Modeling the Immune Response with Machine Learning: Validation and Computational Cost Analysis of PINNs Compared to Neural Networks and Finite Volume Methods

Thiago E. Fernandes[1], **Rodrigo W. dos Santos**[2], **Marcelo Lobosco**[1]

[1]Graduate Program on Computational Modelling, Federal University of Juiz de Fora (UFJF)
Juiz de Fora – MG – Juiz de Fora – Brazil

[2]Department of Computer Science, Federal University of Juiz de Fora (UFJF)
Juiz de Fora – MG – Juiz de Fora – Brazil

thiago.esterci@estudante.ufjf.br, rodrigo.weber@ufjf.br

marcelo.lobosco@ufjf.br

Abstract. *This work investigates deep neural networks as surrogate models for simulating the spatiotemporal dynamics of the immune response in infectious myocarditis. A reduced PDE model describing pathogen–leukocyte interactions in myocardial tissue is solved using a finite volume method (FVM), whose solutions are used to train conventional neural networks (NNs) and physics-informed neural networks (PINNs). The approaches are systematically compared in terms of predictive accuracy and computational performance. Results show that conventional NNs reproduce the reference FVM solution more accurately than PINNs, which exhibit excessive smoothing and reduced ability to capture sharp gradients. Despite these differences, neural networks enable significantly faster inference, achieving approximately* $6.6\times$ *acceleration relative to the GPU-accelerated FVM. The computational advantage becomes most relevant in multi-execution scenarios such as parametric studies, sensitivity analyses, and inverse problems.*

1. Introduction

Infectious myocarditis is an inflammatory disease of the cardiac muscle frequently associated with myocardial edema, which can compromise electrical conduction and contractile function and may lead to arrhythmias, heart failure, or sudden death [Lourenço et al. 2022]. The condition is triggered by pathogen invasion, which activates the innate immune response, promotes leukocyte recruitment, increases vascular permeability, and induces fluid extravasation into the interstitial space. Epidemiological data indicate a 62% increase in global myocarditis cases between 1990 and 2019, reaching approximately 1.26 million cases and 324,000 deaths in 2019 alone [Wang et al. 2023], underscoring its clinical relevance.

Mathematical models based on ordinary and partial differential equations have been widely used to describe immune–tissue interactions in a mechanistic and interpretable manner. However, these approaches depend on accurate parameter estimation and well-defined initial and boundary conditions, which are often uncertain in biological systems and may limit the practical applicability of traditional numerical solvers

[Regazzoni et al. 2024]. In parallel, machine learning (ML) methods enabled by advances in high-performance computing have demonstrated the ability to learn complex dynamical systems directly from data. Despite their flexibility, purely data-driven models often lack interpretability and physical consistency, raising concerns in biomedical and *in silico* applications [Viceconti et al. 2019]. Physics-informed neural networks (PINNs) aim to bridge this gap by embedding governing equations into the training process, thereby combining physical consistency with the representational flexibility of neural networks. Nevertheless, enforcing these constraints increases computational complexity, and it remains unclear whether PINNs consistently outperform conventional neural networks (NNs), particularly in problems involving sharp gradients and localized phenomena.

Therefore, this work investigates whether NN-based surrogate models can reproduce the spatiotemporal dynamics of immune response models with competitive accuracy while providing meaningful computational advantages, and whether the explicit enforcement of physical constraints in PINNs justifies the associated computational overhead. The main contributions of this work are: a) a controlled benchmark comparing GPU-parallelized finite volume simulations, conventional neural networks, and physics-informed neural networks for an immune response PDE model [1]; b) a fair comparison in which PINNs and conventional NNs share identical architectures, isolating the impact of physics-based constraints on performance; c) demonstration that neural surrogates enable substantial inference acceleration through fully parallel spatio-temporal evaluation; and d) the development of *FisiocomPINN*, a unified and modular framework for defining and training PINNs within reproducible and GPU-compatible workflows.

2. Immune System Model

This study considers a reduced model of the innate immune response describing the interaction between pathogens (C_p) and leukocytes (C_l) within a homogeneous, isotropic, fluid-saturated porous medium representing myocardial tissue. The formulation follows the model proposed in [Lourenço et al. 2022], restricting the analysis to immune dynamics. The spatiotemporal evolution of pathogen concentration is governed by

$$\begin{cases} \dfrac{\partial(\phi C_p)}{\partial t} = \nabla \cdot (D_b \nabla C_p) - r_b + q_b & \text{in } \Omega \times I, \\ D_b \nabla C_p \cdot \mathbf{n} = 0 & \text{on } \partial\Omega \times I, \\ C_p(x,0) = C_{p0}(x) & \text{in } \Omega, \end{cases} \tag{1}$$

where $\Omega \subset \mathbb{R}^2$, $I = [0, t_f)$, ϕ denotes tissue porosity, and D_b is the pathogen diffusion coefficient. Pathogen reproduction and leukocyte-mediated removal are given by

$$q_b = c_p C_p, \qquad r_b = \lambda_{nb} C_l C_p, \tag{2}$$

where c_p is the pathogen reproduction rate and λ_{nb} is the phagocytosis rate.

Leukocyte dynamics are described by

[1] Source codes and examples: https://github.com/Esterci/2D_imune_edema_pinn.

$$\begin{cases} \dfrac{\partial(\phi_f C_l)}{\partial t} = \nabla \cdot (D_n \nabla C_l - \chi_{nb} C_l \nabla C_p) - r_n + q_n & \text{in } \Omega \times I, \\ (D_n \nabla C_l - \chi_{nb} C_l \nabla C_p) \cdot \mathbf{n} = 0 & \text{on } \partial\Omega \times I, \\ C_l(x, 0) = C_{l0}(x) & \text{in } \Omega, \end{cases} \tag{3}$$

where D_n is the leukocyte diffusion coefficient and χ_{nb} denotes the chemotactic sensitivity. The leukocyte recruitment and decay terms are

$$q_n = \gamma_n C_p \left(C_{n,\max} - C_l\right), \tag{4}$$

$$r_n = \lambda_{bn} C_l C_p + \mu_n C_l, \tag{5}$$

where γ_n represents vascular permeability, $C_{n,\max}$ is the maximum circulating leukocyte concentration, λ_{bn} denotes the pathogen-induced leukocyte death rate, and μ_n is the natural leukocyte decay rate.

Homogeneous Neumann boundary conditions are imposed for both variables, assuming no flux across the boundary and representing functional isolation of the inflamed region.

3. Numerical Implementation and Results

The governing equations were solved using an explicit finite volume method (FVM) on a uniform grid. Diffusive terms were approximated with second-order centered differences, while time integration employed a first-order forward Euler scheme. To prevent non-physical oscillations, an upwind scheme was applied to the advective fluxes.

The FVM solver was parallelized on GPU using CUDA, exploiting the spatial independence of control volumes within each time step. The resulting FVM solutions were used as reference data for training neural networks. Conventional neural networks (NNs) and physics-informed neural networks (PINNs) shared identical architectures, ensuring that performance differences arise solely from the inclusion of physics-based constraints in the loss function. The dataset was divided into 90% for training and 10% for validation.

All experiments were implemented using the *FisiocomPINN* framework [2], which provides modular loss construction, stochastic sampling utilities, and reproducible training pipelines across CPU and GPU backends.

3.1. Results

To address the research questions, the immune response model was solved using three approaches: the FVM (implemented in both sequential CPU and CUDA-accelerated GPU versions), conventional neural networks, and PINNs. All computational accelerations reported in this study are measured relative to the sequential CPU-based FVM implementation, which serves as the baseline reference.

[2]Fisiocompinn: A physics-informed neural network framework for scientific computing in pytorch. https://github.com/ ybwerneck/Pinn-Torch.

Multiple NN architectures were evaluated by varying the number of hidden layers and neurons per layer. Each configuration was repeated 33 times. A grid search over 600 configurations showed that accuracy improves with network depth, stabilizing at approximately five hidden layers, while increasing network width reduces error up to about 96 neurons with marginal improvements thereafter. Considering the trade-off between accuracy and computational efficiency, the architecture with seven hidden layers and 32 neurons per layer was selected.

Figure 1 compares the solutions obtained with the three approaches. Although the PINN partially reconstructs the spatiotemporal dynamics of pathogens (C_p) and leukocytes (C_l), it fails to preserve important features of the reference FVM solution, particularly stationary regimes and sharp spatial gradients, exhibiting premature dissipation and over-smoothing. In contrast, the conventional NN shows significantly better agreement with the FVM results, accurately capturing quasi-stationary pathogen behavior and reproducing leukocyte activation peaks and propagation patterns with low absolute error across the domain.

Table 1 summarizes the quantitative comparison between the GPU-accelerated FVM, the PINN, and the conventional NN. The NN achieved the lowest RMSE among the ML approaches, while the PINN presented higher errors and longer training times. Although the NN's MaxAE remains higher than its RMSE, indicating localized deviations, its overall predictive accuracy remains superior.

Table 1. Comparison between the GPU-accelerated FVM, the PINN, and the conventional Neural Network (NN), considering error metrics (RMSE and MaxAE), acceleration relative to the CPU-based FVM, total training time, and inference time.

Method	RMSE	MaxAE	Acceleration	Training Time (s)	Inference Time (s)
FVM-CUDA	–	–	483 ± 8	–	0.01840 ± 0.00029
PINN	0.065	1.799	3216 ± 23	2050.09 ± 2.59	0.00277 ± 0.00005
NN	0.010	0.795		267.11 ± 0.46	

The neural network models achieved an inference acceleration of approximately 6.6 times ($3216 \div 483$) relative to the GPU-accelerated FVM. This improvement arises from the structural difference between the approaches: neural networks treat time as an independent input variable and evaluate the solution fully in parallel across space and time, whereas the explicit FVM requires sequential time integration, limiting its parallel efficiency despite its lower per-step computational cost.

However, when training cost is considered, this advantage must be interpreted carefully. If the total computational cost is expressed as

$$\frac{t_t}{n_e} + t_i \leq t_{\mathrm{FVM}}, \tag{6}$$

where t_t is the total training time, t_i is the inference time per execution, t_{FVM} is the execution time of the FVM, and n_e is the number of executions, then

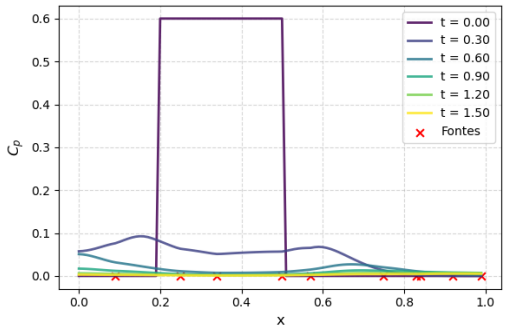

(a) FVM – pathogen concentration (C_p).

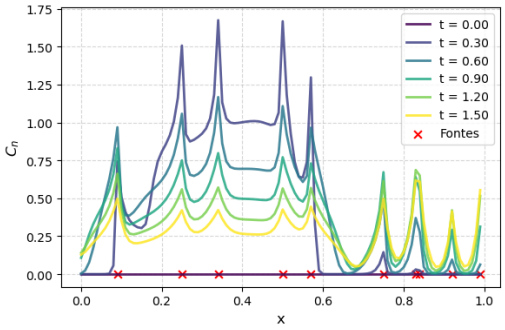

(b) FVM – leukocyte concentration (C_l).

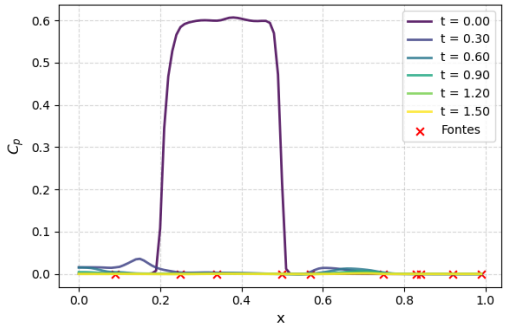

(c) PINN – pathogen concentration (C_p).

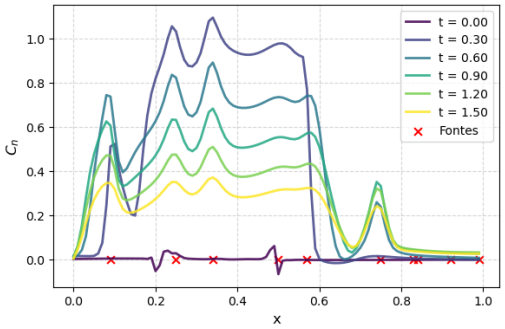

(d) PINN – leukocyte concentration (C_l).

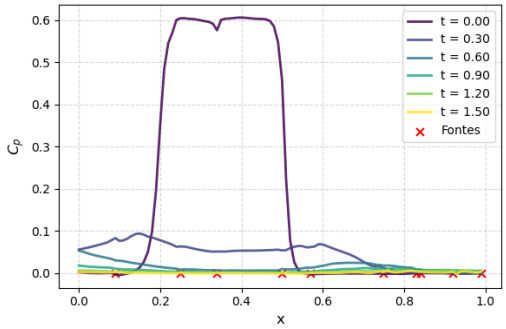

(e) Classical neural network – pathogen concentration (C_p).

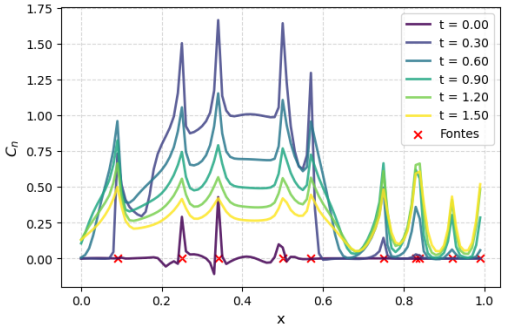

(f) Classical neural network – leukocyte concentration (C_l).

Figure 1. Temporal comparison of pathogen (C_p) and leukocyte (C_l) concentration curves generated by the FVM, PINN, and NN methods.

$$n_e \geq \frac{t_t}{(t_{\mathrm{FVM}} - t_i)}. \tag{7}$$

Based on the measured runtimes, more than 1.7×10^4 executions are required for the conventional NN to achieve a lower total computational cost than direct FVM simulation. This threshold indicates that the practical advantage of the NN emerges primarily in multi-execution scenarios, such as parametric studies, sensitivity analyses, or inverse problems, where the initial training cost can be amortized.

4. Discussion

This work investigated the applicability of deep neural networks for modeling the spatiotemporal dynamics of the local immune response in infectious myocarditis, focusing on the evolution of pathogen and leukocyte concentrations. Physics-informed neural networks (PINNs) and conventional data-driven neural networks (NNs) were implemented and systematically compared against a validated mathematical model solved using the finite volume method (FVM), considering predictive accuracy, computational efficiency, and robustness under limited data availability.

The results show that, for the immune response model considered, conventional NNs consistently outperform PINNs in terms of predictive accuracy and agreement with the reference FVM solution. In particular, NNs successfully reproduced stationary regimes, propagation fronts, and localized activation patterns when trained with a sufficient number of temporal samples. In contrast, PINNs exhibited a systematic tendency toward excessive smoothing, which led to underestimation of concentration peaks and limited their ability to capture localized or quasi-stationary solution features. Moreover, the training cost of PINNs was substantially higher due to the repeated evaluation of differential operators during optimization. Both neural network approaches achieved significant inference acceleration relative to the FVM, including its GPU-accelerated implementation, by allowing fully parallel evaluation in space and time without requiring sequential temporal integration. However, when only a single simulation is required, the training cost of neural networks outweighs the benefits of direct numerical simulation. In contrast, in multi-execution scenarios, such as parametric studies, sensitivity analyses, or inverse problems, neural networks enable rapid inference at arbitrary spatial and temporal locations, allowing efficient exploration of the solution manifold and compensating for the higher upfront training cost.

Although PINNs remain a promising framework for problems characterized by strong physical constraints or severely limited data availability, their additional complexity and computational overhead were not justified for the immune response model analyzed in this study.

References

Lourenço, W. d. J., Reis, R. F., Ruiz-Baier, R., Rocha, B. M., Dos Santos, R. W., and Lobosco, M. (2022). A poroelastic approach for modelling myocardial oedema in acute myocarditis. *Frontiers in Physiology*, 13:888515.

Regazzoni, F., Pagani, S., Salvador, M., Dede', L., and Quarteroni, A. (2024). Learning the intrinsic dynamics of spatio-temporal processes through latent dynamics networks. *Nature Communications*, 15(1):1834.

Viceconti, M., Juárez, M. A., Curreli, C., Pennisi, M., Russo, G., and Pappalardo, F. (2019). Credibility of in silico trial technologies—a theoretical framing. *IEEE journal of biomedical and health informatics*, 24(1):4–13.

Wang, Y.-W.-Y., Liu, R.-B., Huang, C.-Y., Li, H.-Y., Zhang, Z.-X., Li, X.-Z., Liu, J.-L., Zhang, C., Xiong, X., and Niu, Y.-M. (2023). Global, regional, and national burdens of myocarditis, 1990–2019: systematic analysis from gbd 2019: Gbd for myocarditis. *BMC Public Health*, 23(1):714.