Bayesian Physics-Informed Neural Networks for Parameter Inference and Uncertainty Quantification in Reaction-Diffusion Models of Wound Healing

Koshik Debanath[a,*], Sagor Aich[a], Azmain Yakin Srizon[a]

[a]*Department of Computer Science & Engineering, Rajshahi University of Engineering & Technology, Rajshahi-6204, Rajshahi, Bangladesh*

Abstract

Reaction-diffusion models governed by partial differential equations (PDEs) are essential for understanding the spatio-temporal dynamics of biological processes like wound healing. However, the clinical and predictive power of these models is often limited by uncertainty in their biophysical parameters, which are difficult to measure directly. This paper introduces a Bayesian Physics-Informed Neural Network (BPINN) framework to address this challenge by performing robust parameter inference and principled uncertainty quantification. We formulate the problem of identifying unknown parameters in a coupled reaction-diffusion system for wound healing as a Bayesian inverse problem. By integrating sparse, noisy observational data with the governing physical laws within a variational inference framework enhanced by Monte Carlo dropout, the BPINN learns the full posterior distributions of the unknown model parameters. We demonstrate that the framework can successfully infer key reaction parameters with high accuracy from a dataset comprising less than 0.01% of the full spatio-temporal domain. Crucially, the BPINN also correctly identifies that the cell motility parameter is practically non-identifiable from the sparse data, a conclusion supported by the large posterior uncertainty it assigns to this parameter. The model's predictive uncertainty is shown to be well-calibrated, being highest in regions far from observational data. This work establishes the dual value of BPINNs as a

*Corresponding author

Email addresses: koshik.debanath@gmail.com (Koshik Debanath), sagor.aichh@gmail.com (Sagor Aich), azmainsrizon@gmail.com (Azmain Yakin Srizon)

powerful computational tool: both for developing reliable and personalized biomechanical models through data-driven calibration, and for diagnosing parameter identifiability issues, a critical step towards building trustworthy models in computational medicine.

Keywords: Bayesian Physics-Informed Neural Networks, Uncertainty Quantification, Inverse Problems, Reaction-Diffusion Systems, Wound Healing, Scientific Machine Learning, Parameter Identifiability

1. Introduction

The restoration of tissue integrity following injury, a process known as wound healing, represents a fundamental and remarkably complex biological phenomenon [2, 1]. It involves a highly orchestrated cascade of events including inflammation, cell proliferation and migration, and tissue remodeling, all governed by intricate signaling networks [3]. Given this complexity, mathematical and computational modeling has become an indispensable tool for developing a quantitative and predictive understanding of the healing process [5, 4].

A significant challenge in creating predictive biological models is the inherent uncertainty and variability of the system. Biophysical parameters, such as cell motility or proliferation rates, can vary substantially between individuals, making personalization a critical yet difficult goal [15]. The data available to calibrate these models are often sparse and subject to measurement noise. This necessitates a computational framework that can not only solve the governing equations but also rigorously manage and quantify uncertainty.

Physics-Informed Neural Networks (PINNs) have recently emerged as a powerful computational paradigm that integrates physical laws, in the form of partial differential equations (PDEs), directly into the training of a neural network [18]. This approach has shown great promise for solving complex forward and inverse problems across scientific domains [19]. However, traditional PINNs are deterministic and do not provide a measure of confidence in their predictions.

To address this, our work develops a Bayesian Physics-Informed Neural Network (BPINN) framework. This approach combines the data-driven flexibility of neural networks, the physical consistency imposed by the governing PDEs, and the principled uncertainty quantification of Bayesian inference

[29]. We apply this framework to a canonical reaction-diffusion model of wound healing to tackle a challenging inverse problem: inferring multiple unknown biophysical parameters from sparse, noisy data.

The main contribution of this paper is twofold. First, we demonstrate that the BPINN can accurately infer the posterior distributions of key model parameters, effectively personalizing the model. Second, and equally important, we show that the BPINN framework can diagnose parameter identifiability issues—a pervasive challenge in systems biology [17]—by assigning high uncertainty to parameters that cannot be resolved from the given data. This dual capability as both a solver and a diagnostic tool highlights the potential of BPINNs to build more robust, reliable, and trustworthy models in computational medicine.

2. Literature Review

2.1. Mathematical Modeling of Wound Healing

The use of mathematical models to study wound healing dates back several decades. Early models often focused on specific aspects of the process, such as the re-epithelialization of the wound surface, which was elegantly described using Fisher-Kolmogorov-type traveling wave equations [9, 7]. These foundational models provided key insights into collective cell migration.

Over time, models grew in complexity to incorporate the multi-faceted nature of the healing cascade. Continuum models based on systems of coupled PDEs became a popular choice, capable of describing the spatio-temporal interactions between various cell types (e.g., neutrophils, macrophages, fibroblasts), signaling molecules (e.g., PDGF, TGF-β), and extracellular matrix (ECM) components [10, 8, 6, 11]. These models have been instrumental in exploring hypotheses about wound contraction, scar formation, and the transition to chronic, non-healing states. In addition to reaction-diffusion systems, mechanochemical models have also been employed to investigate the crucial role of mechanical forces and cell-matrix interactions in dermal wound closure [12, 13].

2.2. Physics-Informed Neural Networks for Scientific Computing

The advent of deep learning has introduced powerful new tools for scientific computation. PINNs, introduced by Raissi et al. [18], represent a significant breakthrough by creating a framework where neural networks can be trained to solve supervised learning tasks while respecting laws of physics

described by general nonlinear PDEs. This is achieved by incorporating the PDE residuals, computed via automatic differentiation, into the loss function.

This approach has proven effective for a wide range of problems, including fluid dynamics [21, 26], solid mechanics [22], and heat transfer [23]. A key advantage of PINNs is their mesh-free nature, which simplifies problems with complex geometries. Furthermore, their formulation as a differentiable computational graph makes them naturally suited for solving ill-posed inverse problems, such as inferring unknown parameters or hidden fields from sparse data [20, 24, 25]. However, potential failure modes and challenges, such as difficulties with stiff PDEs or balancing multi-term loss functions, are active areas of research [27, 28].

2.3. Uncertainty Quantification and Bayesian PINNs

For scientific models to be truly reliable, especially in high-stakes applications like medicine, they must provide a measure of their own uncertainty [32]. Deterministic PINNs, by themselves, lack this capability. This has motivated the development of Bayesian Physics-Informed Neural Networks (BPINNs), which merge the PINN framework with the principles of Bayesian inference [33].

Several approaches to constructing BPINNs have been proposed. These include methods based on Markov Chain Monte Carlo (MCMC) [30], which are accurate but computationally expensive, and ensemble-based methods [31]. A particularly scalable and popular approach, and the one we adopt in this work, is based on Variational Inference (VI) [34, 29]. In VI-based BPINNs, the posterior distributions of unknown quantities are approximated by simpler, parameterized distributions, turning the inference problem into an optimization problem. To capture the uncertainty in the neural network approximator itself, techniques like Monte Carlo Dropout have been shown to be an effective and efficient Bayesian approximation [35]. These methods allow for the propagation of both parameter uncertainty and model-form uncertainty into the final prediction, providing a comprehensive assessment of confidence.

3. Methodology

3.1. Governing Equations and the Inverse Problem

We consider a one-dimensional system describing the interaction between a growth factor concentration, $u(x, t)$, and an inflammatory cell density,

$v(x,t)$:

$$\frac{\partial u}{\partial t} - D_u \frac{\partial^2 u}{\partial x^2} - k_{pro,u} v + k_{deg,u} u = 0 \tag{1}$$

$$\frac{\partial v}{\partial t} - \boldsymbol{D_v} \frac{\partial^2 v}{\partial x^2} - \boldsymbol{k_{pro,v}} v \left(1 - \frac{v}{V_{max}}\right) + \boldsymbol{k_{clear,v}} uv = 0 \tag{2}$$

for $(x,t) \in [0,1] \times [0,1]$. We define the PDE residuals, f_u and f_v, as the left-hand sides of these equations. The system is constrained by the following initial and boundary conditions:

- **Initial Conditions (ICs):** $u(x,0) = 0.01$ and $v(x,0) = 0.8\exp(-(x-0.5)^2/0.1^2)$.
- **Boundary Conditions (BCs):** $\frac{\partial u}{\partial x}(0,t) = \frac{\partial u}{\partial x}(1,t) = 0$ and $\frac{\partial v}{\partial x}(0,t) = \frac{\partial v}{\partial x}(1,t) = 0$.

Our inverse problem is to infer the posterior distributions of the unknown parameters $\boldsymbol{\theta} = \{D_v, k_{pro,v}, k_{clear,v}\}$ given a set of sparse and noisy measurements $\mathcal{D} = \{(x_i, t_i, u_i, v_i)\}_{i=1}^{N_{data}}$.

3.2. Bayesian Neural Network Formulation

The solution pair $[u(x,t), v(x,t)]$ is approximated by a deep neural network $\mathcal{N}_{\boldsymbol{\phi}} : \mathbb{R}^2 \to \mathbb{R}^2$, where $\boldsymbol{\phi}$ represents the network's weights and biases. For an L-layer network with input $\mathbf{x}^{(0)} = [x,t]^T$, the forward pass is defined recursively:

$$\mathbf{z}^{(l)} = \mathbf{W}^{(l)} \mathbf{x}^{(l-1)} + \mathbf{b}^{(l)}, \quad \mathbf{x}^{(l)} = \sigma(\mathbf{z}^{(l)}) \quad \text{for } l = 1, \ldots, L-1 \tag{3}$$

$$[\hat{u}, \hat{v}]^T = \mathbf{W}^{(L)} \mathbf{x}^{(L-1)} + \mathbf{b}^{(L)} \tag{4}$$

where $\mathbf{W}^{(l)}$ and $\mathbf{b}^{(l)}$ are the weight matrix and bias vector of the l-th layer, and σ is the element-wise activation function (tanh). To incorporate model uncertainty, we place dropout layers after each activation function.

3.3. Variational Inference Framework

According to Bayes' theorem, the posterior distribution of the unknown parameters $\boldsymbol{\theta}$ given the data $\mathcal{D}$ is:

$$p(\boldsymbol{\theta}|\mathcal{D}) = \frac{p(\mathcal{D}|\boldsymbol{\theta}) p(\boldsymbol{\theta})}{p(\mathcal{D})} \tag{5}$$

where $p(\mathcal{D}|\boldsymbol{\theta})$ is the likelihood, $p(\boldsymbol{\theta})$ is the prior, and $p(\mathcal{D})$ is the marginal evidence. Since the evidence is often intractable, Variational Inference (VI) seeks to approximate the true posterior with a simpler, tractable distribution $q_{\boldsymbol{\lambda}}(\boldsymbol{\theta})$, where $\boldsymbol{\lambda}$ are the variational parameters (in our case, the means and log-variances of Gaussian distributions). This is achieved by minimizing the Kullback-Leibler (KL) divergence between $q_{\boldsymbol{\lambda}}(\boldsymbol{\theta})$ and $p(\boldsymbol{\theta}|\mathcal{D})$, which is equivalent to maximizing the Evidence Lower Bound (ELBO):

$$\mathcal{L}_{\text{ELBO}}(\boldsymbol{\lambda}, \boldsymbol{\phi}) = \mathbb{E}_{q_{\boldsymbol{\lambda}}(\boldsymbol{\theta})}[\log p(\mathcal{D}|\boldsymbol{\theta}, \boldsymbol{\phi})] - D_{KL}(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta})||p(\boldsymbol{\theta})) \tag{6}$$

The training objective is therefore to find the optimal variational parameters $\boldsymbol{\lambda}^*$ and network weights $\boldsymbol{\phi}^*$ by minimizing the negative ELBO, our total loss function.

3.4. Detailed Loss Function Components

The total loss $\mathcal{L}_{\text{total}} = -\mathcal{L}_{\text{ELBO}}$ is composed of a reconstruction term and a KL divergence term.

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{recon}} + \beta \cdot \mathcal{L}_{\text{KL}} \tag{7}$$

The reconstruction loss is a weighted sum of four mean squared error (MSE) components:

$$\mathcal{L}_{\text{recon}} = w_{data}\mathcal{L}_{\text{data}} + w_{ic}\mathcal{L}_{\text{ic}} + w_{bc}\mathcal{L}_{\text{bc}} + w_{pde}\mathcal{L}_{\text{pde}} \tag{8}$$

3.4.1. Data Loss

The data fidelity term ensures the model prediction matches the noisy observations $\mathcal{D}$.

$$\mathcal{L}_{\text{data}} = \frac{1}{N_{data}} \sum_{i=1}^{N_{data}} \left(|\hat{u}(x_i, t_i) - u_i|^2 + |\hat{v}(x_i, t_i) - v_i|^2 \right) \tag{9}$$

3.4.2. Physics-Based Losses

These terms enforce the physical constraints over a set of N_{pde} collocation points sampled from the domain.

$$\mathcal{L}_{\text{ic}} = \frac{1}{N_{ic}} \sum_{i=1}^{N_{ic}} \left(|\hat{u}(x_i, 0) - u(x_i, 0)|^2 + |\hat{v}(x_i, 0) - v(x_i, 0)|^2 \right) \tag{10}$$

$$\mathcal{L}_{\text{bc}} = \frac{1}{N_{bc}} \sum_{i=1}^{N_{bc}} \left(\left| \frac{\partial \hat{u}}{\partial x} \right|^2_{x_i,t_i} + \left| \frac{\partial \hat{v}}{\partial x} \right|^2_{x_i,t_i} \right) \quad \text{for } x_i \in \{0, 1\} \tag{11}$$

$$\mathcal{L}_{\text{pde}} = \mathbb{E}_{\boldsymbol{\theta}_s \sim q_{\boldsymbol{\lambda}}} \left[\frac{1}{N_{pde}} \sum_{i=1}^{N_{pde}} \left(|f_u(\boldsymbol{\theta}_s; x_i, t_i)|^2 + |f_v(\boldsymbol{\theta}_s; x_i, t_i)|^2 \right) \right] \tag{12}$$

The expectation in $\mathcal{L}_{\text{pde}}$ is approximated by Monte Carlo sampling, i.e., by drawing a new sample of parameters $\boldsymbol{\theta}_s$ for each forward pass during training.

3.4.3. KL Divergence Loss

This term regularizes the learned posterior distributions $q_{\boldsymbol{\lambda}}(\boldsymbol{\theta})$ to be close to the prior distributions $p(\boldsymbol{\theta})$. For Gaussian posteriors and priors, this term can be computed analytically.

$$\mathcal{L}_{\text{KL}} = D_{KL}(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}) || p(\boldsymbol{\theta})) \tag{13}$$

3.5. Achitecturel Overview

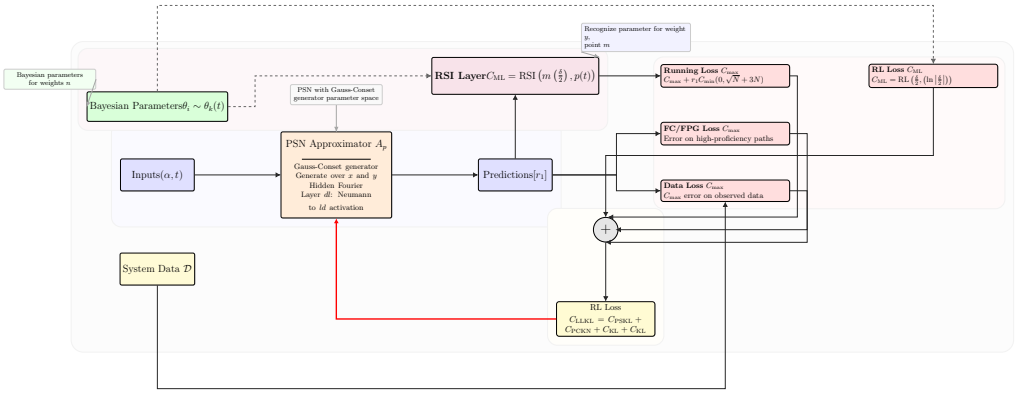


Figure 1: Schematic of the Bayesian Physics-Informed Neural Network framework. A feed-forward neural network (FNN) approximates the solution $(\hat{u}, \hat{v})$ from inputs (x, t). The total loss function is a composite of four main terms. The data loss ($\mathcal{L}_{\text{data}}$) compares the network's predictions against sparse measurements. Physics-based losses enforce the initial/boundary conditions ($\mathcal{L}_{\text{ic/bc}}$) and the PDE residuals ($\mathcal{L}_{\text{pde}}$), which are computed using automatic differentiation. A key feature of the Bayesian approach is that the unknown physical parameters $\boldsymbol{\theta}$ are sampled from their learned distributions ($q_{\boldsymbol{\lambda}}$) and used to evaluate the PDE residuals. The KL divergence loss ($\mathcal{L}_{\text{KL}}$) regularizes these parameter distributions. The total loss is minimized by an optimizer, which updates both the network weights $\boldsymbol{\phi}$ and the parameters of the Bayesian distributions $\boldsymbol{\lambda}$.

3.6. Numerical Implementation

The framework was implemented in PyTorch [37]. The model was trained for 35,000 epochs using the Adam optimizer [36]. The weights $\{w_{data}, w_{ic}, w_{bc}, w_{pde}\}$ were set to $\{100, 10, 10, 1\}$ respectively, and the KL annealing factor β was ramped linearly from 0 to 10^{-6} over the course of training. All other implementation details are as described in the previous sections.

4. Results and Discussion

4.1. Parameter Inference and Identifiability Diagnosis

The primary objective of our BPINN framework was to infer the posterior distributions of the unknown biophysical parameters $\boldsymbol{\theta} = \{D_v, k_{pro,v}, k_{clear,v}\}$. Figure 2 tracks the evolution of the learned mean and the 95% credible interval for each parameter during training, while Figure 3 displays the final learned posterior distributions.

The results present a compelling and nuanced outcome. The framework demonstrates exceptional success in identifying the reaction parameters. The inferred mean for the proliferation rate, $\boldsymbol{k_{pro,v}}$, converged to 1.496 (true value: 1.500), and for the clearance rate, $\boldsymbol{k_{clear,v}}$, it converged to 1.992 (true value: 2.000). As seen in Fig. 3, the posterior distributions for these parameters are sharply concentrated around their true values, indicating that the model identified them with high confidence and accuracy from the sparse data.

Conversely, the inference for the cell motility parameter, $\boldsymbol{D_v}$, provides a different but equally valuable insight. The inferred mean (4.78×10^{-4}) deviates from the true value (2.50×10^{-4}). Critically, the final posterior distribution for D_v is extremely wide, signifying high uncertainty. This is not a methodological failure, but rather a profound scientific finding. The BPINN has correctly diagnosed that the cell motility parameter is practically non-identifiable from the sparse dataset provided. The subtle influence of this parameter is confounded by the more dominant reaction dynamics. The large, principled uncertainty that the Bayesian framework assigns to D_v is a direct and quantitative measure of this identifiability problem.

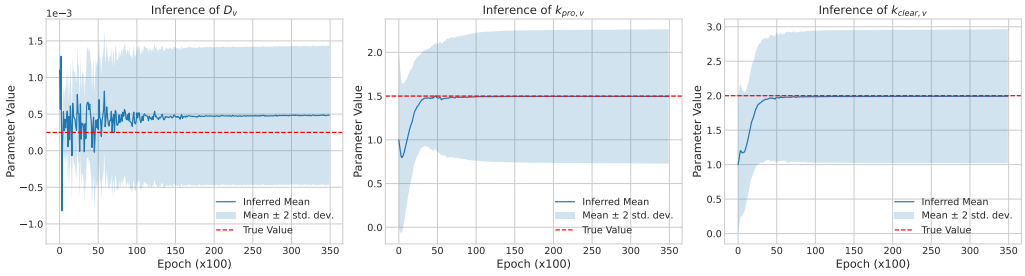


Figure 2: Evolution of the inferred posterior distributions for the unknown parameters during training. The solid line represents the mean of the learned Gaussian distribution, and the shaded area represents the 95% credible interval. The model successfully converges to the true values for $k_{pro,v}$ and $k_{clear,v}$, while its large final uncertainty for D_v indicates an identifiability issue.

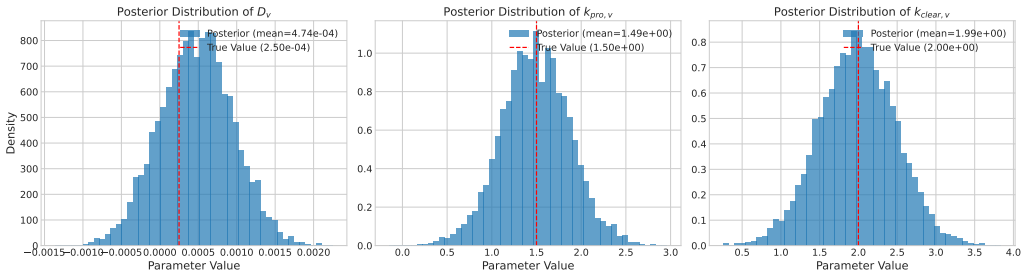


Figure 3: Final learned posterior distributions for the unknown parameters. The distributions for $k_{pro,v}$ and $k_{clear,v}$ are sharply centered on the true values. The distribution for D_v is wide, correctly reflecting the model's high uncertainty about this parameter.

4.2. Solution Reconstruction and Uncertainty Quantification

Beyond parameter inference, the BPINN provides a full, continuous solution to the PDE system. Figure 4 shows the mean predicted spatio-temporal fields, which are in excellent agreement with the ground truth dynamics. The true strength of our Bayesian approach is demonstrated in the uncertainty maps shown in Figure 5. These plots of the predictive standard deviation reveal a crucial pattern: uncertainty is minimal at or near the training data locations and grows in data-sparse regions. This confirms that the UQ is well-calibrated, providing a reliable measure of the model's local confidence.

This is further corroborated by the time-slice comparisons in Figure 6. The mean prediction of the BPINN accurately tracks the ground truth, and the 95% credible interval (shaded region) consistently envelops the true solution. The width of this uncertainty band visibly shrinks near the data points, providing an intuitive visualization of the model's confidence. The point-wise error maps in Figure 7 complete the picture, showing that the regions of highest error coincide with the regions of highest model-reported uncertainty.

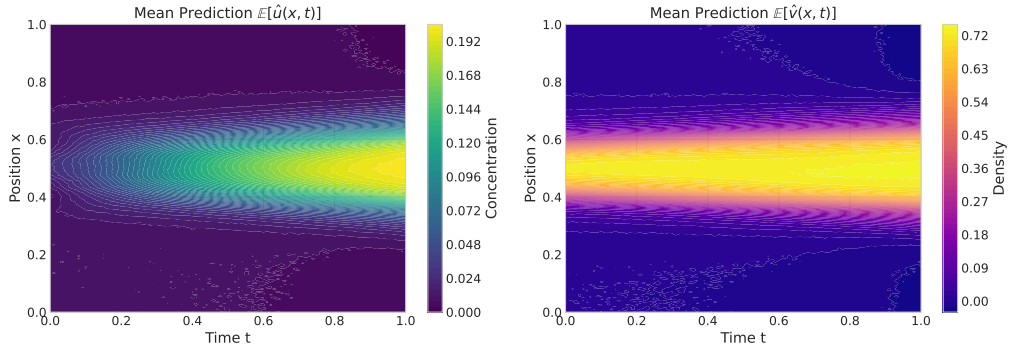


Figure 4: Mean spatio-temporal solutions for the growth factor concentration $\hat{u}(x,t)$ (left) and inflammatory cell density $\hat{v}(x,t)$ (right) predicted by the trained BPINN.

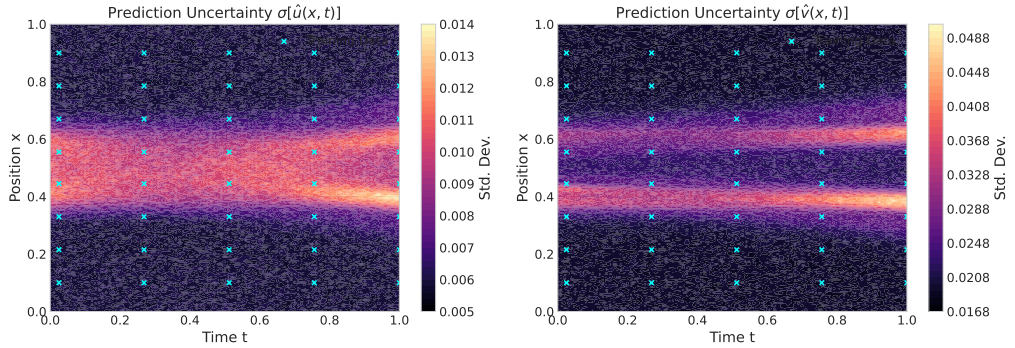


Figure 5: Predictive uncertainty (standard deviation) for $\hat{u}(x,t)$ (left) and $\hat{v}(x,t)$ (right). The uncertainty is lowest near the training data locations (cyan markers) and increases in regions where no data was provided, demonstrating a well-calibrated UQ.

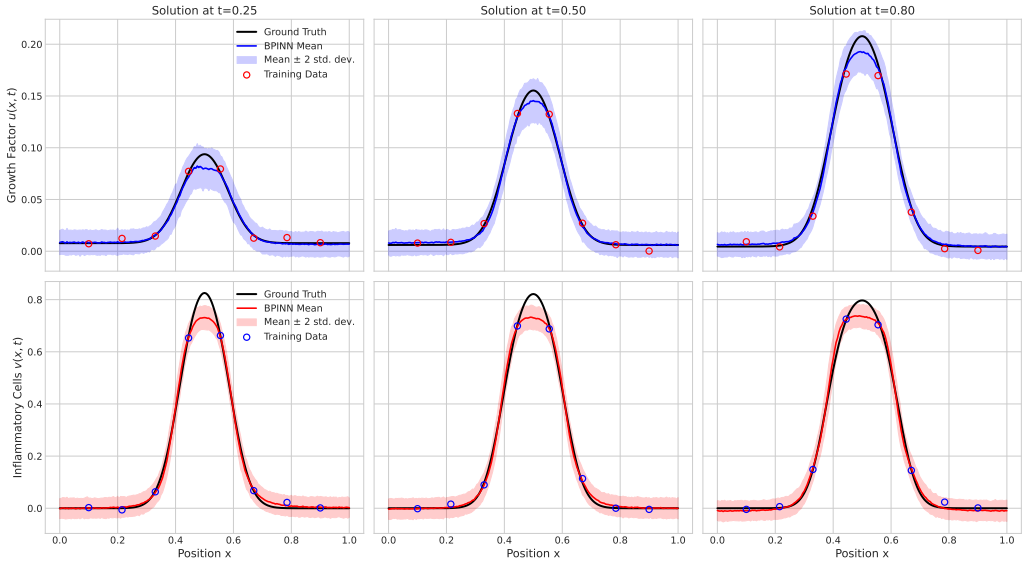


Figure 6: Comparison of BPINN predictions against ground truth at three time snapshots. The mean prediction (blue/red lines) accurately tracks the ground truth (black lines). The 95% credible interval (shaded regions) provides a reliable uncertainty bound and correctly contains the true solution.

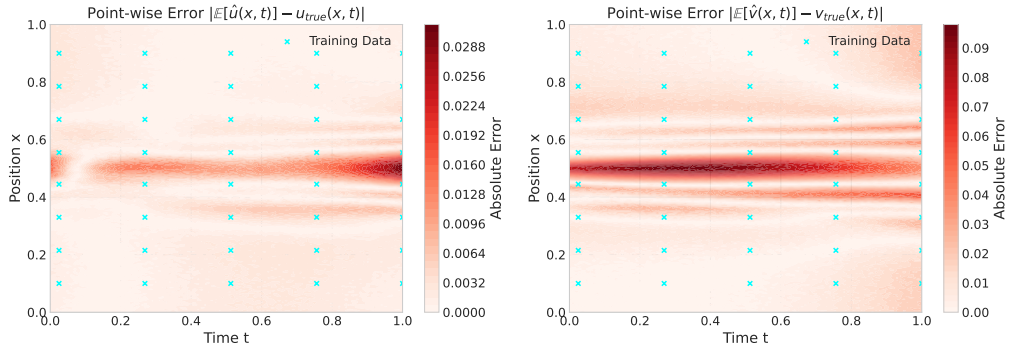


Figure 7: Point-wise absolute error between the mean BPINN prediction and the ground truth solution. The error is minimal across most of the domain and correlates with the predictive uncertainty shown in Fig. 5.

5. Conclusion

In this paper, we have developed and validated a Bayesian Physics-Informed Neural Network framework for the challenging inverse problem of

parameter inference in a PDE-based model of wound healing dynamics. By integrating Variational Inference and Monte Carlo Dropout, our approach successfully moves beyond deterministic point estimates to provide full posterior distributions for unknown biophysical parameters and principled uncertainty quantification for model predictions.

Our results demonstrate the dual strengths of this methodology. First, the BPINN framework proved capable of accurately inferring the key reaction parameters of the model from a dataset that was both extremely sparse and corrupted by noise. This highlights its potential as a powerful tool for calibrating complex biological models against realistic experimental data. Second, and perhaps more importantly, the framework correctly diagnosed a practical non-identifiability issue with the cell motility parameter. Instead of converging to a wrong value with false confidence, the BPINN assigned a large, honest uncertainty to this parameter, effectively flagging it as unresolvable from the given data. This diagnostic capability is invaluable for building robust and trustworthy models, preventing the propagation of unacknowledged parameter uncertainty into clinical or scientific conclusions.

The well-calibrated predictive uncertainty, which was shown to be lowest near data points and highest in unexplored regions of the domain, further underscores the reliability of the approach. By providing not just a prediction but also a measure of confidence, BPINNs can guide future experimental design, suggesting where new measurements would be most effective at reducing model uncertainty. This work establishes BPINNs as a sophisticated and practical tool for the computational sciences, with significant potential to advance the development of personalized and predictive models in biology and medicine.

6. Future Work

While this study successfully demonstrates the dual capability of BPINNs, it also illuminates pathways for future research. The current work is based on a simplified one-dimensional, two-species model; extending this framework to higher-dimensional systems that incorporate more complex biological features, such as additional cell types or mechanical forces, is a critical next step. Furthermore, our approach relies on Variational Inference with a mean-field Gaussian approximation and MC Dropout. Exploring more expressive posterior approximations, such as those based on normalizing flows [38], or employing more computationally intensive but asymptotically exact MCMC

methods like Hamiltonian Monte Carlo [39], could provide even more accurate uncertainty estimates, especially for problems with highly correlated or multi-modal posteriors. Finally, a particularly exciting direction is to leverage the quantified uncertainty from the BPINN to create a closed-loop, active learning system. The uncertainty maps could guide a Bayesian optimal experimental design (BOED) strategy [40, 41], which would intelligently select the next set of experimental measurement points to maximally reduce model uncertainty, thereby making the data acquisition process itself more efficient and targeted.

Acknowledgments

The authors acknowledge the developers of the PyTorch library and the broader open-source scientific computing community, whose tools were indispensable for this research.

References

[1] A. J. Singer, R. A. F. Clark, Cutaneous wound healing, New England Journal of Medicine, 341(10) (1999) 738-746.

[2] G. C. Gurtner, S. Werner, Y. Barrandon, M. T. Longaker, Wound repair and regeneration, Nature, 453(7193) (2008) 314-321.

[3] P. Martin, Wound healing—aiming for perfect skin regeneration, Science, 276(5309) (1997) 75-81.

[4] N. B. Menke, K. R. Ward, T. M. Witten, D. G. Bonchev, R. F. Diegelmann, Impaired wound healing, Clinics in Dermatology, 25(1) (2007) 19-25.

[5] A. F. Jørgensen, K. A. H. van der A, M. G. D. M. Follon, The 'art' of modeling in wound healing: a review, Mathematical Biosciences, 335 (2021) 108573.

[6] J. Cumming, E. McElwain, G. Upton, A mathematical model of wound healing and subsequent scarring, Journal of The Royal Society Interface, 4(12) (2007) 19-34.

[7] J. D. Murray, Mathematical Biology II: Spatial Models and Biomedical Applications, Springer, 2003.

[8] C. Valero, M. A. García-López, J. L. Grima, et al., A new mathematical model of wound contraction and scarring, Mathematical and Computer Modelling, 43(3-4) (2006) 404-420.

[9] J. A. Sherratt, J. D. Murray, Models of epidermal wound healing, Proceedings of the Royal Society of London. Series B: Biological Sciences, 241(1300) (1990) 29-36.

[10] L. Olsen, J. A. Sherratt, P. K. Maini, A mechanochemical model for adult dermal wound contraction and the permanence of the contracted state, Journal of Theoretical Biology, 184(4) (1997) 341-355.

[11] E. Javierre, F. J. Vermolen, C. W. J. Oomens, J. M. G. van der Slagt, A mathematical model for wound healing and its application to diabetic ulcers, Mathematical and Computer Modelling, 49(9-10) (2009) 1849-1864.

[12] R. Adam, A. Dale, A. D. C. K. Maini, A mechanochemical model of collagen-driven dermal wound healing, Bulletin of Mathematical Biology, 79(3) (2017) 538-566.

[13] K. E. Murphy, E. K. Hall, S. L. D. S. McElwain, A mathematical model of the role of growth factors and matrix metalloproteinases in dermal wound healing, Journal of Theoretical Biology, 313 (2012) 65-81.

[14] A. Tarantola, Inverse Problem Theory and Methods for Model Parameter Estimation, SIAM, 2005.

[15] H. T. Banks, S. Dediu, S. L. Ernstberger, A. L. Fuller, S. L. Hobbs, M. C. M. Thompson, Modeling and parameter estimation for a minimal wound healing model, Mathematical Biosciences and Engineering, 11(2) (2014) 243.

[16] R. C. Aster, B. Borchers, C. H. Thurber, Parameter Estimation and Inverse Problems, Elsevier, 2018.

[17] A. Raue, C. Kreutz, T. Maiwald, et al., Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood, Bioinformatics, 25(15) (2009) 1923-1929.

[18] M. Raissi, P. Perdikaris, G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, Journal of Computational Physics, 378 (2019) 686-707.

[19] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, et al., Physics-informed machine learning, Nature Reviews Physics, 3(6) (2021) 422-440.

[20] L. Lu, P. Jin, G. E. Karniadakis, DeepXDE: A deep learning library for solving differential equations, SIAM Review, 63(1) (2021) 208-228.

[21] M. Raissi, A. Yazdani, G. E. Karniadakis, Hidden fluid mechanics: A physics-informed deep-learning framework for inferring lift and drag on a cylinder, Science, 367(6481) (2020) 1026-1030.

[22] E. Haghighat, R. Raesi, T. T. T. Bui, et al., A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics, Computer Methods in Applied Mechanics and Engineering, 379 (2021) 113741.

[23] S. Cai, Z. Wang, S. Wang, et al., Physics-informed neural networks for heat transfer problems, Journal of Heat Transfer, 143(6) (2021) 060801.

[24] G. Kissas, Y. Yang, E. Hwu, et al., Machine learning in cardiovascular flows: Predicting arterial pressures from non-invasive 4D-flow MRI, Nature Machine Intelligence, 2(5) (2020) 273-281.

[25] F. Sahli Costabal, E. K. Yang, P. Perdikaris, et al., Physics-informed neural networks for cardiac activation mapping, Frontiers in Physics, 8 (2020) 42.

[26] L. Sun, H. Gao, S. Pan, J. Wang, Surrogate modeling for fluid flows based on physics-constrained deep learning, Computer Methods in Applied Mechanics and Engineering, 365 (2020) 113024.

[27] A. Krishnapriyan, A. Gholami, S. Z. R. Bekam, et al., Characterizing possible failure modes in physics-informed neural networks, in: Advances in Neural Information Processing Systems 34, 2021, pp. 26548-26560.

[28] S. Wang, Y. Teng, P. Perdikaris, Understanding and mitigating gradient flow pathologies in physics-informed neural networks, SIAM Journal on Scientific Computing, 43(5) (2021) A3055-A3081.

[29] L. Yang, X. Meng, G. E. Karniadakis, B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data, Journal of Computational Physics, 425 (2021) 109913.

[30] J. Zhang, Y. Li, R. K. L. Su, et al., A Bayesian deep learning framework for solving forward and inverse problems of partial differential equations, arXiv preprint arXiv:1911.07229 (2019).

[31] Y. Zhu, N. Zabaras, P. Koutsourelakis, P. Perdikaris, Physics-informed-GANs for coastal flood visualization, Journal of Computational Physics, 453 (2022) 110952.

[32] A. F. Psaros, K. Kawaguchi, G. E. Karniadakis, Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons, Journal of Computational Physics, 477 (2023) 111902.

[33] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, D. B. Rubin, Bayesian Data Analysis, Third Edition, CRC Press, 2013.

[34] D. M. Blei, A. Kolda, M. I. Jordan, Variational inference: A review for statisticians, Journal of the American Statistical Association, 112(518) (2017) 859-877.

[35] Y. Gal, Z. Ghahramani, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning, in: Proceedings of the 33rd International Conference on Machine Learning, 2016, pp. 1050-1059.

[36] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).

[37] A. Paszke, S. Gross, F. Massa, et al., PyTorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems 32, 2019.

[38] D. J. Rezende, S. Mohamed, Deep generative models, in: Proceedings of the 32nd International Conference on Machine Learning, 2015, pp. 1530-1538.

[39] S. Duane, A. D. Kennedy, B. J. Pendleton, D. Roweth, Hybrid Monte Carlo, Physics Letters B, 195(2) (1987) 216-222.

[40] K. Chaloner, I. Verdinelli, Bayesian experimental design: A review, Statistical Science, 10(3) (1995) 273-304.

[41] D. V. Lindley, On a measure of the information provided by an experiment, The Annals of Mathematical Statistics, 27(4) (1956) 986-1005.