeffrey's prior [1].

5. Results

We find MAP of the unknown PDE parameter and PDE solutions $A(x, t)$ and $q(x, t)$ by optimising their joint posterior distribution (while keeping the noise variance σ^2 fixed at its true value) in eq (4). In Figure 1(a), we show the optimisation trajectory of the PDE (stiffness) parameter. We notice that the parameter converges towards the true value, and its final value (7999) is only decimals away from the true value (80000), which is expected given the low noise variance considered.

To quantify the uncertainty in the estimation, we run HMC, and we infer the PDE parameter, the PDE solutions and the noise variance jointly. In Figure 1(b), we display the marginal posterior distribution of the PDE parameter, which contains the true parameter value. It also becomes obvious that compared to the prior distribution (a uniform within the imposed parameter range), the posterior is much narrower (posterior uncertainty is reduced compared to prior uncertainty), suggesting informativeness of the data about the parameter.

In addition, in Figures 2 and 3 we show the PIGP predictions for the PDE solutions of A and q at various test locations x , and time t as given by the posterior mean. We also show the 95% credible interval for the PDE solutions. Firstly, we notice that both A and q are well predicted, except A at the final x , i.e. at the end of the vessel. We can understand why the area estimation is poor here by analysing Figure 4, which shows the PIGP predictions against the training data for both area and flow, as well as the difference between the observations and predictions. We see that at the end point the area is almost constant, which contrasts with the dynamic behaviour in the rest of the locations. The GP is less accurate at the end point because of the lower signal amplitude relative to the other locations, which is a consequence of the constant pressure outlet

boundary conditions. We expect that the accuracy would increase if boundary conditions (i.e. outlet constant area and inlet dynamic flow) were included in the PIGP. This constitutes future work.

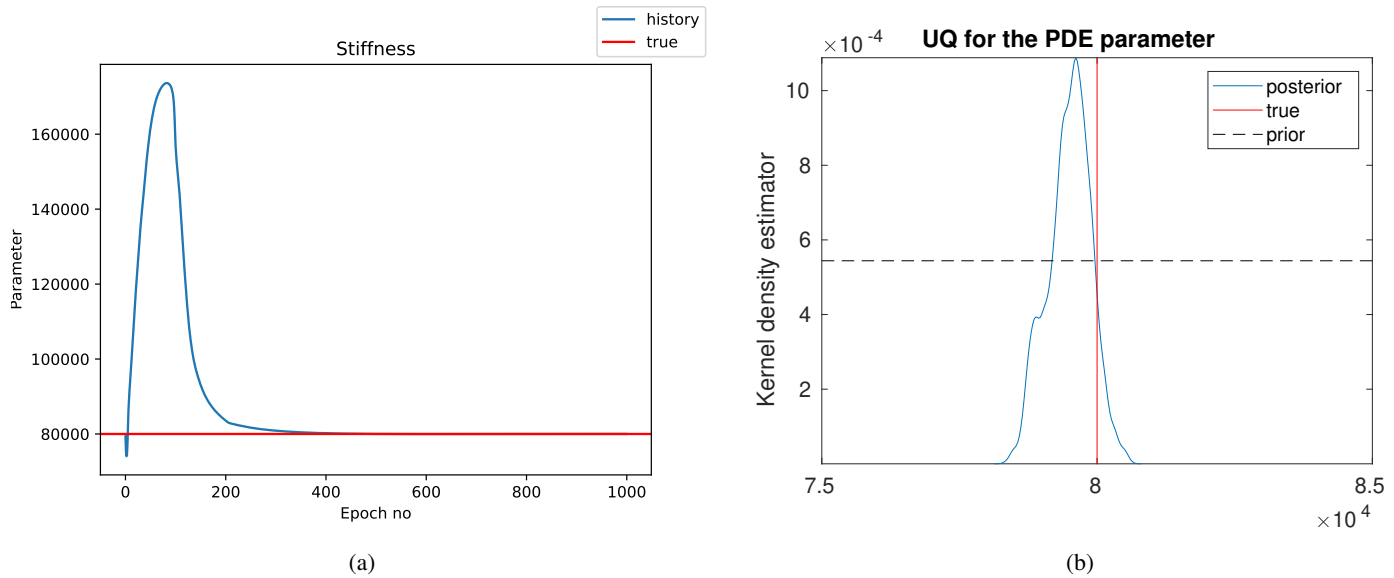


Fig. 1: Panel (a): Optimisation history of the PDE (stiffness) parameter: the maximum aposteriori point is seen to converge to the true parameter value, indicated in red continuous (horizontal) line. Panel(b): Kernel density estimation showing the uncertainty in the estimated PDE parameter. We superimpose in red vertical line the true parameter value, and also show the prior as a uniform distribution.

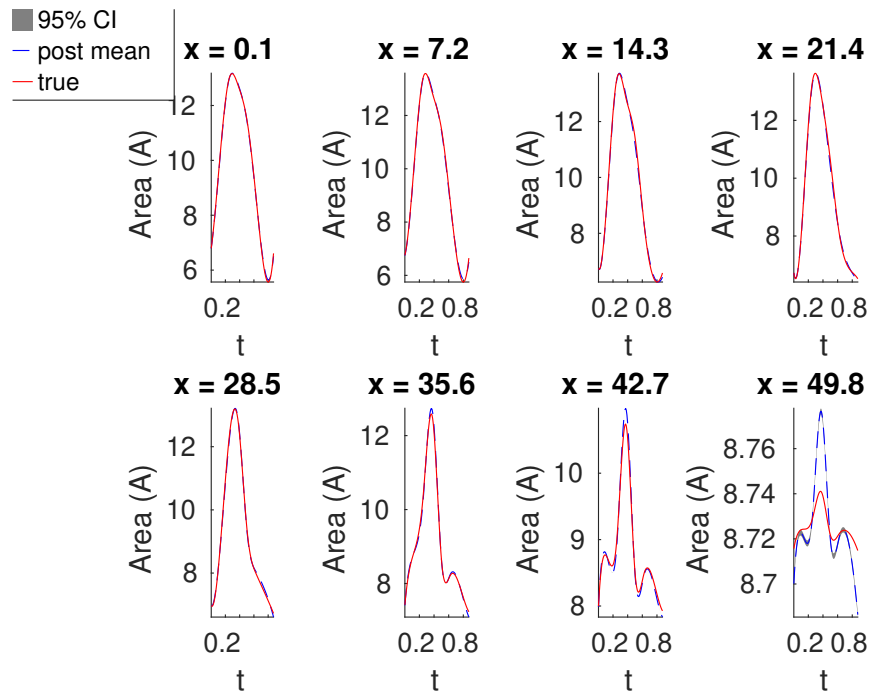


Fig. 2: Out-of-sample estimation and uncertainty of PDE solution for area obtained with the PIGP method. We show the true signal in red continuous line, the PIGP-prediction (posterior mean) in blue dashed line and the 95% credible interval as grey-shaded, obtained with HMC sampling.




Fig. 3: Out-of-sample estimation and uncertainty of PDE solution for flow obtained with the PIGP method. We show the true signal in red continuous line, the PIGP-prediction (posterior mean) in blue dashed line and the 95% credible interval as grey-shaded, obtained with HMC sampling.

6. Conclusions and future work

In this work we have investigated the novel application of a state-of-the-art physics-informed Gaussian Process method to a computational haemodynamics model described by nonlinear PDEs, commonly used to describe blood flow in a network of blood vessels. The method estimates the PDE solution and unknown PDE parameters without computationally burdensome evaluations of the PDEs, while also quantifying the uncertainty of the estimation in a Bayesian framework. Our results indicate that with low noise, both the PDE parameter and solutions are well inferred.

Future work extensions include: (1) investigations into a larger noise variance, as well as understand why the noise variance is underestimated; (2) incorporation of partial boundary conditions into the PIGP framework, such as the inflow $q(0, t)$, as well as more realistic distal boundary conditions, such as Windkessel or structured tree boundary conditions [3, 4]; (3) increasing the vessel network size and incorporating the bifurcation boundary conditions into the PIGP framework; (4) incorporation of partial observations (either flow or area, but not both), consistent with actual data collection; (5) inference of multiple parameters, such as k_1 or k_2 , or boundary condition parameters. We will also compare these methods with physics-informed neural networks (PINNs) in terms of (i) accuracy in solving inverse problems and (ii) their ability to capture posterior uncertainty in both parameter and output space [7].

Acknowledgements

I thank my newborn son Dominic for giving me the time to write this conference article.




Fig. 4: In-sample PIGP predictions of area (panel (a)) and flow (panel (c)) with training data superimposed, showing good agreement. We show the signals concatenated for all locations, e.g. the first signal ranging from 0 to 15 on the x-axis corresponds to the first location ($x = 0$) for all time points t , while the last signal ranging from 211 to 225 on the x-axis corresponds to the last location ($x = L$) for all time points t . Panels (b) and (d) show the difference between observations and predictions corresponding to panels (a) and (c), respectively.

References

- [1] Z. Li, S. Yang, and C. F. J. Wu, “Parameter Inference Based on Gaussian Processes Informed by Nonlinear Partial Differential Equations,” *SIAM/ASA Journal on Uncertainty Quantification*, vol. 12, no. 3, pp. 964–1004, 2024.
- [2] A. Taylor-LaPole, L. M. Paun, D. Lior, J. D. Weigand, C. Puelz, and O. S. Mette, “Parameter selection and optimization of a computational network model of blood flow in single-ventricle patients,” *J R Soc Interface*, vol. 22, no. 223, p. 20240663, 2025.
- [3] L. M. Paun, M. J. Colebank, M. S. Olufsen, N. A. Hill, and D. Husmeier, “Assessing model mismatch and model selection in a Bayesian uncertainty quantification analysis of a fluid-dynamics model of pulmonary blood circulation,” *Journal of The Royal Society Interface*, vol. 17, p. 20200886, 12 2020.
- [4] L. M. Paun, M. J. Colebank, A. Taylor-LaPole, M. S. Olufsen, W. Ryan, I. Murray, J. M. Salter, V. Applebaum, M. Dunne, J. Hollins, L. Kimpton, V. Volodina, X. Xiong, and D. Husmeier, “SECRET: Statistical Emulation for Computational

- Reverse Engineering and Translation with applications in healthcare,” *Computer Methods in Applied Mechanics and Engineering*, vol. 430, p. 117193, 2024.
- [5] S. Lan, T. Bui-Thanh, M. Christie, and M. Girolami, “Emulation of higher-order tensors in manifold Monte Carlo methods for Bayesian Inverse Problems,” *Journal of Computational Physics*, vol. 308, pp. 81–101, 2016.
- [6] L. M. Paun, M. J. Colebank, and D. Husmeier, “A comparison of Gaussian processes and polynomial chaos emulators in the context of haemodynamic pulse-wave propagation modelling,” *Philos Trans A Math Phys Eng Sci.*, vol. 383, no. 2292, p. 20240222, 2025.
- [7] G. Pang and G. E. Karniadakis, *Physics-Informed Learning Machines for Partial Differential Equations: Gaussian Processes Versus Neural Networks*, pp. 323–343. Cham: Springer International Publishing, 2020.
- [8] D. Dalton, A. Lazarus, H. Gao, and D. Husmeier, “Boundary constrained gaussian processes for robust physics-informed machine learning of linear partial differential equations,” *Journal of Machine Learning Research*, vol. 25, no. 272, pp. 1–61, 2024.