

Bayesian Physics Informed Neural Networks for Reliable Transformer Prognostics

Ibai Ramirez¹, Jokin Alcibar¹, Joel Pino¹, Mikel Sanz^{2,5,6}, David Pardo^{3,5,6}, and Jose I. Aizpurua^{4,6*}

¹ *Mondragon University, Electronics & Computing Department, Mondragon, 20500, Spain*
iramirezg@mondragon.edu
jalcibar@mondragon.edu
jpino@mondragon.edu

² *University of the Basque Country, Department of Physical Chemistry, Leioa, 48080, Spain*
mikel.sanz@ehu.eus

³ *University of the Basque Country, Department of Mathematics, Leioa, 48080, Spain*
david.pardo@ehu.eus

⁴ *University of the Basque Country, Department of Computer Science & Artificial Intelligence, Donostia-San Sebastian, 20018, Spain*
joxe.aizpurua@ehu.eus

⁵ *Basque Centre for Applied Mathematics, Bilbao, 48009, Spain*

⁶ *Ikerbasque, Basque Foundation for Science, Bilbao, 48011, Spain*
*Corresponding Author

ABSTRACT

Scientific Machine Learning (SciML) integrates physics and data into the learning process, offering improved generalization compared with purely data-driven models. Despite its potential, applications of SciML in prognostics remain limited, partly due to the complexity of incorporating partial differential equations (PDEs) for ageing physics and the scarcity of robust uncertainty quantification methods. This work introduces a Bayesian Physics-Informed Neural Network (B-PINN) framework for probabilistic prognostics estimation. By embedding Bayesian Neural Networks into the PINN architecture, the proposed approach produces principled, uncertainty-aware predictions. The method is applied to a transformer ageing case study, where insulation degradation is primarily driven by thermal stress. The heat diffusion PDE is used as the physical residual, and different prior distributions are investigated to examine their impact on predictive posterior distributions and their ability

to encode *a priori* physical knowledge. The framework is validated against a finite element model developed and tested with real measurements from a solar power plant. Results, benchmarked against a dropout-PINN baseline, show that the proposed B-PINN delivers more reliable prognostic predictions by accurately quantifying predictive uncertainty. This capability is crucial for supporting robust and informed maintenance decision-making in critical power assets.

1. INTRODUCTION

Scientific Machine Learning (SciML) is an emerging interdisciplinary field that integrates physics and data within the learning process (Karniadakis et al., 2021). Several SciML solutions have been proposed to address scientific and engineering problems, including Physics-Informed Neural Networks (PINNs) (Raissi, Perdikaris, & Karniadakis, 2019a), Kolmogorov Arnold Networks (KANs) (Liu et al., 2025), Neural operators (Kovachki et al., 2023), Physics-informed Neural Operators (Z. Li et al., 2024), and their variants (Toscano et al., 2025).

In the area of Prognostics & Health Management (PHM),

Ibai Ramirez et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

hybrid prognostics models have been widely explored (Aizpurua & Catterson, 2015; Guo, Li, & Li, 2020; Zio, 2022), which are conceptually aligned with SciML principles (Xu, Kohtz, Boakye, Gardoni, & Wang, 2023). These approaches integrate physics-of-failure models with machine learning (ML) methods through (i) sequential configurations, *i.e.* connect physics model outcomes with ML models (Arias Chao, Kulkarni, Goebel, & Fink, 2022) or vice-versa (Daigle, Roychoudhury, & Bregon, 2015), and (ii) parallel configurations, where physics and ML models are used simultaneously and their outcomes are fused, *e.g.* for error-correction configurations (Alcibar, Aizpurua, Zugasti, & Peñagarikano, 2025).

The focus of this work is on SciML approaches that learn the dynamics of the physics model along with the ML model. Namely, PINNs focus on the use of Neural Network (NN) architectures, and they learn to solve partial differential equations (PDE) through modifications in the NN loss function (Raissi et al., 2019a). Despite their promise, PINN-based solutions for PHM remain limited (Xu et al., 2023). There have been solutions focused on battery prognostics through imposing a monotonic degradation constraint in the loss function (F. Wang, Zhai, Zhao, Di, & Chen, 2024) and transformer prognostics through heat diffusion PDEs (Ramirez et al., 2025).

Other relevant hybrid prognostics methods have focused on the integration of principled models in a Recurrent Neural Network for Li-Ion battery prognostics (Nascimento, Viana, Corbetta, & Kulkarni, 2023). The approach captures epistemic uncertainty through variational layers to estimate reduced-order model parameters. Similarly, (Fernández, Chiachío, Chiachío, Barros, & Corbetta, 2023) proposed a physics-guided Bayesian Neural Network (BNN) for reinforced concrete columns under lateral loading, using Approximate Bayesian Computation for robust uncertainty quantification.

It can be observed that PINN-based solutions for PHM have been focused on deterministic predictions. This is aligned with the observation that, while the progress and application of ML and AI models is rapidly increasing in different scientific and engineering areas, in contrast, the development of ML reliability assessment methods are relatively underdeveloped. In the context of failure prognostics, uncertainty quantification (UQ) is essential for reliable future ageing predictions (Sankararaman, 2015). To this end, different UQ methods have been developed, including state-space models incorporating Bayesian principles (Daigle et al., 2015), Gaussian Process based approaches (Biggio, Wieland, Chao, Kastanis, & Fink, 2021), Conformal Prediction for prognostics (Javanmardi & Hüllermeier, 2023), and Bayesian Convolutional NNs (Alcibar, Aizpurua, & Zugasti, 2024). For

a complete picture of UQ for prognostics, refer to an extensive tutorial (Nemani et al., 2023) and benchmarking studies (Basora, Viens, Chao, & Olive, 2025).

Within the SciML community, extensions of PINNs with UQ have been explored by replacing deterministic NNs with BNNs, leading to Bayesian-PINNs (B-PINNs) (Linka et al., 2022). More broadly, Bayesian deep learning integrated with physics principles has shown solid and reliable results, not only for PINNs (Flores, Graf, Protopapas, & Pichara, 2025), but also for Neural Operators (Garg & Chakraborty, 2023) and KANs (Gao & Karniadakis, 2025). Other principled UQ alternatives for PINNs have focused on the modification of the output layer with a Gaussian process (J. Li et al., 2024). Refer to (Psaros, Meng, Zou, Guo, & Karniadakis, 2023) for an extensive review.

Additionally, there have been studies that focused on the development of wrapper models that operate as an add-on to existing PINN models for UQ, such as Epistemic-PINN (EPINN), which are trained through Hamiltonian Monte Carlo (HMC) methods (Nair, Jacob, Howard, Drgona, & Stinis, 2025) or Conformal Prediction PINNs (C-PINNs), which provides guaranteed uncertainty estimates (Podina, Rad, & Kohandel, 2024).

Simpler yet effective UQ modelling techniques include Monte Carlo dropout (Gal & Ghahramani, 2016) and deep ensemble methods (Lakshminarayanan, Pritzel, & Blundell, 2017), also used for PINNs, *i.e.* dropout-PINN (dPINN) (Zhang, Lu, Guo, & Karniadakis, 2019) and deep ensemble PINNs (EnsPINN) (Jiang, Wang, Wen, Li, & Wang, 2023; Zou, Wang, & Karniadakis, 2025).

Table 1 summarizes representative SciML and hybrid PHM methods, highlighting if they incorporate UQ, adhere to SciML principles, their underlying method, and the specific PHM application considered (if addressed).

Table 1. Comparison of SciML and hybrid PHM methods, including relevant features for this work.

Method	UQ	SciML	Method	PHM
(Fernández et al., 2023)	✓	✗	BNN	Material
(Nascimento et al., 2023)	✓	✗	RNN	Battery
(F. Wang et al., 2024)	✗	✓	PINN	Battery
(Ramirez et al., 2025)	✗	✓	PINN	Transformer
(Linka et al., 2022)	✓	✓	B-PINN	✗
(Podina et al., 2024)	✓	✓	C-PINN	✗
(Nair et al., 2025)	✓	✓	EPINN	✗
(Zhang et al., 2019)	✓	✓	dPINN	✗
(Zou et al., 2025)	✓	✓	EnsPINN	✗
Ours	✓	✓	B-PINN	Transformer

In this context, it can be observed that, focused on prognostics activities, existing PINN models have been based on deterministic models. In the broader SciML community, uncertainty-aware PINN methods have been pro-

posed, but to the best of authors' knowledge, tested on synthetic controlled problems without addressing PHM problems.

Accordingly, the main contribution of this work is the development of a Bayesian-PINN model for electrical transformer prognostics. The proposed approach predicts a full posterior distribution for each spatiotemporal coordinate, which enables the inference of the maximum likelihood estimate along with the confidence of the model on the prediction.

The influence of different sources of uncertainty is thoroughly investigated to determine the most appropriate hyperparameters. Furthermore, the robustness of the designed B-PINN model is evaluated with respect to noise of the input and residual data points. The methodology is validated on a real transformer case study using probabilistic evaluation metrics. Obtained results demonstrate that the proposed B-PINN (i) delivers more accurate results than the dropout-PINN, with the added flexibility of specifying prior knowledge, and (ii) produces more reliable predictions than vanilla PINN model due to the probabilistic predictions and inferred uncertainty levels.

The remainder of this article is organized as follows. Section 2 provides background theory to understand transformer thermal and ageing modelling. Section 3 introduces the proposed approach. Section 4 presents the case study. Section 5 presents the results, and finally, Section 6 concludes.

2. POWER TRANSFORMER THERMAL AND AGEING MODELLING

Transformers are key assets for the reliable operation of the power grid. The increasing penetration of renewable energy sources (RESs) to the grid affects the transformer's health (Aizpurua et al., 2023). The main insulating material for oil-immersed transformers is paper immersed oil, and their main failure mode is the insulation degradation. The insulating paper is made of cellulose polymer and the degree of polymerization determines the strength of the insulating paper (International Electrotechnical Commission, 2018). The insulation paper degradation is directly caused by the thermal stress.

Accordingly, this section reviews the main transformer thermal modelling steps, including an oil temperature estimation stage (Section 2.1) and subsequent winding temperature estimation stage, which is used to calculate the hottest-spot temperature (HST), *i.e.* highest insulation temperature (Section 2.2). Finally, the thermal model is connected to the insulation aging assessment model, which is used to estimate the loss of life of the insulation (Section 2.3).

2.1. Spatiotemporal Oil Temperature Modelling

The spatial distribution of oil and winding temperature is key for the cost-effective transformer health management. To capture this, a thermal modelling approach is developed based on partial differential equations (PDE). The heat diffusion PDE is considered to model the temporal and spatial evolution of the transformer oil temperature ($\Theta_O(x, t)$). Due to the transformer oil characteristics, radiative heat diffusion is considered and not convection. The general form of the one-dimensional heat diffusion equation is defined as (Ramirez et al., 2025):

$$\begin{aligned} k \frac{\partial^2 \Theta_O(x, t)}{\partial x^2} + q(x, t) &= \rho c_p \frac{\partial \Theta_O(x, t)}{\partial t} \\ \frac{\partial^2 \Theta_O(x, t)}{\partial x^2} + \frac{1}{k} q(x, t) &= \frac{1}{\alpha} \frac{\partial \Theta_O(x, t)}{\partial t} \end{aligned} \quad (1)$$

where $x, t \in \mathbb{R}$ are the independent variables, which denote position [m] and time [s], respectively, $\Theta_O(x, t)$ is given in Kelvin [K], k is the thermal conductivity [W/m.K], c_p is the specific heat capacity [J/kg.K], ρ is the density [kg/m³], $q(x, t)$ is the rate of heat generation [W/m³], and $\alpha = \frac{k}{\rho c_p}$ is the thermal diffusivity [m²/s].

Figure 1 shows the thermal parameters of the transformer of the heat diffusion model (Ramirez et al., 2025), which considers the heat source, $q(x, t)$, and the convective heat transfer, $h(\Theta_O(x, t) - \Theta_A(t))$.




Figure 1. Transformer heat diffusion model.

The evolution of the heat source in space and time, $q(x, t)$, is defined as follows:

$$q(x, t) = P_0 + P_K(t) - h(\Theta_O(x, t) - \Theta_A(t)) \quad (2)$$

where $\Theta_A(t)$ is the ambient temperature, h is the convective heat transfer coefficient, P_0 is the no load losses [W], and $P_K(t)$ is the load losses, defined as follows:

$$P_K(t) = K(t)^2 \mu, \quad (3)$$

where $K(t)$ is the load factor [p.u.], and μ is the rated load losses [W].

The challenge is to accurately model the transformer oil temperature vertically along its height H (Figure 1). It

can be observed that the spatial distribution is considered along the vertical axis (x). It is assumed that $\Theta_O(x, t)$ is equal to $\Theta_A(t)$ at the bottom ($x=0$) and $\Theta_{TO}(t)$ at the top ($x=H$). Namely, these are the Dirichlet boundary conditions of the PDE that is going to be solved:

$$\begin{aligned}\Theta_O(0, t) &= \Theta_A(t) \\ \Theta_O(H, t) &= \Theta_{TO}(t)\end{aligned}\quad (4)$$

2.2. Hotspot Temperature (HST) Modelling

The HST is the most critical thermal indicator. However, direct measurements are difficult and expensive. Generally the HST value, $\Theta_H(t)$, is estimated indirectly from the top-oil temperature (TOT), $\Theta_{TO}(t)$, defined as (International Electrotechnical Commission, 2018):

$$\Theta_H(t) = \Theta_{TO}(t) + \Delta\Theta_H(t) \quad (5)$$

where $\Delta\Theta_H(t)$ is HST rise over TOT and $t \in \mathbb{R}$ is time.

With the spatiotemporal oil temperature estimate $\hat{\Theta}_O(x, t)$, the winding temperature distribution, $\hat{\Theta}_W(x, t)$, is defined as follows:

$$\hat{\Theta}_W(x, t) = \hat{\Theta}_O(x, t) + \Delta\Theta_H(t) \quad (6)$$

where, $x, t \in \mathbb{R}$ are the position and time, respectively, $\hat{\Theta}_O(x, t)$ is the spatiotemporal oil temperature estimate. The hottest spatial temperature in Eq. (6) at each time instant t , corresponds to the HST in Eq. (5), i.e. $\max(\hat{\Theta}_W(\cdot, t)) = \Theta_H(t)$.

$\Delta\Theta_H(t)$ is the HST rise over TOT, which is given by:

$$\Delta\Theta_H(t) = \Delta\Theta_{H_1}(t) - \Delta\Theta_{H_2}(t) \quad (7)$$

where $\Delta\Theta_{H_1}(t)$ and $\Delta\Theta_{H_2}(t)$ model the oil heating considering the HST variations defined as follows (International Electrotechnical Commission, 2018):

$$d\Delta\Theta_{H_i}(t) = v_i [\beta_i K(t)^y - \Delta\Theta_{H_i}(t)] \quad (8)$$

where $K(t)$ is the load factor [p.u.], y is the winding exponent constant, which models the loading exponential power with the heating of the windings, $i=\{1, 2\}$, $v_1 = \Delta t / k_{22} \tau_w$, $\beta_1 = k_{21} \Delta\Theta_{H,R}$ both for $i=1$, and $v_2 = k_{22} \Delta t / \tau_{TO}$, $\beta_2 = (k_{21} - 1) \Delta\Theta_{H,R}$ both for $i=2$. $\Delta t = t - t'$, τ_w and τ_{TO} are the winding and oil time constants, k_{21} and k_{22} are the transformer thermal constants, and $\Delta\Theta_{H,R}$ is the HST rise at rated load. The operator d denotes a difference operation on Δt , such that $d\Delta\Theta_{H_i}(t) = \Delta\Theta_{H_i}(t) - \Delta\Theta_{H_i}(t')$ also for $i=\{1, 2\}$. To guarantee numerical stability, Δt should be small, never greater than half of the smaller time constant.

Under steady state, the initial condition, $\Theta_H(0)$, can be defined as:

$$\Theta_H(0) = \Theta_{TO}(0) + k_{21} \Delta\Theta_{H,R} K(0)^y - (k_{21} - 1) \Delta\Theta_{H,R} K(0)^y \quad (9)$$

Eq. (9) allows iteratively estimating the next HST values, $\Theta_H(n\Delta t)$, $n \in \mathbb{Z}^+$, using Eqs. (5), (7), and (8).

2.3. Ageing Assessment

The IEC 60076-7 standard defines insulation ageing acceleration factor at time t , $V(t)$, as (International Electrotechnical Commission, 2018):

$$V(t) = 2^{(\Theta_H(t) - 98)/6} \quad (10)$$

Using the spatiotemporal winding temperature $\hat{\Theta}_W(x, t)$, the ageing acceleration factor at time t and position x , $V(x, t)$, can be defined as:

$$V(x, t) = 2^{(\hat{\Theta}_W(x, t) - 98)/6} \quad (11)$$

The IEC 60076-7 assumes an expected life of 30 years, with a reference HST of 98°C (International Electrotechnical Commission, 2018). The loss of life (LOL) at location x and time t can be defined as:

$$LOL(x, t) = \int_0^t V(x, t) dt \quad (12)$$

Consequently, the LOL at discrete time $L\Delta t$ and location x can be obtained by summing the ageing (cf. Eq. (11)) evaluated at the same time instants:

$$LOL(x, L\Delta t) = \sum_{n=0}^L V(x, n\Delta t) \quad (13)$$

where $n, L \in \mathbb{Z}^+$.

LOL can be converted into a recurrence relation for remaining useful life estimation (Aizpurua et al., 2023).

3. PROBABILISTIC AGEING APPROACH BASED ON BAYESIAN PINNS

The proposed transformer ageing estimation approach is based on a probabilistic spatiotemporal thermal stress model coupled with an ageing estimation model, as illustrated in Figure 2. The thermal model is formulated using a Bayesian-PINN, which estimates the spatiotemporal oil temperature $\hat{\Theta}_O(x, t)$. This estimate is fed into an empirical winding temperature estimation model, that computes the spatiotemporal winding temperature, $\hat{\Theta}_W(x, t)$. Finally, the winding temperature is connected to an empirical insulation ageing model, which calculates the probabilistic ageing estimate, $p(\hat{V}(x, t))$.

For clarity, this section focuses on the B-PINN approach. Figure 3 shows the proposed method, which integrates PINNs with Bayesian inference by replacing the deter-




Figure 2. Overall framework for probabilistic transformer insulation ageing estimation. The Bayesian PINN-based oil temperature model is coupled with an empirical winding temperature model and an empirical insulation ageing model.




Figure 3. Proposed Bayesian-PINN approach for the probabilistic spatiotemporal transformer thermal model.

ministic NNs in PINNs with BNNs. Accordingly, Section 3.1 introduces PINN basics, Section 3.2 defines BNNs, and Section 3.3 describes the B-PINN approach.

3.1. PINN basics

PINNs were introduced with the goal of encoding PDE-based physics in ML models (Raissi, Perdikaris, & Karniadakis, 2019b), taking advantage of the ability of NNs to act as universal approximators (Wright et al., 2022). A general PDE can be written as (Anagnostopoulos, Toscano, Stergiopoulos, & Karniadakis, 2024):

$$d\{u(x, t)\} = f(x, t) \quad (14)$$

where $x, t \in \mathbb{R}$ denote position and time, $u(x, t)$ is the unknown solution, d is the differential operator, and $f(x, t)$ is a forcing function introducing external influences.

PINNs consist of the NN part, in which the inputs define temporal ($t \in \mathbb{R}$) and spatial coordinates ($x \in \mathbb{R}$ for one-dimensional cases) for the initial conditions (IC) and boundary conditions (BC). The NN output is an approximated PDE solution at the space and time coordinates, denoted $\hat{u}(x, t)$. This is calculated through the iterative application of weights (w), biases (b), and non-linear activation functions (σ) over the input. Namely, the inputs are connected through neurons, where they are multi-

plied with the weights and summed with the bias term. Finally, the weighted sum is passed through an activation function (σ). Subsequently, the outcome $\hat{u}(x, t)$ is post-processed via automatic differentiation to compute the derivatives in space and time at certain collocation points (CP), generated via random sampling in the interior of the domain. The PINN is trained at these CPs by minimizing the residuals of the underlying PDE.

The loss function, $\mathcal{L}(\theta, \lambda_0, \lambda_b, \lambda_r)$, incorporates the prediction error of the NN at IC and BC, and the residual of the PDE estimated via automatic differentiation at CP:

$$\mathcal{L}(\theta, \lambda_0, \lambda_b, \lambda_r) = \mathcal{L}_0(\theta, \lambda_0) + \mathcal{L}_b(\theta, \lambda_b) + \mathcal{L}_r(\theta, \lambda_r) \quad (15)$$

where $\theta = \{w, b\}$ are the weights and bias terms of the NN, and $\mathcal{L}_0(\theta, \lambda_0)$, $\mathcal{L}_b(\theta, \lambda_b)$, and $\mathcal{L}_r(\theta, \lambda_r)$ are, respectively, the loss terms corresponding to IC, BC, and the residual of the PDE with their corresponding weights, λ_0 , λ_b , and λ_r , defined as follows:

$$\mathcal{L}_0(\theta, \lambda_0) = \lambda_0 \frac{1}{N_0} \sum_{i=1}^{N_0} |\hat{u}(x_i, 0) - u(x_i, 0)|^2 \quad (16)$$

$$\mathcal{L}_b(\theta, \lambda_b) = \lambda_b \frac{1}{N_b} \sum_{i=1}^{N_b} |\hat{u}(x_i, t_i) - u(x_i, t_i)|^2 \quad (17)$$

$$\mathcal{L}_r(\boldsymbol{\theta}, \lambda_r) = \lambda_r \frac{1}{N_r} \sum_{i=1}^{N_r} |r(x_i, t_i)|^2 \quad (18)$$

where N_0 , N_b , and N_r are, respectively, the number of IC, BC, and residue points, while $u(x_i, t_i)$ and $r(x_i, t_i)$ denote the known solution and the residual of PDE, for each training point i defined at the coordinates (x_i, t_i) . According to the PDE defined in Eq. (14), the residual $r(x, t)$ is defined as follows:

$$r(x, t) = d\{u(x, t)\} - f(x, t) \quad (19)$$

Unlike classical NNs, which only minimize the prediction error at measurement points, PINNs enforce governing physics throughout the domain.

Minimizing the loss function in Eq. (15) using a suitable optimization algorithm provides an optimal set of NN parameters $\boldsymbol{\theta} = \{\mathbf{w}, \mathbf{b}\}$. That is, approximating the PDE becomes equivalent to finding $\boldsymbol{\theta}$ values that minimize the loss with a predefined accuracy. Overall, the key training parameters include: number of neurons and number of layers, number of CP, activation function, and the optimizer. Finding the correct solution requires knowing IC and BC. Additionally, random locations (x_i, t_i) , named CP, are used to evaluate the residual loss in Eq. (18).

The main motivation for using PINNs over numerical methods is the computational effort and adaptability to different solutions. Numerical models require a mesh of parameters to model and evaluate the PDE. As for PINNs, there is no need to define the whole mesh.

3.2. Bayesian Neural Networks basics

Bayesian Neural Networks (BNNs) aim to estimate the posterior parameter distribution $P(\boldsymbol{\theta}|\mathcal{D})$, from a training dataset $\mathcal{D} = \{\mathbf{x}^{(i)}, \mathbf{y}^{(i)}\}$, of a set of NN parameters $\boldsymbol{\theta}$:

$$P(\boldsymbol{\theta}|\mathcal{D}) = \frac{P(\mathcal{D}|\boldsymbol{\theta})P(\boldsymbol{\theta})}{P(\mathcal{D})} \quad (20)$$

where $P(\mathcal{D}|\boldsymbol{\theta})$ is the likelihood, $P(\boldsymbol{\theta})$ is the prior, *i.e.* prior NN parameters knowledge expressed as a PDF over $\boldsymbol{\theta}$, and $P(\mathcal{D})$ is the marginal likelihood.

The likelihood is often estimated the individual product of pointwise estimated likelihoods $p_i(\mathcal{D}|\boldsymbol{\theta})$ based on the normal distribution $N(\mu, \sigma)$ defined as follows:

$$P(\mathcal{D}|\boldsymbol{\theta}) = \prod_{i=0}^N p(y^{(i)}|x^{(i)}, \boldsymbol{\theta}) \quad (21)$$

$$p(y^{(i)}|x^{(i)}, \boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\|x_i - x(t_i)\|^2}{2\sigma^2}\right) \quad (22)$$

The selection of a good prior distribution, $P(\boldsymbol{\theta})$, for BNNs is challenging (Fortuin, Garriga-Alonso, van der Wilk, & Aitchison, 2021). The most common choice is to use a non-informative prior, which follows a Gaussian distribution with mean zero and unit variance, $\mathcal{N}(0, 1)$, *i.e.* isotropic Gaussian prior. In addition to the isotropic Gaussian prior, alternative priors implemented in this work include *Spike-and-Slab* (SS) prior (Williams, 1995), which is defined as follows:

$$P(\boldsymbol{\theta})_{SS} = \prod_{i=1}^d [\pi \mathcal{N}(\theta_i; 0, \sigma_1^2) + (1-\pi) \mathcal{N}(\theta_i; 0, \sigma_2^2)] \quad (23)$$

where $\pi \in (0, 1)$ is the mixing coefficient, σ_1 and σ_2 correspond to the variance of the two Gaussian distributions, with $\sigma_1^2 \gg \sigma_2^2$, and d is the dimensionality of $\boldsymbol{\theta}$. The first component (*slab*) is a wide Gaussian that allows for larger parameter values, while the second (*spike*) is a narrow Gaussian centered at zero, encouraging sparsity. This prior combines flexibility with regularization by allowing most weights to be close to zero while allowing a few to have large magnitudes.

In addition, the *Laplace prior* will be also implemented (Williams, 1995), defined as follows:

$$P(\boldsymbol{\theta})_{Laplace} = \prod_{i=1}^d \frac{\lambda}{2} \exp(-\lambda|\theta_i|) \quad (24)$$

where $\lambda > 0$ is the scale parameter, and d is the dimensionality of $\boldsymbol{\theta}$. The Laplace prior also promotes sparsity due to its sharp peak at zero and heavier tails compared to the non-informative Gaussian prior.

The analytical posterior in BNNs, $p(\boldsymbol{\theta}|\mathcal{D})$, is intractable, and therefore approximation methods are required. In this work, variational inference (VI) is used to approximate the true posterior with a variational distribution $q(\boldsymbol{\theta}|\phi)$ of known functional form by minimizing the Kullback-Leibler (KL) divergence between them. Since the true posterior cannot be computed directly, this minimization is performed indirectly by maximizing the Evidence Lower Bound (ELBO), which includes a KL divergence term between the variational distribution and the prior $p(\boldsymbol{\theta})$. The corresponding optimization objective is defined as follows:

$$\mathcal{L}(\mathcal{D}, \phi) = \text{KL}[q(\boldsymbol{\theta} | \phi) \parallel p(\boldsymbol{\theta})] - \mathbb{E}_{q(\boldsymbol{\theta}|\phi)}[\log p(\mathcal{D} | \boldsymbol{\theta})] \quad (25)$$

The first term is the KL divergence between the variational distribution $q(\boldsymbol{\theta}|\phi)$ and the prior $p(\boldsymbol{\theta})$ (known as the complexity cost). The second term is the expected

value of the likelihood with respect to the variational distribution (known as the likelihood cost). The cost function can be also written as:

$$\mathcal{L}(\mathcal{D}, \phi) = \mathbb{E}_{q(\theta|\phi)} [\log q(\theta|\phi)] - \mathbb{E}_{q(\theta|\phi)} [\log p(\theta)] - \mathbb{E}_{q(\theta|\phi)} [\log p(\mathcal{D}|\theta)] \quad (26)$$

It can be observed from Eq. (26) that loss terms are expectations with respect to the variational distribution $q(\theta|\phi)$. Therefore, the cost function can be approximated by Monte Carlo sampling, drawing N samples $\theta^{(i)}$ from $q(\theta|\phi)$:

$$\mathcal{L}(\mathcal{D}, \phi) \approx \frac{1}{N} \sum_{i=1}^N [\log q(\theta^{(i)}|\phi) - \log p(\theta^{(i)}) - \log p(\mathcal{D}|\theta^{(i)})] \quad (27)$$

Throughout this paper, the variational posterior is assumed to follow a Gaussian distribution $\phi = (\mu, \sigma)$, where μ is the mean vector of the distribution and σ is the standard deviation vector.

3.3. Bayesian PINNs

Bayesian PINNs (B-PINNs) extend standard PINNs by replacing the deterministic NN with a Bayesian NN. This enables the estimation of epistemic uncertainty, which reflects the model's lack of complete knowledge (Kendall & Gal, 2017). In practice, epistemic uncertainty is quantified by placing probability density functions (PDFs) over the BNN parameters.

The main outcome of interest in the B-PINN approach is the posterior distribution. This is calculated following the Bayes rule (cf. Eq. (20)). Namely, the prior distribution is assumed to be given based on *a priori* knowledge (i.e., Gaussian, Spike-and-Slab, or Laplace). The likelihood, $P(\mathcal{D}|\Theta)$, is estimated for the initial point, boundary condition points, and residual samples assuming a Gaussian distribution (cf. Eq. (22)), defined as follows:

$$\begin{aligned} p(u_0|x_0, \theta^{(i)}) &= \prod_{i=1}^{N_0} \frac{1}{\sqrt{2\pi}\sigma_0} \exp\left(-\frac{\|\hat{u}(x_0, 0; \theta^{(i)}) - u_0\|^2}{2\sigma_0^2}\right) \\ p(u_{bc}|x_{bc}, t_{bc}, \theta^{(i)}) &= \prod_{i=1}^{N_{bc}} \frac{1}{\sqrt{2\pi}\sigma_{bc}} \exp\left(-\frac{\|\hat{u}(x_{bc}, t_{bc}; \theta^{(i)}) - u_{bc}\|^2}{2\sigma_{bc}^2}\right) \\ p(r(x_f, t_f; \theta^{(i)})) &= \prod_{i=1}^{N_f} \frac{1}{\sqrt{2\pi}\sigma_f} \exp\left(-\frac{\|r(x_f, t_f; \theta^{(i)})\|^2}{2\sigma_f^2}\right) \end{aligned} \quad (28)$$

where σ_0 , σ_{bc} , and σ_f denote the standard deviation of initial, boundary, and residual points.

For computational tractability and efficiency, the log-

likelihood terms are considered by taking the logarithm of Eq. (28). Subsequently, the Monte Carlo ELBO loss defined for a generic BNN [cf. Eq. (27)], is adapted for B-PINN posterior inference:

$$\begin{aligned} \mathcal{L}^{(i)} &= \log q(\theta^{(i)}|\mathbf{w}) - \log p(\theta^{(i)}) - \lambda_0 \log p(u_0|x_0, \theta^{(i)}) \\ &\quad - \lambda_b \log p(u_{bc}|x_{bc}, t_{bc}, \theta^{(i)}) - \lambda_r \log p(r|x_f, t_f, \theta^{(i)}) \end{aligned} \quad (29)$$

This process results in the approximation of the variational posterior distribution $q(\theta|\mathbf{w})$. The training process of the B-PINN approach is shown in Algorithm 1.

Algorithm 1 B-PINN Training via Variational Inference

- 1: **Input:** Collocation points (x_f, t_f) , initial condition data (x_0, u_0) , boundary condition data (x_{bc}, t_{bc}, u_{bc}) , prior distribution $p(\theta)$
 - 2: Initialize variational parameters $\mathbf{w} = \{\mu, \sigma\}$ for θ
 - 3: **while** not converged **do**
 - 4: Sample $\theta^{(i)} \sim q(\theta|\mathbf{w})$ via reparameterization trick
 - 5: Evaluate predicted solution: $\hat{u}(x, t; \theta^{(i)})$
 - 6: Compute residual: $r(x_f, t_f; \theta^{(i)}) = \mathcal{N}[\hat{u}](x_f, t_f; \theta^{(i)}) - f(x_f, t_f)$
 - 7: Evaluate log-likelihood terms taking the log of Eq. (28)
 - 8: Evaluate prior log-probability: $\log p(\theta^{(i)})$
 - 9: Evaluate variational density: $\log q(\theta^{(i)}|\mathbf{w})$
 - 10: Compute Monte Carlo estimate of ELBO loss via Eq. (29)
 - 11: Update $\mathbf{w} = \{\mu, \sigma\}$ using gradient $\nabla_{\mathbf{w}} \mathcal{L}^{(i)}$
 - 12: **Return:** Learned variational posterior $q(\theta|\mathbf{w})$
-

4. CASE STUDY & EXPERIMENTS

The case study examines the lifetime of a distribution transformer installed in a floating solar power plant in Spain (Aizpurua et al., 2023). The transformer's nameplate parameters are summarized in Table 2.

Table 2. Transformer parameter values.

Parameter	Value
Rating [kVA], V_1/V_2	1100, 22000/400
R=Load losses/No load losses [W]	9800/842
$\Delta\Theta_{H,R}$ [$^{\circ}\text{C}$]	15.1
k_{21}, k_{22}	2.32, 2.05
τ_0, τ_w [min.]	266.8, 9.75

The dataset was preprocessed by removing variables without measurements, duplicate entries, and incorrect sensor readings. Missing values were imputed with average values. Figure 4 shows the available minutely sampled time-series for ambient temperature, oil temperature, and load with a total of 5760 samples over 4 days of operation. Oil and winding temperature profiles analysed

in this study are based on this minute-level sampling.




Figure 4. Available load (K), ambient temperature (Θ_A), and top-oil temperature (Θ_{TO}) over four days of operation with one-minute resolution.

Since the PDE solution is unknown, no validation dataset is available for hyperparameter tuning. This makes balancing the loss function particularly challenging (S. Wang, Sankaran, Wang, & Perdikaris, 2023). To assess learning and generalization capabilities, the model is evaluated under varying numbers of initial condition samples N_i (known values) and residual points N_r (PDE evaluation points, unknown values), as detailed in Section 4.3.

4.1. Evaluation Metrics

The accuracy of probabilistic predictions is evaluated using three complementary metrics.

Negative Log Likelihood (NLL) evaluates how well a probabilistic model explains a given set of observations. It is defined as the negative logarithm of the likelihood function, which measures the probability of the observed data under a particular model (Murphy, 2012). For a model with parameters θ and observed data $X = \{x_1, \dots, x_n\}$, NLL is defined as:

$$\text{NLL}(\theta) = - \sum_{i=1}^n \log p(x_i | \theta) \quad (30)$$

where $p(x_i | \theta)$ is the probability of observation x_i given the model parameters. Minimizing the NLL is equivalent to maximizing the likelihood function, as the logarithm is a monotonically increasing function. NLL is particularly useful because it avoids numerical underflow with small probabilities and transforms the product of probabilities into a sum of log probabilities, which is more computationally stable (Bishop & Nasrabadi, 2006).

Continuous Ranked Probability Score (CRPS) measures the discrepancy between the predicted Cumulative Distribution Function (CDF), $F(\cdot)$, and the observed empirical CDF for a given scalar observation y (Zamo & Naveau, 2018):

$$\text{CRPS}(F, y) = \int (F(x) - \mathbb{1}(x \geq y))^2 dx, \quad (31)$$

where $\mathbb{1}(x \geq y_i)$ is the indicator function, which models the empirical CDF.

In order to obtain a single score value from Eq. (31), a weighted average is computed for each individual observation of the test set (Gneiting, Raftery, Westveld, & Goldman, 2005):

$$\text{CRPS} = \frac{1}{N} \sum_{i=1}^N \text{CRPS}(F_i, y_i) \quad (32)$$

where N denotes the total number of predictions.

Prediction Interval Coverage Probability (PICP) quantifies the reliability of prediction intervals by evaluating the fraction of true observations that fall within the predicted interval. PICP evaluates the effectiveness of the prediction interval in capturing the variability of the data, reflecting the accuracy and validity of the predictions (González-Sopeña, Pakrashi, & Ghosh, 2021):

$$\text{PICP} = \frac{1}{N} \sum_{i=1}^N \hat{\Theta}_{O_i} \quad (33)$$

where N denotes the total number of predictions and $\hat{\Theta}_{O_i}$ individual observations in the prediction interval.

4.2. Benchmarking

The proposed B-PINN approach is benchmarked against a probabilistic hybrid method. Namely, dropout based PINN implementation has been selected due to its computational efficiency (Zhang et al., 2019).

Dropout-PINN

Dropout-PINN extends the vanilla PINN (Section 3.1) by inserting dropout layers after each hidden layer, with dropout rate ρ . During training, neurons are randomly deactivated, improving generalization and reducing overfitting (Gal & Ghahramani, 2016). During inference, K stochastic forward passes are performed, each with different neuron subsets deactivated. The resulting ensemble of predictions captures predictive uncertainty (Zhang et al., 2019).

The dropout rate strongly affects the uncertainty quality (Alcibar et al., 2025) and therefore a sensitivity analysis is conducted to select the best dropout rate.

4.3. Experiments and Implementation Details

The base-case B-PINN uses two layers with 50 neurons each, $N_i=100$ initial points, $N_{bc}=11520$ boundary condition points, and $N_r=10000$ residual points. Input data noise and residual data noise follow $\sigma_i=\mathcal{N}(0, 0.01)$ and $\sigma_r=\mathcal{N}(0, 0.01)$, respectively.

Starting from the base-case, different hyperparameters are evaluated including the number of neurons and the number of initial and residual points. After identifying the best-performing configuration, experimental results are obtained for the selected configuration. In addition to the hyperparameters, three different priors are tested: Gaussian, Spike-and-Slab (cf. Eq. (23)), and Laplace (cf. Eq. (24)). Table 3 summarizes the configurations and corresponding hyperparameters.

Table 3. Base-case and tested hyperparameter ranges for B-PINN models.

Parameters	Configurations	
	Base-Case	Range
Neurons	50	[10,25,50,100]
N_i	100	[5,100,200]
N_r	10000	[5000,10000, 20000]
σ_i	0.01	[0.01, 0.05, 0.1]
σ_r	0.01	[0.01, 0.05, 0.1]

All experiments use variational inference to approximate the posterior distribution. Each simulation is repeated five times for numerical stability, and mean and standard deviation values are reported. All models are implemented from scratch in PyTorch. The optimization algorithm is Adam (learning rate = 0.01, batch size = 16) and all B-PINN models are trained with 15000 epochs. The weights of the individual loss terms have been manually weighted, with final values of $\lambda_0=1$, $\lambda_b=1$, and $\lambda_r=10^{-6}$.

The vanilla PINN model is implemented using the Adam optimization algorithm for the first 20000 iterations, followed by the L-BFGS optimizer for the last 10000 iterations. The terms of individual loss terms have been manually weighted with final values of $\lambda_0=1$, $\lambda_b=1$, and $\lambda_r=10^{-6}$. Finally, the dropout-PINN model is designed taking the vanilla PINN architecture as the reference model, including additional dropout layers. The dropout rate is tested for a range of values $\rho=[0.05, 0.1, 0.15, 0.2, 0.25]$, with $K=200$ Monte Carlo iterations.

5. RESULTS & DISCUSSION

Figure 5 shows the numerical solution of Eqs. (1)–(4), solved using finite element method (FEM) with Matlab’s pdepe solver (The MathWorks, Inc., 2023). The solution is experimentally validated against fiber optic sensor measurements installed at different transformer heights. Refer to (Ramirez et al., 2025) for more details.




Figure 5. Spatiotemporal transformer oil temperature estimation obtained through a finite element method.

The oil temperature values, $\hat{\Theta}_O(x, t)$, follow the seasonality of solar energy, *i.e.* peak values coincide with periods of maximum solar irradiation and vice-versa (cf. Figure 4). As for vertical temperature resolution, the highest temperatures are observed near the top position ($x \rightarrow 1$).

The validated FEM model in Figure 5 serves as ground truth. Next, starting from the base-case configuration, a detailed hyperparameter tuning evaluation will be performed to select an appropriate B-PINN configuration.

5.1. Hyperparameter tuning

Neurons and prior distributions. The influence of the network width (number of neurons) and prior distribution is first examined with the remaining parameters fixed according to the base-case configuration (cf. Subsection 4.3). Table 4 reports the results.

Table 4 shows that the best performance is obtained with 50 neurons, with the Laplace prior distribution. This setting achieves the lowest CRPS (0.1156 ± 0.0242) and NLL (0.0179 ± 0.3536), indicating both accurate and sharp predictive distributions. At the same time, it maintains a high PICP (0.3501 ± 0.0327), confirming that predictive intervals are well-calibrated and consistent with the observed data. These results suggest that the Laplace prior, in combination with a moderate network width, provides an adequate balance between predictive accuracy and reliable UQ.

Across all priors, increasing the number of neurons improves performance up to a point. Specifically, moving from 10 to 50 neurons leads to substantial improvements across all UQ metrics, producing sharper distributions, better likelihood scores, and higher coverage. However, further increasing the width to 100 neurons does not consistently improve results. While PICP occasionally increases, this often comes at the cost of higher CRPS and NLL values, indicating a deterioration in the overall quality of the predictive distributions.

Table 4. Comparison of different number of neurons and prior distributions for different UQ metrics with remaining parameters fixed (2 layers, $N_i=100$, $N_r=10000$, $\sigma_i=\sigma_r=\mathcal{N}(0, 0.01^2)$). Best results highlighted.

#N	Gaussian			Spike-and-Slab			Laplace		
	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)
10	0.078 $\pm$.0225	0.287 $\pm$.051	5.774 $\pm$ 2.850	0.067 $\pm$.012	0.365 $\pm$.096	15.898 $\pm$ 11.827	0.071 $\pm$.018	0.297 $\pm$.075	11.933 $\pm$ 5.840
25	0.197 $\pm$.028	0.180 $\pm$.045	0.445 $\pm$.608	0.159 $\pm$.043	0.130 $\pm$.021	-0.063 $\pm$ 0.136	0.170 $\pm$.016	0.153 $\pm$.034	0.136 $\pm$.345
50	0.377 $\pm$.034	0.159 $\pm$.042	0.194 $\pm$.323	0.351 $\pm$.036	0.139 $\pm$.054	-0.023 $\pm$.336	0.350 $\pm$.032	0.115 $\pm$.024	0.017 $\pm$.353
100	0.465 $\pm$.018	0.238 $\pm$.125	0.927 $\pm$ 1.066	0.495 $\pm$.058	0.278 $\pm$.205	0.654 $\pm$.776	0.463 $\pm$.083	0.204 $\pm$.048	0.402 $\pm$.255

Legend. #N: Number of neurons.

Table 5. Comparison of different number of samples and prior distributions for different UQ metrics with remaining parameters fixed (2 layers, 50 neurons, $\sigma_i = \sigma_r = \mathcal{N}(0, 0.01^2)$). Best results highlighted.

# Samples		Gaussian			Spike-and-Slab			Laplace		
N_i	N_r	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)
I_1	R_1	0.320 $\pm$.027	0.208 $\pm$.052	0.451 $\pm$.255	0.334 $\pm$.060	0.233 $\pm$.083	0.641 $\pm$.569	0.390 $\pm$.059	0.178 $\pm$.030	0.338 $\pm$.127
	R_2	0.409 $\pm$.052	0.202 $\pm$.022	0.390 $\pm$.083	0.363 $\pm$.040	0.184 $\pm$.058	0.261 $\pm$.237	0.450 $\pm$.018	0.177 $\pm$.045	0.331 $\pm$.178
	R_3	0.387 $\pm$.099	0.192 $\pm$.043	0.325 $\pm$.216	0.511 $\pm$.070	0.175 $\pm$.038	0.227 $\pm$.203	0.541 $\pm$.043	0.169 $\pm$.022	0.208 $\pm$.150
I_2	R_1	0.282 $\pm$.044	0.157 $\pm$.035	0.173 $\pm$.183	0.322 $\pm$.042	0.187 $\pm$.048	0.358 $\pm$.226	0.305 $\pm$.032	0.113 $\pm$.027	-0.096 $\pm$.208
	R_2	0.363 $\pm$.022	0.163 $\pm$.046	0.236 $\pm$.349	0.351 $\pm$.036	0.139 $\pm$.054	-0.023 $\pm$.336	0.360 $\pm$.027	0.108 $\pm$.022	-0.145 $\pm$.150
	R_3	0.448 $\pm$.097	0.119 $\pm$.023	-0.202 $\pm$.126	0.375 $\pm$.076	0.150 $\pm$.050	0.080 $\pm$.346	0.377 $\pm$.040	0.128 $\pm$.033	-0.112 $\pm$.194
I_3	R_1	0.281 $\pm$.044	0.125 $\pm$.032	-0.058 $\pm$.307	0.234 $\pm$.022	0.109 $\pm$.061	-0.063 $\pm$.720	0.289 $\pm$.036	0.134 $\pm$.046	-0.009 $\pm$.247
	R_2	0.416 $\pm$.041	0.092 $\pm$.026	-0.374 $\pm$.218	0.348 $\pm$.083	0.131 $\pm$.045	0.127 $\pm$.733	0.324 $\pm$.034	0.077 $\pm$.009	-0.478 $\pm$.071
	R_3	0.435 $\pm$.072	0.166 $\pm$.051	0.455 $\pm$.590	0.390 $\pm$.082	0.121 $\pm$.048	-0.113 $\pm$.414	0.420 $\pm$.042	0.137 $\pm$.034	-0.050 $\pm$.336

Legend. # of initial samples (N_i): $I_1=5$, $I_2=100$, $I_3=200$. # of residual samples (N_r): $R_1=500$, $R_2=10000$, $R_3=20000$.

In summary, the results highlight the need for a trade-off when selecting the network width with respect to the prior distribution. In this case, the configuration with 50 neurons under the Laplace prior provides the most robust compromise, offering strong predictive performance together with well-calibrated uncertainty estimates.

Number of initial (N_i), residual samples (N_r), and prior distributions. Table 5 presents the sensitivity of the B-PINN model to the number of initial conditions samples, residual points, and boundary conditions, evaluated under different prior distributions using probabilistic metrics (cf. Subsection 4.3).

The results indicate that increasing the number of initial condition samples generally improves the quality of uncertainty estimates across all priors. For instance, under the Gaussian prior, increasing N_i from 5 to 200 (with fixed $N_r=10000$) reduces the NLL from 0.390 ± 0.083 to -0.374 ± 0.218 , while also lowering the CRPS. This pattern is consistent across priors, suggesting that larger N_i values leads to more informative predictive distribution, reflected in higher coverage (PICP) and improved probabilistic accuracy (CRPS, NLL).

In contrast, increasing the number residual points does not consistently enhance performance. For example, under the spike-and-slab prior with $N_i=5$, increasing N_r from 10000 to 20000 results in only a marginal NLL improvement (0.261 ± 0.237 to 0.227 ± 0.203), with similar trends observed for other priors. These results suggest that excessively increasing the residual samples may reduce the relative influence of initial and boundary data,

potentially degrading uncertainty calibration.

Under the adopted loss weighing scheme (cf. Eq. (29), Algorithm 1), this trade-off highlights the need for a balanced allocation between residual points and initial/boundary condition samples. Residual points remain essential to enforce PDE dynamics, but their dominance over data-driven terms can compromise calibration. Hence, an appropriate trade-off between N_r and N_i is critical to achieve reliable predictive performance.

The best results are achieved with 200 initial samples and 10000 residual points, using the Laplace prior distribution. Consequently, subsequent experiments will be performed with the selected configuration.

5.2. Experimental Results

First, deterministic spatiotemporal transformer oil temperature estimates are obtained using the vanilla PINN (Section 3.1) following the implementation in (Ramirez et al., 2025). Figure 6(a) shows the spatiotemporal temperature estimate and Figure 6(b) shows the error with respect to the ground truth (Figure 5).

To compare the B-PINN outcome with a probabilistic estimate, a dropout-PINN model has been implemented following the PINN architecture (3 layers, 20 neurons each). Table 6 shows the results for different dropout rates.

The best-performing dropout-PINN is obtained with a dropout rate of 0.05. Higher dropout rates achieve better PICP values, but at the cost of degraded CRPS and




Figure 6. Spatiotemporal temperature estimation results obtained using a vanilla PINN. (a) Predicted temperature distribution across the transformer oil domain and (b) corresponding absolute error.

Table 6. Sensitivity of the dropout rate in dropout-PINN. Best results highlighted.

Dropout Rate	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)
0.05	0.714 $\pm$ 0.045	0.722 $\pm$ 0.044	1.955 $\pm$ 0.159
0.1	0.776 $\pm$ 0.061	0.853 $\pm$ 0.031	2.010 $\pm$ 0.143
0.15	0.805 $\pm$ 0.071	0.967 $\pm$ 0.060	2.248 $\pm$ 0.265
0.2	0.787 $\pm$ 0.013	1.069 $\pm$ 0.036	2.386 $\pm$ 0.157
0.25	0.724 $\pm$ 0.048	1.214 $\pm$ 0.019	2.731 $\pm$ 0.226

NLL performance. Accordingly, the best dropout-PINN and B-PINN models were selected for comparison. Figure 7 shows the resulting probabilistic spatiotemporal temperature estimates parameterized through the mean and standard deviation.

Compared with the vanilla PINN (Figure 6), the B-PINN is able to capture variation in the mean estimates, thereby quantifying prediction confidence at each spatiotemporal coordinate. The mean B-PINN predictions exhibit higher accuracy than the dropout-PINN, as also reflected in the error maps of Figure 8(a) and (c). Furthermore, the variance of B-PINN predictions is consistently lower than

dropout-PINN, confirmed in Figures 8(b) and (d), which illustrate the standard deviation of the prediction error relative to the ground truth (cf. Figure 5).

Figure 9 compares probabilistic spatiotemporal temperature estimates for the B-PINN, dropout-PINN, and vanilla PINN at selected time instants. The B-PINN consistently provides more accurate estimates than both alternatives. Importantly, the B-PINN not only delivers mean predictions, but also quantifies the uncertainty of the estimates. In most cases, the ground truth falls within the confidence intervals, in line with the PICP metric. The uncertainty bands are adaptive, reflecting prediction confidence that varies across the spatiotemporal domain. In contrast, the dropout-PINN produces wider and less consistent uncertainty bands, which would lead to overly conservative decision-making.

Building on the probabilistic temperature field, Figure 10 shows the spatiotemporal ageing estimates $p(\hat{V}(x, t))$ obtained with the B-PINN, including both the mean and standard deviation. After 4 days of operation under the loading profile of Figure 4, the maximum transformer loss-of-life (cf. Eq. (12)) is 8.964 minutes according to the B-PINN mean estimate (Figure 10(a)), with a maximum deviation of 0.985 minutes (Figure 10(b)). This probabilistic ageing estimation framework supports informed maintenance decisions by balancing worst-case scenarios with less conservative, risk-adjusted strategies.

The ageing estimation errors of different models were further evaluated relative to the FEM-based ageing estimates. Figure 11 shows the results for vanilla PINN, dropout-PINN, and B-PINN models. The B-PINN achieved the lowest error, with a worst-case ageing estimation error of 0.35 minutes (4.06%) (Figure 11(a)). The dropout-PINN and vanilla PINN achieved worst-case errors of 1.25 minutes (14.22%) and 1.28 minutes (14.57%), respectively (Figure 11(b)). The small standard deviation of the B-PINN ageing estimates (Figure 10(b)) further indicates that prediction variability is limited, avoiding the need for conservative decisions.

Overall, the probabilistic spatiotemporal ageing estimates produced by the B-PINN provide a more informative basis for asset health monitoring. On one hand, the spatiotemporal ageing maps enable maintenance planning that accounts for localized ageing effects, improving conventional practices based on worst-case, spatially-agnostic estimates. On the other hand, the explicit quantification of predictive uncertainty allows operators to assess reliability and adapt decisions accordingly, ranging from conservative interventions in uncertain scenarios to optimized actions in reliable regimes.




Figure 7. Spatiotemporal transformer oil temperature estimation using the B-PINN model (a) posterior mean; (b) standard deviation; and dropout-PINN model (c) posterior mean, (d) posterior standard deviation.

Table 7. Comparison of different noise levels and prior distributions for different UQ metrics with remaining parameters fixed (2 layers, 50 neurons, $N_i=200$, $N_r=10000$). Best results highlighted.

Noise Level	Gaussian			Spike-and-Slab			Laplace		
	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)	PICP ($\uparrow$)	CRPS ($\downarrow$)	NLL ($\downarrow$)
$\sigma_i=\sigma_r=\mathcal{N}(0,0.01^2)$	0.416 $\pm$.041	0.092 $\pm$.026	-0.374 $\pm$.218	0.348 $\pm$.083	0.131 $\pm$.045	0.127 $\pm$.733	0.324 $\pm$.034	0.077 $\pm$.009	-0.478 $\pm$.071
$\sigma_i=\sigma_r=\mathcal{N}(0,0.05^2)$	0.242 $\pm$.029	0.151 $\pm$.050	0.192 $\pm$.237	0.266 $\pm$.048	0.225 $\pm$.024	0.482 $\pm$.161	0.245 $\pm$.025	0.179 $\pm$.014	0.285 $\pm$.105
$\sigma_i=\sigma_r=\mathcal{N}(0,0.1^2)$	0.191 $\pm$.037	0.179 $\pm$.041	0.351 $\pm$.125	0.215 $\pm$.026	0.275 $\pm$.097	0.957 $\pm$.805	0.216 $\pm$.043	0.202 $\pm$.011	0.441 $\pm$.050

5.3. Robustness to Noise

After observing the probabilistic prediction results, the effect of different noise levels on the best-performing model (*i.e.* B-PINN) will be examined. Namely, noise was introduced in the initial condition samples (σ_i) and residual samples (σ_r), and the resulting performance was evaluated under three prior configurations (Gaussian, spike-and-slab, and Laplace), while keeping all other base-case configuration parameters fixed (cf. Subsection 4.3). The results are summarized in Table 7.

The results indicate a deterioration in probabilistic performance as the noise level increases. Higher noise reduces the PICP values, indicating a smaller proportion of true values fall within the predicted intervals, which compromises the reliability of the uncertainty estimates. At the same time, CRPS and NLL increase with noise,

indicating a decline in the quality and calibration of the predictive distributions. Specifically, higher CRPS indicates less accurate probabilistic forecasts, while higher NLL values show that the model assigns lower likelihood to observed outcomes.

These trends are consistent across all priors, but their magnitudes are different. The Gaussian prior achieves the highest PICP values, but simultaneously produces worse NLL scores compared to the Laplace prior. This suggests that the Gaussian prior generates wider predictive intervals, capturing more true values but at the expense of less accurate probability assignment. Conversely, the Laplace prior provides better calibration (lower NLL and CRPS) despite offering narrower intervals and lower coverage. This illustrates the trade-off between interval coverage (PICP) and calibration qual-




Figure 8. Spatiotemporal transformer oil temperature estimation error of the B-PINN model (a) mean absolute error and (b) standard deviation; and dropout PINN model (c) mean absolute error and (d) standard deviation.

ity (NLL, CRPS), showing that a higher PICP may not imply superior probabilistic performance.

Overall, the results confirm that increasing noise in both initial and residual data degrades the probabilistic performance and calibration of the B-PINN model. These findings emphasize the importance of carefully accounting for data noise and quality in UQ tasks.

5.4. Computational Cost

Table 8 reports the computational cost of the different models, obtained on an AMD 4900HS processor with 32GB of RAM for 466560 evaluation points (81 in space, 5760 in time). The dropout-PINN achieves the shortest training time because randomly deactivating neurons reduces the effective model size during optimization. In contrast, the vanilla PINN requires more training time, and the B-PINN incurs additional cost due to uncertainty quantification. The FEM does not require training.

During evaluation, the vanilla PINN serves as the baseline without UQ and therefore avoids additional computational overhead. The dropout-PINN requires multiple stochastic forward passes to estimate predictive distributions, leading to higher evaluation time than vanilla

Table 8. Computational cost of the evaluated models. Train and evaluation times averaged over repeated runs.

Model	Train [s]	Evaluation [s]
Vanilla-PINN	719.62 ± 8.64	3.73 ± 0.05
Dropout-PINN	685.66 ± 4.06	110.22 ± 0.64
B-PINN	785.72 ± 24.12	498.38 ± 2.87
FEM	25.83 ± 3.69	

PINN, but lower than B-PINN. The B-PINN incurs the highest evaluation cost because posterior inference requires repeated sampling across weight distributions. The FEM is the most efficient when training and evaluation are considered jointly. However, for more complex PDEs, the computational cost of FEM increases rapidly and PINNs provide a scalable and flexible alternative.

6. CONCLUSIONS

Scientific Machine Learning (SciML) combines physical knowledge with machine learning (ML), resulting in predictive models that generalize better than purely data-driven approaches and require less data. Hybrid Prognostics and Health Management (PHM) methods share similarities with SciML, as both integrate physics-of-failure models and data. However, SciML provides a more prin-




Figure 9. Spatiotemporal temperature estimates at different time instants for vanilla PINN, dropout-PINN, and Bayesian PINN (mean and standard deviation) models and ground truth values.

cipled learning framework that, to the best of the authors' knowledge, has not been fully explored within PHM.

This work introduces a Bayesian Physics-Informed Neural Network (B-PINN) framework for probabilistic ageing estimation in electrical transformers. The method integrates the heat diffusion partial differential equation with transformer loading and temperature data, while modeling neural network weights probabilistically to capture epistemic uncertainty. Several prior distributions (Gaussian, spike-and-slab, Laplace) were evaluated, along with a range of B-PINN configurations (varying network width, number of initial and residual samples, and noise levels), and a dropout-PINN baseline for comparison.

Numerical results demonstrate that B-PINNs improve reliability over deterministic PINNs by providing both accurate predictions and well-calibrated uncertainty estimates. Compared with dropout-PINNs, B-PINNs yield more consistent and sharper predictive intervals, reducing the tendency toward overly conservative estimates. Among the priors, the Laplace prior consistently delivers the best performance across different settings, while the

Gaussian prior performs robustly under high noise conditions. The spike-and-slab prior, in contrast, is more sensitive to data quality. Sensitivity analysis indicates that noise levels and the number of initial condition samples have the greatest influence on uncertainty quantification, whereas the number of residual samples has a comparatively lower influence.

Beyond methodological contributions, the proposed B-PINN framework offers practical benefits for transformer asset management. By quantifying and propagating epistemic uncertainty, the method supports informed decision-making under data scarcity, mitigating the risks of overly conservative or unsafe maintenance strategies.

The present implementation employs variational inference for posterior estimation. Future research will explore alternative Bayesian inference methods and physics-informed priors to further enhance uncertainty quantification in SciML-based PHM frameworks.

ACKNOWLEDGEMENTS

This work has been partially funded by the Spanish State Research Agency (grant No. PID2024-156284OA-




Figure 10. Probabilistic spatiotemporal transformer ageing $p(\hat{V}(x, t))$: (a) mean values (b) standard deviation.

I00) and the Basque Government (grants IT1504-22, KK-2024/00030). Jose I. Aizpurua is funded by the Ramón y Cajal Fellowship, Spanish State Research Agency (grant No. RYC2022-037300-I), co-funded by MCIU/AEI/10.13039/501100011033 and FSE+.

REFERENCES

- Aizpurua, J. I., & Catterson, V. M. (2015). Towards a methodology for design of prognostic systems. In *Annual Conference of the PHM Society* (Vol. 7). doi: 10.36001/phmconf.2015.v7i1.2646
- Aizpurua, J. I., Ramirez, I., Lasa, I., Rio, L. d., Ortiz, A., & Stewart, B. G. (2023). Hybrid transformer prognostics framework for enhanced probabilistic predictions in renewable energy applications. *IEEE Trans. Pow. Del.*, 38(1), 599-609. doi: 10.1109/TPWRD.2022.3203873
- Alcibar, J., Aizpurua, J. I., & Zugasti, E. (2024). Towards a probabilistic fusion approach for robust battery prognostics. In *PHM Society European Conference* (Vol. 8). doi: 10.36001/phme.2024.v8i1.3981
- Alcibar, J., Aizpurua, J. I., Zugasti, E., & Peñagarikano, O. (2025). A hybrid probabilistic battery health management approach for robust inspection drone operations. *Engineering Applications of Artificial Intelligence*, 146, 110246. doi: 10.1016/j.engappai.2025.110246
- Anagnostopoulos, S. J., Toscano, J. D., Stergiopoulos, N., & Karniadakis, G. E. (2024). Residual-based attention in physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421, 116805. doi: 10.1016/j.cma.2024.116805
- Arias Chao, M., Kulkarni, C., Goebel, K., & Fink, O. (2022). Fusing physics-based and deep learning models for prognostics. *Reliability Engineering & System Safety*, 217, 107961. doi: 10.1016/j.ress.2021.107961
- Basora, L., Viens, A., Chao, M. A., & Olive, X. (2025). A benchmark on uncertainty quantification for deep learning prognostics. *Reliability Engineering & System Safety*, 253, 110513.
- Biggio, L., Wieland, A., Chao, M. A., Kastanis, I., & Fink, O. (2021). Uncertainty-aware prognosis via deep gaussian process. *IEEE Access*, 9, 123517–123527.
- Bishop, C. M., & Nasrabadi, N. M. (2006). *Pattern recognition and machine learning* (Vol. 4) (No. 4). Springer. doi: 10.1117/1.2819119
- Daigle, M., Roychoudhury, I., & Bregon, A. (2015). Model-based prognostics of hybrid systems. In *Annual Conference of the PHM Society* (Vol. 7). doi: 10.36001/phmconf.2015.v7i1.2586
- Fernández, J., Chiachío, J., Chiachío, M., Barros, J., & Corbetta, M. (2023). Physics-guided bayesian neural networks by abc-ss: Application to reinforced concrete columns. *Engineering Applications of Artificial Intelligence*, 119, 105790. doi: 10.1016/j.engappai.2022.105790
- Flores, P., Graf, O., Protopapas, P., & Pichara, K. (2025). Improved uncertainty quantification in physics-informed neural networks using error bounds and solution bundles. In *The Conference on Uncertainty in Artificial Intelligence*. doi: 10.48550/arXiv.2505.06459
- Fortuin, V., Garriga-Alonso, A., van der Wilk, M., & Aitchison, L. (2021). Bnnpriors: A library for bayesian neural network inference with different prior distributions. *Software Impacts*, 9, 100079. doi: 10.1016/j.simpa.2021.100079
- Gal, Y., & Ghahramani, Z. (2016, 20–22 Jun). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In M. F. Balcan & K. Q. Weinberger (Eds.), *Proceedings of the 33rd international conference on machine learning* (Vol. 48, pp. 1050–1059). New York, New York,




Figure 11. Transformer ageing estimation error for (a) mean B-PINN, (b) mean dropout-PINN, and (c) vanilla PINN.

- USA: PMLR. doi: 10.48550/arXiv.1506.02142
- Gao, Z., & Karniadakis, G. E. (2025). Scalable bayesian physics-informed kolmogorov-arnold networks. *arXiv preprint arXiv:2501.08501*. doi: 10.48550/arXiv.2501.08501
- Garg, S., & Chakraborty, S. (2023). Vb-deepnet: A bayesian operator learning framework for uncertainty quantification. *Engineering Applications of Artificial Intelligence*, 118, 105685. doi: 10.1016/j.engappai.2022.105685
- Gneiting, T., Raftery, A. E., Westveld, A. H., & Goldman, T. (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133(5), 1098–1118. doi: 10.1175/MWR2904.1
- González-Sopeña, J., Pakrashi, V., & Ghosh, B. (2021). An overview of performance evaluation metrics for short-term statistical wind power forecasting. *Renewable and Sustainable Energy Reviews*, 138, 110515. doi: 10.1016/j.rser.2020.110515
- Guo, J., Li, Z., & Li, M. (2020). A review on prognostics methods for engineering systems. *IEEE Transactions on Reliability*, 69(3), 1110–1129. doi: 10.1109/TR.2019.2957965
- International Electrotechnical Commission. (2018). Loading guide for oil-immersed power transformers. *IEC 60076-7*. Retrieved from <https://webstore.iec.ch/>
- Javanmardi, A., & Hüllermeier, E. (2023). Conformal prediction intervals for remaining useful lifetime estimation. *International Journal of Prognostics and Health Management*, 14(2).
- Jiang, X., Wang, X., Wen, Z., Li, E., & Wang, H. (2023). Practical uncertainty quantification for space-dependent inverse heat conduction problem via ensemble physics-informed neural networks. *International Communications in Heat and Mass Transfer*, 147, 106940. doi: 10.1016/j.icheatmasstransfer.2023.106940
- Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422–440. doi: 10.1038/s42254-021-00314-5
- Kendall, A., & Gal, Y. (2017). What uncertainties do we need in bayesian deep learning for computer vision? In I. Guyon et al. (Eds.), *Advances in neural information processing systems* (Vol. 30, p. 1–11). Curran Associates, Inc. doi: 10.48550/arXiv.1703.04977
- Kovachki, N., Li, Z., Liu, B., Aizzadenesheli, K., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2023, January). Neural operator: learning maps between function spaces with applications to pdes. *J. Mach. Learn. Res.*, 24(1). doi: 10.5555/3648699.3648788
- Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *International Conference on Neural Information Processing Systems* (p. 6405–6416). Red Hook, NY, USA: Curran Associates Inc. doi: 10.48550/arXiv.1612.01474
- Li, J., Long, X., Deng, X., Jiang, W., Zhou, K., Jiang, C., & Zhang, X. (2024). A principled distance-aware uncertainty quantification approach for enhancing the reliability of physics-informed neural network. *Reliability Engineering & System Safety*, 245, 109963. doi: 10.1016/j.ress.2024.109963
- Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., ... Anandkumar, A. (2024, May). Physics-informed neural operator for learning partial differential equations. *ACM / IMS J. Data Sci.*, 1(3). doi: 10.1145/3648506
- Linka, K., Schäfer, A., Meng, X., Zou, Z., Karniadakis, G. E., & Kuhl, E. (2022). Bayesian physics informed neural networks for real-world nonlinear dynamical systems. *Computer Methods in Applied Mechanics and Engineering*, 402, 115346. doi: 10.1016/j.cma.2022.115346
- Liu, Z., Wang, Y., Vaidya, S., Ruehle, F., Halverson, J., Soljačić, M., ... Tegmark, M. (2025). Kan:

- Kolmogorov-arnold networks. In *International Conference on Learning Representations (ICLR)*. doi: <https://doi.org/10.48550/arXiv.2404.19756>
- Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.
- Nair, A. S., Jacob, B., Howard, A. A., Drgona, J., & Stinis, P. (2025). *E-pinns: Epistemic physics-informed neural networks*. doi: [10.48550/arXiv.2503.19333](https://doi.org/10.48550/arXiv.2503.19333)
- Nascimento, R. G., Viana, F. A., Corbetta, M., & Kulrmi, C. S. (2023). A framework for li-ion battery prognosis based on hybrid bayesian physics-informed neural networks. *Scientific Reports*, *13*(1), 13856. doi: [10.1038/s41598-023-33018-0](https://doi.org/10.1038/s41598-023-33018-0)
- Nemani, V., Biggio, L., Huan, X., Hu, Z., Fink, O., Tran, A., ... Hu, C. (2023). Uncertainty quantification in machine learning for engineering design and health prognostics: A tutorial. *Mechanical Systems and Signal Processing*, *205*, 110796.
- Podina, L., Rad, M. T., & Kohandel, M. (2024). Conformalized physics-informed neural networks. In *ICLR 2024 Workshop on AI4Differential Equations In Science*. doi: [10.48550/arXiv.2405.08111](https://doi.org/10.48550/arXiv.2405.08111)
- Psaros, A. F., Meng, X., Zou, Z., Guo, L., & Karniadakis, G. E. (2023). Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons. *Journal of Computational Physics*, *477*, 111902. doi: [10.1016/j.jcp.2022.111902](https://doi.org/10.1016/j.jcp.2022.111902)
- Raissi, M., Perdikaris, P., & Karniadakis, G. (2019a). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, *378*, 686-707. doi: [10.1016/j.jcp.2018.10.045](https://doi.org/10.1016/j.jcp.2018.10.045)
- Raissi, M., Perdikaris, P., & Karniadakis, G. (2019b). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, *378*, 686-707. doi: [10.1016/j.jcp.2018.10.045](https://doi.org/10.1016/j.jcp.2018.10.045)
- Ramirez, I., Pino, J., Pardo, D., Sanz, M., del Rio, L., Ortiz, A., ... Aizpurua, J. I. (2025). Residual-based attention physics-informed neural networks for spatio-temporal ageing assessment of transformers operated in renewable power plants. *Engineering Applications of Artificial Intelligence*, *139*, 109556. doi: [10.1016/j.engappai.2024.109556](https://doi.org/10.1016/j.engappai.2024.109556)
- Sankararaman, S. (2015). Significance, interpretation, and quantification of uncertainty in prognostics and remaining useful life prediction. *Mechanical Systems and Signal Processing*, *52*, 228-247.
- The MathWorks, Inc. (2023). Pdepe toolbox [Computer software manual]. Natick, Massachusetts, USA. Retrieved from <https://mathworks.com/>
- Toscano, J. D., Oommen, V., Varghese, A. J., Zou, Z., Daryakenari, N. A., Wu, C., & Karniadakis, G. E. (2025). From pinns to pikans: Recent advances in physics-informed machine learning. *Mach. Learn. Comput. Sci. Eng.*, *1*(15). doi: [10.1007/s44379-025-00015-1](https://doi.org/10.1007/s44379-025-00015-1)
- Wang, F., Zhai, Z., Zhao, Z., Di, Y., & Chen, X. (2024). Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis. *Nature Communications*, *15*(1), 4332. doi: <https://doi.org/10.1038/s41467-024-48779-z>
- Wang, S., Sankaran, S., Wang, H., & Perdikaris, P. (2023). *An expert's guide to training physics-informed neural networks*. doi: [10.48550/arXiv.2308.08468](https://doi.org/10.48550/arXiv.2308.08468)
- Williams, P. M. (1995). Bayesian regularization and pruning using a laplace prior. *Neural Computation*, *7*(1), 117-143. doi: [10.1162/neco.1995.7.1.117](https://doi.org/10.1162/neco.1995.7.1.117)
- Wright, L. G., Onodera, T., Stein, M. M., Wang, T., Schachter, D. T., Hu, Z., & McMahon, P. L. (2022). Deep physical neural networks trained with backpropagation. *Nature*, *601*(7894), 549-555. doi: [10.1038/s41586-021-04223-6](https://doi.org/10.1038/s41586-021-04223-6)
- Xu, Y., Kohtz, S., Boakye, J., Gardoni, P., & Wang, P. (2023). Physics-informed machine learning for reliability and systems safety applications: State of the art and challenges. *Reliability Engineering & System Safety*, *230*, 108900. doi: [10.1016/j.ress.2022.108900](https://doi.org/10.1016/j.ress.2022.108900)
- Zamo, M., & Naveau, P. (2018, February). Estimation of the Continuous Ranked Probability Score with Limited Information and Applications to Ensemble Weather Forecasts. *Mathematical Geosciences*, *50*(2), 209-234. doi: [10.1007/s11004-017-9709-7](https://doi.org/10.1007/s11004-017-9709-7)
- Zhang, D., Lu, L., Guo, L., & Karniadakis, G. E. (2019). Quantifying total uncertainty in physics-informed neural networks for solving forward and inverse stochastic problems. *Journal of Computational Physics*, *397*, 108850. doi: [10.1016/j.jcp.2019.07.048](https://doi.org/10.1016/j.jcp.2019.07.048)
- Zio, E. (2022). Prognostics and health management (phm): Where are we and where do we (need to) go in theory and practice. *Reliability Engineering & System Safety*, *218*, 108119. doi: [10.1016/j.ress.2021.108119](https://doi.org/10.1016/j.ress.2021.108119)
- Zou, Z., Wang, Z., & Karniadakis, G. E. (2025). Learning and discovering multiple solutions using physics-informed neural networks with random initialization and deep ensemble. *arXiv preprint arXiv:2503.06320*. doi: [10.48550/arXiv.2503.06320](https://doi.org/10.48550/arXiv.2503.06320)