



# Physics-informed neural networks as surrogate models of hydrodynamic simulators

James Donnelly<sup>a,b,\*</sup>, Alireza Daneshkhah<sup>a</sup>, Soroush Abolfathi<sup>b</sup>

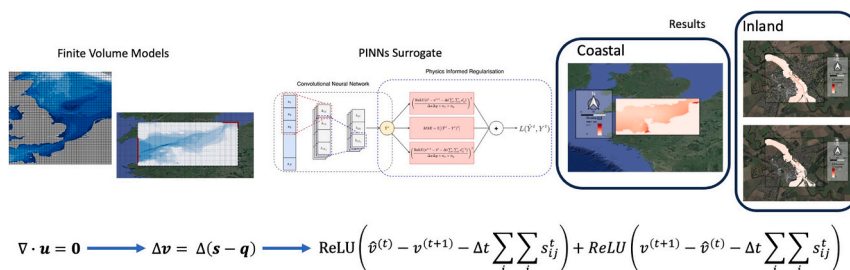
<sup>a</sup> Centre for Computational Science & Mathematical Modelling, Coventry University, UK

<sup>b</sup> School of Engineering, University of Warwick, UK

## HIGHLIGHTS

- Novel flexible PINNs model without requiring direct calculation of derivatives is developed.
- Improvement over the data-driven CNN model without increased computational overheads
- Physics-informed regularisation extends existing data-driven surrogate modelling.
- Novel methodology is developed to improve functional API over previous PINNs models.

## GRAPHICAL ABSTRACT



## ARTICLE INFO

Editor: Fernando A.L. Pacheco

MSC:  
2000  
0000  
1111

### Keywords:

Physics-informed machine learning  
Surrogate modelling  
PINNs  
Hydrodynamic modelling  
Flood modelling

## ABSTRACT

In response to growing concerns surrounding the relationship between climate change and escalating flood risk, there is an increasing urgency to develop precise and rapid flood prediction models. Although high-resolution flood simulations have made notable advancements, they remain computationally expensive, underscoring the need for efficient machine learning surrogate models. As a result of sparse empirical observation and expensive data collection, there is a growing need for the models to perform effectively in ‘small-data’ contexts, a characteristic typical of many scientific problems. This research combines the latest developments in surrogate modelling and physics-informed machine learning to propose a novel Physics-Informed Neural Network-based surrogate model for hydrodynamic simulators governed by Shallow Water Equations. The proposed method incorporates physics-based prior information into the neural network structure by encoding the conservation of mass into the model without relying on calculating continuous derivatives in the loss function. The method is demonstrated for a high-resolution inland flood simulation model and a large-scale regional tidal model. The proposed method outperforms the existing state-of-the-art data-driven approaches by up to 25 %. This research demonstrates the benefits and robustness of physics-informed approaches in surrogate modelling for flood and hydroclimatic modelling problems.

\* Corresponding author at: School of Engineering, University of Warwick, UK.

E-mail addresses: [jamie.donnelly@warwick.ac.uk](mailto:jamie.donnelly@warwick.ac.uk) (J. Donnelly), [Alireza.Daneshkhah@coventry.ac.uk](mailto:Alireza.Daneshkhah@coventry.ac.uk) (A. Daneshkhah), [Soroush.Abolfathi@warwick.ac.uk](mailto:Soroush.Abolfathi@warwick.ac.uk) (S. Abolfathi).

<https://doi.org/10.1016/j.scitotenv.2023.168814>

Received 14 September 2023; Received in revised form 10 November 2023; Accepted 21 November 2023

Available online 26 November 2023

0048-9697/© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

A rapidly changing climate could rapidly render previous forecasts and flood risk assessments obsolete, necessitating a significant recalibration of existing forecasts. The continuous validation and monitoring of existing models and forecasts in light of new data is crucial for establishing robust predictive models for hydroclimatic problems. To facilitate cost-effective re-evaluation and the production of low-cost forecasts, it is imperative that machine learning-based models, such as surrogate models, are more widely adopted in flood modelling and other environmental science and engineering applications.

The data-driven approach to ML involves relying entirely upon a set of training data to inform decisions about model selection and parameterisations. In this paradigm, selecting the best model, for example determining the number of hidden layers, neurons per layer, and type of activation functions in a neural network, generally relies on heuristics, experience and a trial-and-error approach. After selecting the most appropriate model structure, parameters will be learned by minimising a loss function, possibly subject to some constraints on the parameters, which is often equivalent to minimising the ‘data-fit’ error on the training data. For instance, for the supervised learning tasks, and in particular for regression problems, this usually involves selecting a set of parameters,  $\theta^*$ , that minimise a distance metric between the observed and predicted data. For instance, this could be the mean squared error (MSE) on the training data,

$$\theta^* = \underset{\theta}{\operatorname{argmin}} L(\mathbf{y}, f(\mathbf{x}; \theta)) \quad (1)$$

$$L(\mathbf{y}, f(\mathbf{x}; \theta)) = \mathbb{E}[(\mathbf{y} - f(\mathbf{x}; \theta))^2] \quad (2)$$

where,  $\mathbf{x}$  are the selected inputs,  $\mathbf{y}$  are the corresponding observed/computed outputs, and  $f(\cdot; \theta)$  is the selected model using ML approaches, parameterised by  $\theta$ . Optimisation using the loss function outlined in Eq. (2) is a purely data-driven approach, because it only considers how much the target values predicted by the model differ, under a given parameterisation, from the observed values.

Supervised ML tasks often involve the challenge of learning an unknown distribution,  $P(\mathbf{y}|\mathbf{x})$ , which enable us to draw invaluable inferences about some system or real-world problem. In the case of prediction, the task is to learn a predictive distribution,  $P(\mathbf{y}^*|\mathbf{y}, \mathbf{x}, \mathbf{x}^*)$ , given the training dataset  $\mathcal{D} = (\mathbf{x}, \mathbf{y})$  and a new input,  $\mathbf{x}^*$ . The ML and deep learning literature continues to show, experimentally, that statistical learning-based models can be successfully adopted to a variety of scientific and engineering domains (Lateef et al., 2022; Al-Jamimi et al., 2022). Furthermore, the research also suggests that in the case of deep learning, increasing the scale of models and data yields continuously better estimates of  $P(\mathbf{y}^*|\mathbf{y}, \mathbf{x}, \mathbf{x}^*)$  even in highly complex settings. However, in settings where data is sparse, incorporating prior knowledge about the system,  $\pi(\mathbf{y})$ , is an alternative approach to improving generalisation performance of machine learning models.

When applying ML approaches to address problems in engineering and environmental science, the distribution being learned,  $P(\mathbf{y}|\mathbf{x})$ , is often a description of a real physical system with known laws governing its behaviour. Many of these complex systems have a great deal of prior knowledge available, but it is often ignored during the modelling process because it is challenging to robustly incorporate it into the model structure. As a result, data-driven approaches are often favoured over those that explicitly consider the underlying physics of other domain-specific knowledge.

The problem with applying data-driven ML models to complex problems is that they often require a large number of free parameters in order to accurately capture the underlying complexity of the data. In order to appropriately optimise a large number of parameters, a significant amount of data may be required. Obtaining large amounts of training data is often infeasible due to sparsity and/or the high costs

associated with empirical data collection. Furthermore, purely data-driven approaches are susceptible to extrapolation and generalisation problems. These approaches, trained on a finite set of observations, may struggle to effectively generalise to out-of-sample data, potentially resulting in reduced performance. Additionally, when treating ML models as black boxes, it is often impossible to explain the reasoning and the underlying mechanisms behind a specific prediction.

Physics Informed Machine Learning (PIML; (Karniadakis et al., 2021)) was developed to tackle the shortcomings of data-driven ML methods by incorporating information about the physical laws governing a system in the development and training of the ML models. Specifically, Physics-Informed Neural Networks (PINNs) (Raissi et al., 2019; Cuomo et al., 2022) introduce physics-based regularisation to neural networks. They constrain the parameter space of learnable models by modifying the loss function to include terms that encode knowledge of the underlying physics. One common approach is to incorporate the residuals of partial differential equations (PDEs) into the loss function, aiming to create models that are more physically consistent.

By using loss functions that increase the error for predictions that violate the underlying physics, the training process produces a model which simultaneously minimises the data-fit error and remains consistent with the underlying physics. As long as the additional terms in the loss function are differentiable then gradients can still be propagated through the model and training can take place as normal. This approach to physics-informed machine learning has been applied successfully to a variety of modelling tasks such as forward problems, i.e., inferring latent solutions to systems governed by known equations.

Examples of applications in the literature involve a variety of known PDE systems such as Schrodinger’s equations (Raissi et al., 2019), heat transfer problems (Cai et al., 2021), the Euler equations for aerodynamic flows (Mao et al., 2020), the Eikonal equation for cardiac function (Sahli Costabal et al., 2020) and a physics-guided NNs for lake temperature modelling (Karpatne et al., 2017). There are examples in the hydrodynamic modelling literature for tasks such as predicting fluid flow through porous media (Almajid and Abu-Al-Saud, 2022), flood modelling (Liu et al., 2022), forecasting water levels (Qian et al., 2019; Donnelly et al., 2023; Fanous et al., 2023a) and river flows (Feng et al., 2023), and visualisation tasks like visualising spatial inundation (Lütjens et al., 2020; Lütjens et al., 2021). Applications of PINNs and PIML can be seen in atmospheric and climate modelling applications as well (Kashinath et al., 2021; Yang et al., 2023). PINNs methods have also been demonstrated for inverse problems, using observed data to make inferences about the underlying physics (Jagtap et al., 2022; Depina et al., 2022; Chen et al., 2020).

Many previous PINNs works rely heavily on data from computational fluid dynamics (CFD) simulations. CFD modelling has seen growing adoption of ML methods in recent years (Kutz, 2017; Brunton et al., 2020; Kochkov et al., 2021). However, many of the examples involving PINNs models are highly idealised toy problems and may not be representative of real-world problems involving fluid dynamics. Numerical CFD-based modelling and simulation of flooding scenarios is a vital part of flood risk management and climate resilience enhancement schemes (Donnelly et al., 2022). Flood modelling relies on accurate and timely forecasting of complex flow-structure interactions which can be used to guide policy, infrastructure planning, and decisions. However, due to the associated high computational costs of the underlying CFD-based models in flood modelling and other applications, ML-based methods have been increasingly adopted (Mosavi et al., 2018; Yang and Chang, 2020; Lin et al., 2020; Kabir et al., 2020).

Surrogate models (Kabir et al., 2020; O’Hagan, 2006; Sudret et al., 2017; Conti et al., 2009) have been widely adopted as cost-effective alternatives to expensive numerical simulators. Surrogate models involve training an ML model to approximate the input-output patterns of another, usually expensive to evaluate, model. By utilising rapid ML models, it becomes feasible to make nearly instantaneous predictions of the outputs of a highly complex numerical model. A variety of ML

methods have been used in the context of surrogate modelling, such as neural networks (Garca-Alba et al., 2019; Wang et al., 2019; Shi et al., 2018), Polynomial Chaos Expansions (Massoud, 2019; Moreno-Rodenas et al., 2018), boosting trees (Yan et al., 2023) and Gaussian Processes (Conti et al., 2009; Chang et al., 2015; Longobardi et al., 2020; Yang et al., 2018). Existing approaches to surrogate modelling almost entirely rely on data-driven methodologies, minimising a distance metric between the surrogate's predictions and the numerical model's true outputs. This approach makes no consideration for the governing behaviour of the numerical model and in turn the real physical system the numerical model approximates.

There are existing issues surrounding the use of purely data-driven approaches to surrogate modelling such as poor data efficiency. Some machine learning models require extensive data for reliable out-of-sample predictive performance. However, in many real-world scenarios, data is either sparse or very prohibitively expensive to collect, necessitating models that can generalise effectively on smaller datasets. Data-driven models are generally treated as black-boxes, where the rationale for prediction is unknown, and their explainability is typically low. This lack of transparency introduces a significant level of uncertainty regarding the basis for certain predictions. Introducing stricter regularisation into ML models through robust incorporation of prior information can address the issues related to data efficiency. This would allow models to be trained effectively on smaller datasets while enhancing their explainability, as it becomes possible to study the behaviour of additional variables into the models. These developments would be highly beneficial for surrogate modelling, offering a robust hybrid approach to forecasting that combines the strengths of ML and numerical modelling.

This study extends the application of PINNs to complex realistic problems in fluid dynamics by constructing a PINNs surrogate model for numerical hydrodynamic simulators in the context of flood modelling. The proposed PINNs surrogate is developed for fast inference of latent solutions of fluid depth (i.e., surface water elevations) at high spatio-temporal resolutions, as generated by numerical simulators solving the two-dimensional Navier-Stokes equations. Furthermore, instead of directly incorporating PDEs' residuals into the loss function, a novel approach to physics-based regularisation approach has been developed to account for the conservation of mass in predictions made by the surrogate model.

The proposed approach avoids relying on auto-differentiation to estimate partial derivatives, thus reducing computational complexity and eliminating the implicit reliance on realistic gradients throughout the model. In surrogate modelling tasks, boundary conditions that are known a priori are effectively utilised by the proposed model, replacing implicit computation (such as obtaining auto-diff gradients) with empirical discrete comparisons to quantify violations from true physics. The efficacy of the proposed approach is demonstrated by constructing surrogate PINNs models for two widely used hydrodynamic simulators, LISFLOOD-FP (LFP) and Delft3D, and the performance is assessed against a data-driven approach to surrogate modelling.

## 2. Methods

### 2.1. Shallow water equations

Numerical hydrodynamic simulators implement discretised versions of the Navier-Stokes equations. The Navier-Stokes equations are a set of PDEs that describe the motion of fluid substances in three-dimensional space. The Shallow Water Equations (SWE), on the other hand, are a simplified form of the Navier-Stokes equations that describe the motion of fluids that are shallow in depth (e.g. such as those found in rivers, lakes and estuaries) and homogeneous along the depth-axis. Therefore, two-dimensional hydrodynamic simulators usually solve the SWE in two-dimensions rather than the full Navier-Stokes equations. Eqs. (3)–(4) describe the SWE governing equations:

$$\nabla \cdot \mathbf{u} = 0, \quad (3)$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} - \nu \nabla^2 \mathbf{u} = -\nabla \left( \frac{p}{\rho_0} \right) + \mathbf{g}, \quad (4)$$

where Eqs. (3)–(4) are the continuity equation, expressing that mass is conserved, and the momentum equation, expressing that momentum is conserved, respectively.

The PINNs model developed in this study is applied to two different hydrodynamic simulators, LISFLOOD-FP and Delft3D for fluvial flooding and coastal modelling applications, respectively. However, the approach outlined in this study is not limited to these models or applications and could be universally applied to any hydrodynamic simulator partly governed by the continuity equation. In order to derive the discrete aggregated physics-informed loss function, a detailed description of a two-dimensional hydrodynamic simulator is outlined in Section 2.2.

### 2.2. Two-dimensional hydrodynamic simulator

LISFLOOD-FP and Delft3D are independent simulators with different numerical implementation details, however, both models implement the Navier-Stokes equations and so conceptually they operate similarly. Therefore, a conceptual overview of a two-dimensional hydrodynamic simulator is outlined in order to introduce the PINNs model more efficiently.

It is assumed that the hydrodynamic simulators operate on a two-dimensional Eulerian mesh describing the topography, with  $n_x$  cells in the  $x$ -direction and  $n_y$  cells in the  $y$ -direction, with a cell size of  $\Delta x$  and  $\Delta y$  [m]. The model has initial conditions  $y_{ij}^0$ , where  $y_{ij}^t$  [m] denotes the fluid depth in the  $(i, j)$ -th cell in the domain at time  $t$ . There are time and space-varying boundary conditions,  $s_{ij}^t$  [ $m^3/s$ ], which describe a rate of fluid inflow between times  $t$  and  $t + 1$ , where the number of seconds between these timesteps is  $\Delta t$ . Fluid also leaves the domain with a rate  $q_{ij}^t$  [ $m^3/s$ ], through sinks within the domain or at outlets at boundaries. Given the initial conditions,  $\{y_{ij}^0\}$  and collection of boundary conditions  $\{s_{ij}^t\}_{t=1}^T$ , the numerical model then computes the fluid interactions across the domain according to the governing equations, Eqs. (3)–(4).

### 2.3. PINNs surrogate model

When utilising a hydrodynamic simulator, only the boundary conditions and outputs (e.g., fluid surface elevations, velocity fields, pressure) are known post-simulation, such that direct access to the derivatives is usually not possible. Previous PINNs implementations overcome this problem by regressing outputs,  $u(t, x)$ , onto their spatio-temporal coordinates such that  $\hat{u}(t, x) = f(t, x)$ , where  $f(\cdot, \cdot)$  is a neural network. Then by backpropagating to the original inputs, it's possible to calculate  $\frac{\partial \hat{u}}{\partial x}$  or  $\frac{\partial \hat{u}}{\partial t}$  which can then be incorporated into the loss function via PDE residuals. For example, if inferring latent solutions to the one-dimensional linearised wave equation, a simplification of the Euler equations, the governing PDE can be expressed as,

$$\frac{\partial^2 u}{\partial t^2} - c^2 \nabla^2 u = 0. \quad (5)$$

By using a neural network to estimate  $\hat{u}(x, t)$  and then backpropagating to obtain  $\frac{\partial \hat{u}}{\partial t}$ ,  $\frac{\partial \hat{u}}{\partial x}$  and repeating the process  $n$  times to obtain  $n^{\text{th}}$ -order derivatives, an additional term can be included into the loss function,

$$L(u, \hat{u}) = (u - \hat{u})^2 + \left| \frac{\partial^2 \hat{u}}{\partial t^2} - c^2 \nabla^2 \hat{u} \right|^2, \quad (6)$$

where the first term is the data-fit error term and the second is the PDE residual regularisation term. An illustration of this existing type of PINNs model and network architecture can be seen in Fig. 1.

The API, i.e., how the model can be interacted with and queried, of a standard PINNs model like this is limited in that at test-time it can only robustly estimate outputs using the same boundary conditions it was trained on. It does not allow for flexibility in the parameterisation of boundary conditions during inference. (Raissi et al., 2020) states that a PINNs model as outlined is agnostic to boundary conditions however, in general, for new boundary conditions, a new model or additional training data to update an existing model is likely required for good generalisation. The model proposed in this study overcomes this limitation by explicitly parameterising the boundary conditions in the neural network's input, allowing straightforward inference of latent solutions under widely varying boundary conditions and without the requirement to backpropagate to obtain derivatives.

In this study, a novel PINNs-based surrogate model of the following form is proposed,

$$\hat{Y}^t = f(s^t, \dots, s^{t-N}), \quad (7)$$

where  $\hat{Y}^t$  is a matrix representing predicted fluid depth in all cells (i.e.,  $n_x \times n_y$ ) in the domain, and  $s^t$  are the boundary conditions at time  $t$ , and  $f$  is a neural network. Based on the surrogate model presented in Eq. (7), the water surface elevations at time  $t$  are regressed onto the boundary conditions at time  $t$  and the  $N$  previous timesteps. The model regresses outputs onto a small segment of the time series ( $N + 1$  values) in order to understand the behaviour that arises from the temporal differential in the boundary condition. For instance, without encoding that historical behaviour in the model through the time series input, the model might not be able to discriminate between two different scenarios where the boundary condition value at times  $t_{N_1}$  and  $t_{N_2}$  are the same but one value corresponds to a downward trend and the other an upward trend with respect to time. A model of this form is a flexible surrogate model and can act as a replacement to the original numerical model.

For hydrodynamic simulators the continuity equation, Eq. (3), can be informally stated as,

$$\Delta(\text{storage}) = \Delta(\text{sources} - \text{sinks}), \quad (8)$$

which forms the basis of the physics-informed loss function for the PINNs model proposed in this study. Using the conceptual hydrodynamic simulator outlined in Section 2.2, this can be re-expressed for

fluid depths in cell- $(i, j)$  predicted by the PINNs surrogate,  $\hat{y}_{ij}^t$ , and true fluid depths  $y_{ij}^t$ , as,

$$\Delta x \Delta y \sum_i \sum_j \hat{y}_{ij}^t - \Delta x \Delta y \sum_i \sum_j y_{ij}^{t-1} = \Delta t \sum_i \sum_j s_{ij}^t - \Delta t \sum_i \sum_j q_{ij}^t, \quad (9)$$

where  $s_{ij}^t$  and  $q_{ij}^t$  is the inflow and outflow in cell- $(i, j)$  and time  $t$ , respectively. This shows that the aggregate volume of fluid in the domain predicted by the surrogate at time  $t$ , minus the true volume at time  $t - 1$  should equal the aggregate inflow minus the aggregate outflow. This equation can then be re-expressed as,

$$\Delta x \Delta y \sum_i \sum_j \hat{y}_{ij}^t + \Delta t \sum_i \sum_j q_{ij}^t = \Delta x \Delta y \sum_i \sum_j y_{ij}^{t-1} + \Delta t \sum_i \sum_j s_{ij}^t. \quad (10)$$

In practice, the outflows,  $q_{ij}^t$ , are generally latent variables (in the absence of defined per-cell absorption/drainage rates) calculated during simulation, representing the aggregate discharge leaving the domain through open boundaries, but it is always true that  $q_{ij}^t \geq 0$ , therefore, Eq. (10) can be turned into the following inequality,

$$\Delta x \Delta y \sum_i \sum_j \hat{y}_{ij}^t \leq \Delta x \Delta y \sum_i \sum_j y_{ij}^{t-1} + \Delta t \sum_i \sum_j s_{ij}^t. \quad (11)$$

This is equivalent to putting a constraint on the aggregate predicted volume at time  $t$  such that it cannot exceed the storage at  $t - 1$  and the net volumetric inflow between  $t - 1$  and  $t$ .

However, Eq. (11) only satisfies half the continuity equation, stating that mass cannot be created. In order to fully satisfy the continuity equation a second inequality is introduced as follows:

$$\Delta x \Delta y \sum_i \sum_j \hat{y}_{ij}^t \geq \Delta x \Delta y \sum_i \sum_j y_{ij}^{t+1} - \Delta t \sum_i \sum_j s_{ij}^{t+1}, \quad (12)$$

with the full continuity equation being satisfied by the following inequality,

$$\Delta x \Delta y \sum_i \sum_j y_{ij}^{t+1} - \Delta t \sum_i \sum_j s_{ij}^{t+1} \leq \Delta x \Delta y \sum_i \sum_j \hat{y}_{ij}^t \leq \Delta x \Delta y \sum_i \sum_j y_{ij}^{t-1} + \Delta t \sum_i \sum_j s_{ij}^t. \quad (13)$$

Combining the PINNs model outlined in Eq. (7) with the continuity inequality given in Eq. (13), the physics-informed loss function will be defined as follows:

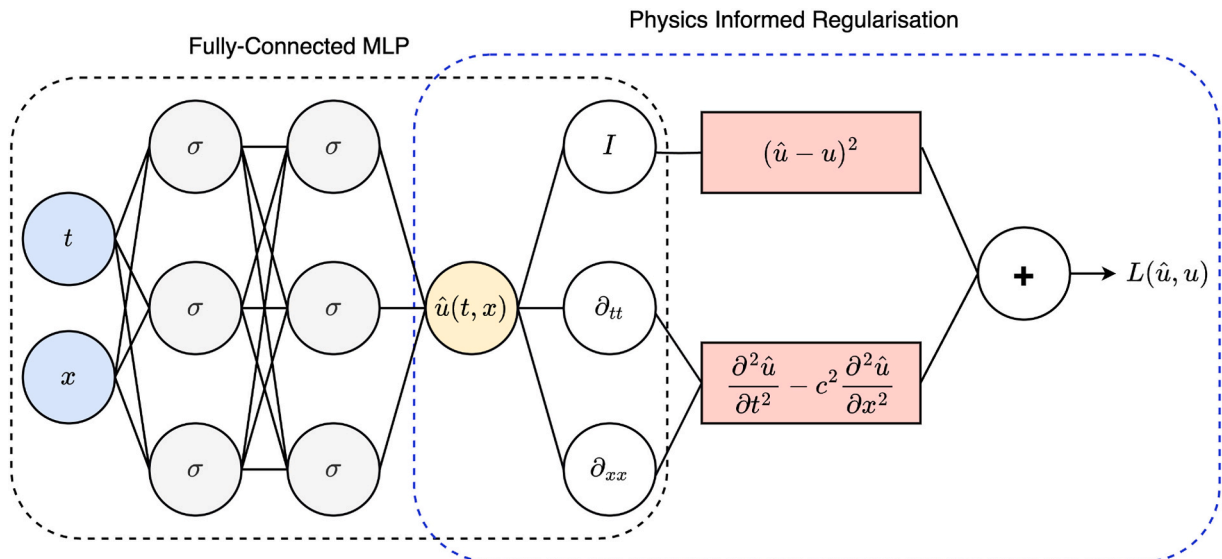


Fig. 1. Example of PINNs architecture for the one-dimensional linearised wave equation, regressing latent solutions  $u$  onto spatiotemporal coordinates,  $(t, x)$ .



$$L(Y^t, \hat{Y}^t) = \mathbb{E} \left[ (Y^t - \hat{Y}^t)^2 \right] + \text{ReLu} \left( \hat{v}^t - v^{t-1} - \Delta t \sum_i \sum_j s_{ij}^t \right) + \text{ReLu} \left( v^{t+1} - \hat{v}^t - \Delta t \sum_i \sum_j s_{ij}^{t+1} \right), \quad (14)$$

where  $v^t = \Delta x \Delta y \sum_i \sum_j y_{ij}^t$  is the true volume in the domain at time  $t$  and  $\hat{v}^t = \Delta x \Delta y \sum_i \sum_j \hat{y}_{ij}^t$  is the predicted volume, and  $\text{ReLu}$  is the Rectified Linear Unit function.  $\text{ReLu}$  functions operate as follows:

$$\text{ReLu}(x) := \max(x, 0). \quad (15)$$

The above equation highlights that if the inequalities expressed in Eq. (13) are violated, additional values are included in the loss function. The first term in Eq. (14) is the Mean Squared Error (MSE) between the true and predicted outputs, referred to as the data-fit term, and the last two components are physics-based terms responsible for penalising predictions that violate the conservation of mass, which are the equivalent representation of the double inequality expressed in Eq. (13).

In order to ensure that Eq. (14) is an effective loss function, its three components need to be appropriately prioritised during the training of the network. The MSE loss term will be represented in  $[m^2]$  whereas the second and third terms will be represented in  $[m^3]$ . Furthermore, the magnitude of these values may vary significantly as volume misprediction could easily exceed  $10^4 m^3$  but a reasonable fluid depth prediction could be expected to have an MSE less than  $10^{-1} m^2$ , meaning the overall loss would be dominated by the second and third terms rather than the data-fit error, leading to undesirable optimisation behaviour. To solve this problem, the second and third terms will be scaled to match the units of the data-fit error. As a result, the final physics-informed loss function with the scaled regularisation terms will be given by,

$$L(Y^t, \hat{Y}^t) = \mathbb{E} \left[ (Y^t - \hat{Y}^t)^2 \right] + \left( \frac{\text{ReLu}(\hat{v}^t - v^{t-1} - \Delta t \sum_i \sum_j s_{ij}^t)}{\Delta x \Delta y \times n_x \times n_y} \right)^2 + \left( \frac{\text{ReLu}(v^{t+1} - \hat{v}^t - \Delta t \sum_i \sum_j s_{ij}^{t+1})}{\Delta x \Delta y \times n_x \times n_y} \right)^2, \quad (16)$$

where the denominator,  $\Delta x \Delta y \times n_x \times n_y$ , is the total area of the modelling domain.

#### 2.4. Neural network architecture

Previous PINNs studies (Cai et al., 2021; Fanous et al., 2023b; Mao et al., 2020; Raissi et al., 2017; Zhang et al., 2019) utilised a dense fully-connected MLP (multi-layer perceptron) architecture as the basis of the model, however, in some of these studies no consideration for alternative architectures was made. In this study, a convolutional neural network (CNN) model is used for the architecture. Preliminary testing in the authors' previous work (Donnelly et al., 2022) found that a CNN model was a more appropriate architecture compared to an MLP. Furthermore, prior to this study, Gaussian Processes (Donnelly et al., 2022; Donnelly et al., 2024) and CNNs (Kabir et al., 2020; Garca-Alba et al., 2019; Bates, 2022) were shown to be the most effective data-driven surrogate model for approximating hydrodynamic simulators. However, physics-informed approaches are currently better suited to the flexibility of neural network-based models, hence why a CNN was adopted rather than GP. The proposed CNN architecture, where convolutions occur in one-dimension over a vector-based input more effectively learns the patterns arising from the input data with a temporal component.

Furthermore, a CNN introduces an inherent structure to the data, learning patterns through local connections, and having highly structure spatial data arising from Eulerian meshes, CNNs parallel this structure

from the governing processes. MLPs do not enforce this same structure and did not learn the patterns as effectively. Moreover, a data-driven CNN was chosen as an appropriate benchmark for this study because of its successes in previous implementations (Kabir et al., 2020). However, it is important to note that the physics-informed loss function proposed in this study is entirely agnostic to the selected neural network architectures and can be readily incorporated into any differentiable neural network model. In Fig. 2, an illustration of the general architecture used for the CNN in this study is provided.

To ensure that any differences in performance between the PINNs and data-driven CNN models are because of the physics-informed regularisation, the respective models were trained and tested by conditioning on the initialisation the network weights. By randomly sampling different weight initialisations,  $\theta_i^0$ , and then training the PINNs and CNN models with their respective loss functions to estimate optimal parameters (as outlined in Eq. (1));  $\theta_i^*$ , the models' test performance was averaged over the subsequent networks to mitigate effects of favourable or unfavourable initial weights. The test performance of the models then becomes:

$$L(\mathbf{y}^*, f(\mathbf{x}^*)) = \mathbb{E}_i [L(\mathbf{y}^*, f(\mathbf{x}^*; \theta_i^*)) | \theta_i^0], \quad (17)$$

where  $(\mathbf{x}^*, \mathbf{y}^*)$  denote the test inputs and outputs.

The physics-informed loss function is likely highly non-convex with respect to the parameters and therefore predictive performance could be sensitive to the initialisation of the weights of the network. Taking expectations with respect to weight initialisations allows for a high level of certainty that any deviations in performance against the CNN model are entirely the result of the model and not due to favourable/unfavourable initial conditions in the weight-space.

#### 2.5. Model evaluation metrics

To assess model predictive performance two metrics are utilised in this study. The first metric used is the Root Mean Squared Error (RMSE) of the true fluid depths,  $Y^t$ , and the fluid depths predicted from the surrogate  $\hat{Y}^t$ ,

$$\text{RMSE} = \sqrt{\mathbb{E} \left[ (Y^t - \hat{Y}^t)^2 \right]}. \quad (18)$$

RMSE is chosen as the primary evaluation metric in this study as it scales the prediction errors back to the original unit of measurement,  $[m]$ , and so gives highly interpretable results. Secondly, a modified version of the RMSE is used,

$$\text{mRMSE} = \sqrt{\frac{\sum_{\Omega} (y_{\Omega}^t - \hat{y}_{\Omega}^t)^2}{|\Omega|}}, \quad (19)$$

where  $\Omega = \{(i, j) | y_{ij}^t \cup \hat{y}_{ij}^t > 0\}$ . The mRMSE metric gives a more realistic estimate of the true error by only considering cells where fluid is either predicted or observed at any time in the dataset, and so does not include cells that remain dry at all times which can misleadingly reduce the mean error when included. This modified metric, especially in the case of the inland flood model where the majority of cells remain dry, is a better measure of how much of the true signal is being captured by the model.

It should be also noted that the evaluation metrics reported in Sections 3.1.2 and 3.2.2, are based on cross-validation metrics. In cross-validation, the full dataset is iteratively partitioned into new training and test sets, and the models' performance is assessed based on aggregate performance across these partitions. In ML applications, this is most robust approach to evaluating predictive performance of a model, ultimately improving the reliability and validity of results.

Physics Informed Regularisation

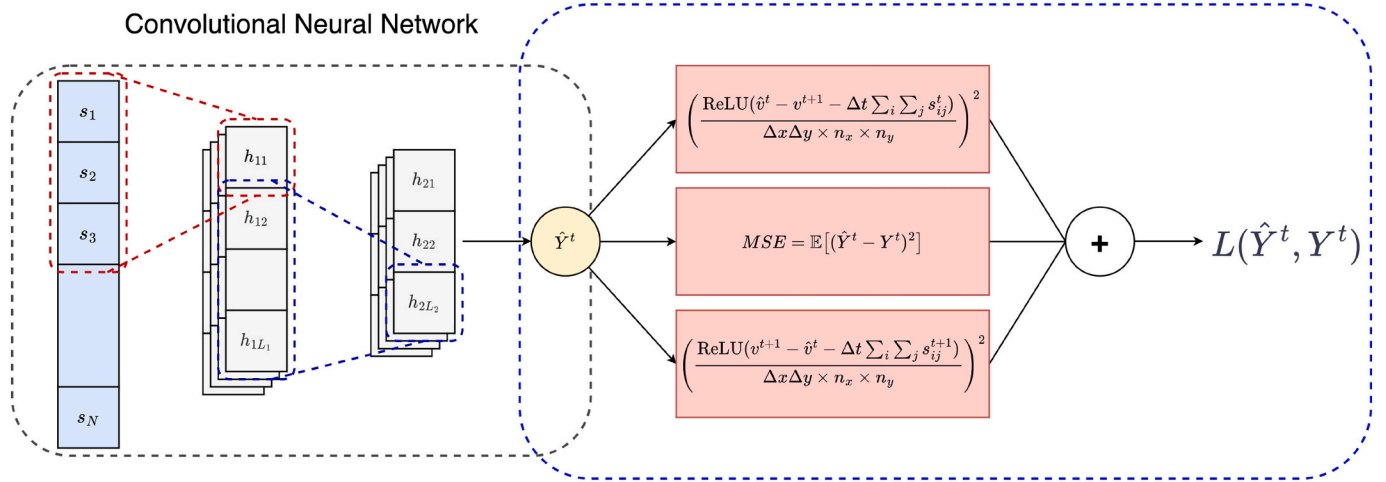


Fig. 2. PINNs with CNN architecture used in this study. The terms in the loss function are the discretised components of the full inequality satisfying the continuity equation, outlined in Eqs. (13)–(16).

3. Experiments

3.1. Inland flood modelling

3.1.1. LISFLOOD-FP

The first two-dimensional hydrodynamic simulator used is LISFLOOD-FP (LFP; (Bates and De Roo, 2000; Neal et al., 2011)). LFP is a raster-based hydraulic inundation model solving various forms of the SWE (using different in-built solvers) and is capable of accurately simulating flow in a channel and over a flood plain at high-resolution in complex topographies. LFP has been validated by many studies with applications including urban inundation, fluvial flooding and to applications in modelling coastal processes (Shustikova et al., 2019; O’Loughlin et al., 2020; Vousdoukas et al., 2016a; Skinner et al., 2015). This study uses the results of an LFP model outlined in previous work by the authors in which a high-resolution urban inundation model was developed for Tadcaster, UK. The model outlined was designed to simulate inundation resulting from river overflow, with the results used to construct a Gaussian Process (GP)-based surrogate models. For

additional details about the model setup the reader should refer to the authors’ previous study (Donnelly et al., 2022) and for additional details on the model’s governing equations and numerical schemes, see (Bates and De Roo, 2000; Neal et al., 2011).

The modelling domain for the proposed inland flood model can be seen in Fig. 3, illustrating a map and satellite image of the study area. On this domain, the simulator was run for 14 synthetic flood events with the boundary conditions for each simulation,  $\{s^1\}_{t=1}^T$ , describing river discharge  $[m^3/s]$  at the upstream boundary, with a focus on extreme climatic events. The proposed inland flood model relied on the generation of synthetic boundary conditions, which would correspond to low-probability, high-impact, short-term flood events in the case study location of Tadcaster. Using synthetic input data (details can be found in (Donnelly et al., 2022)) in this manner means that there is no corresponding observational data available to validate the calibrated model. Therefore, the estimates and data produced from the model cannot be interpreted as real-life event forecasts. However, it is worth noting that the numerical simulator adopted in this study is widely validated and



Fig. 3. Tadcaster LISFLOOD-FP Modelling domain from the authors’ previous work (Donnelly et al., 2022).

without suitable validation data (measured empirical data), most numerical models of this kind face similar limitations.

The model outputs are a collection of time-indexed matrices,  $Y^t \in \mathbb{R}^{n_x \times n_y}$ , describing inundation in each cell of the rectangular domain for timesteps  $t = 1, \dots, T$ . The original LFP model can then be described as a complex mapping,

$$Y^t = f(s^t), \tag{20}$$

where,  $Y^t \in \mathbb{R}^{213 \times 256}$ . The surrogate model however takes the form,

$$\hat{Y}^t = \hat{f}(x^t), \tag{21}$$

where,  $x^t = (s^t, \dots, s^{t-8})$ , meaning the inputs to the surrogate to predict the output at time  $t$  are the boundary conditions at time  $t$  and the 8 preceding timesteps. Similarly, this decision was motivated by the authors' previous work (Donnelly et al., 2022) where for a GP-based surrogate model was found to be an appropriate amount of historical information to encode into the model structure. By aggregating data across 14 discrete simulations into a collection of independent tests, and flattening the matrices of fluid depth,  $Y^t$ , into vectors,  $y^t$ , resulted in a dataset of  $N = 2624$  time-indexed samples. Each time-indexed input is represented by  $x^t \in \mathbb{R}^9$  and each output by  $y^t \in \mathbb{R}^{54528}$ . The boundary conditions for the 14 simulations are illustrated, in Fig. 4, these are a collection of synthetic hydrographs derived from empirical observations. For further detail on the generation of these boundary conditions and the LFP model development, refer to (Donnelly et al., 2022).

### 3.1.2. Results

In order to evaluate the PINNs surrogate model performance, leave-one-out cross-validation (CV) was performed by partitioning the data into subsets, holding one partition out for testing and then training the

model on the remaining partitions. This resulted in 14-fold CV, where the data were partitioned corresponding to the original simulations.

Preliminary testing was done to determine appropriate hyper-parameters for each model, such as learning rates and early stopping criterion. For both networks a learning rate of  $\alpha = 0.001$  was selected, and both models were set to stop training either after 50 epochs, with a batch size of 16, or if the validation RMSE had converged within a range of 0.01 for 3 consecutive epochs. Fig. 5 outlines the training and convergence of each model during CV, with the average converge path highlighted along with the variation among partitions. The CNN model converges faster but also varies more around the validation minima. The PINNs model however takes more weight-updates to reach convergence but appears to reach a more stable minimum that is also very slightly lower, as outlined in Table 1.

The results of the 14-fold CV for both the CNN and PINNs models, illustrated in Table 1, show that when assessing the models by ordinary RMSE, as a naive assessment of the true error (Donnelly et al., 2022), the PINNs model performs very similarly with the CNN obtaining only a 3.5 % improvement, a result which could likely be noise and it could be concluded that the models are equivalent with respect to this metric. However, when assessing the models by the modified RMSE metric (i.e., mRMSE), as a better reflection of the true error rate, the PINNs model significantly outperforms the CNN model by just over 25 %. This result is less likely to be due to the noise in the model and data, and indicates that the PINNs model can provide superior predictive performance compared to the CNN model with respect to this metric. Additionally, the third column in Table 1 includes the average volume discrepancy for each model's prediction where this volume discrepancy is outlined in Eq. (22), where  $\mathcal{T}$  denotes the test-set time indices. This result shows that the PINNs model has a slightly lower volume residual suggesting that it is more physically consistent with the conservation of mass.




Fig. 4. Boundary conditions for the 14 LISFLOOD-FP simulations from the authors previous work (Donnelly et al., 2022).



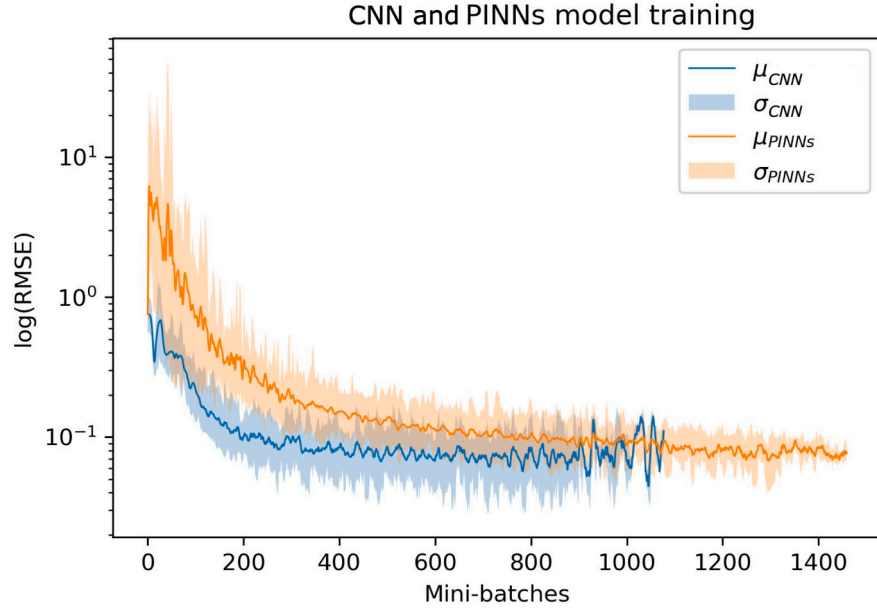


Fig. 5. Training of CNN and PINNs models with the y-axis showing  $\log(\text{RMSE})$  on the test set and x-axis showing the number of mini-batches/weight updates. The  $\sigma$ -bounds shows the variability of convergence across CV partitions.

**Table 1**  
Mean CV scores for the CNN and PINNs models on the LISFLOOD-FP data.

Model	mRMSE	RMSE	Vol. Discrepancy
CNN (MSE)	0.184	0.083	$9.74 \times 10^5$
PINNs	0.136	0.086	$8.70 \times 10^5$

$$\mathbb{E}_{t \in \mathcal{T}} \left[ \text{ReLU} \left( \hat{v}^t - v^{t-1} - \Delta t \sum_i \sum_j s_{ij}^t \right) + \text{ReLU} \left( v^{t+1} - \hat{v}^t - \Delta t \sum_i \sum_j s_{ij}^{t+1} \right) \right]. \quad (22)$$

A full breakdown of the CV results by partition can be seen in Fig. 6. The results for the normal RMSE metric are similar however when observing results for the modified RMSE value it is clear that the PINNs model confers a clear benefit. The CNN still outperforms the PINNs model on a subset of partitions however, when the PINNs model outperforms the CNN, the expected reduction in the test error is approximately 36 %. Whereas the expected reduction in test error in cases where the CNN model outperforms the PINNs is only 18 %. A thorough review of these results demonstrates the improved performance of the PINNs model with respect the chosen error metrics.

A spatial illustration of typical predictive performance of the CNN and PINNs models are outlined in Figs. 7 and 8, showing the  $L_1$  loss for random samples in the test-set for both models, respectively. These

results are indicative of the average performance of each model, where the PINNs model observes a pattern of lower less values generally, especially in the channel and those areas where inundation is higher (e.g., the river channel and surrounding urban areas). Overall the benchmark CNN is still a good predictor, however, in the areas of higher inundation it does not perform as well. These results are illustrated by observing the  $L_1$  loss values of each model's prediction in the river channel and surrounding areas. This result links back to Table 1, which showed that the PINNs model observes a more significant improvement when utilising the less biased mRMSE metric. Additionally,

Fig. 9 outlines aggregated spatial variability in the performance of the PINNs and CNN baseline model by showing the cell-by-cell differences in performance, expressed as  $\text{RMSE}_{\text{CNN}} - \text{RMSE}_{\text{PINNs}}$ , meaning that positive values indicate areas where on average the CNN performed worse (had a higher RMSE) and similarly, negative values indicate the PINNs model performed worse. We can observe that across the majority of the case study domain, the PINNs model is outperforming the CNN model. This is particularly evident within the channel of the river, where in certain areas the average improvement in RMSE is  $\geq 0.01\text{m}$ . These results suggest that the PINNs model is effective at learning a more accurate distribution of fluid volume than the CNN model as a result of the additional physics-based regularisation. By enforcing the conservation laws into the model structure a more accurate volume of total fluid is predicted.






Fig. 6. LFP Model results by CV partition. The plot on the left shows assessment by ordinary RMSE and the right plot by the modified RMSE metric.





(a)  $L_1$  for CNN trained using MSE



(b)  $L_1$  for PINNs

Fig. 7. Illustration of  $L_1$  loss for CNN and PINNs for a random timestep (i).

### 3.2. Coastal flood modelling


#### 3.2.1. Delft3D

Delft3D (Roelvink and Van Banning, 1995; Deltares, 2022) is a modelling suite consisting of several modules, able to accurately simulate hydrodynamics, waves, morphology, and sediment transport among other processes. In this study, the FLOW hydrodynamic module (referred to as Delft3D) of the Delft3D suite is used. Delft3D is a comprehensive standalone model, which is able to robustly simulate storm surges in coastal regions (Vatvani et al., 2012; Vousdoukas et al., 2016b), river flows and flooding (Kumbier et al., 2018), tidal dynamics (Horstman et al., 2013; Waldman et al., 2017), tsunamis, and particle transport. Delft3D is designed to simulate non-steady flows resulting from tidal or meteorological forcing (Deltares, 2022). It can operate in both two and three dimensions, utilising a boundary-fitted grid for accurate


calculations. Delft3D-FLOW uses the Finite Volume Method (FVM) to implement the two or three-dimensional SWE equations. In doing so, Delft3D is able to simulate the hydrodynamics in a wide range of scenarios in which the horizontal length and time scales are much larger than the vertical depth of the model (e.g., shallow seas, coastal areas, estuaries, lagoons, rivers, and lakes) (Deltares, 2022). By default, Delft3D implements a  $k - \epsilon$  turbulence closure model.

The two-dimensional depth-averaged implementation of the SWE is used since the fluid is considered vertically homogeneous and the magnitude of the spatial scales in the  $x$  and  $y$  directions are much greater than the  $z$  direction. The modelling domain chosen in this study is the English Channel, as illustrated in Fig. 10. The modelling domain consists of a 2D grid of  $(287 \times 330)$  cells of size  $1.1\text{km} \times 1.4\text{km}$  with a uniform Manning's roughness coefficient of 0.025 chosen for the domain during calibration. The bathymetry data illustrated in Fig. 10 and used in the





(a)  $L_1$  for CNN trained using MSE



(b)  $L_1$  for PINNs

Fig. 8. Illustration of  $L_1$  loss for CNN and PINNs for a random timestep (ii).

hydrodynamic model is sourced from the GEBCO global bathymetry dataset provided by the British Oceanographic Data Centre (BODC). The developed model was simulated for four discrete 1-month time periods, with the model using time-varying water levels applied at the boundaries. The proposed simulation periods can be seen in Table 2, with simulation periods spread throughout the year to capture some of the temporal variability in tidal dynamics in the dataset.

Likewise to Section 3.1.1, the simulation data is then aggregated across the four simulations, collected at 30-min intervals, to create a full dataset of approximately  $N = 5700$  tests. The tidal model was calibrated and subsequently validated, using empirical data collected from the tide gauge locations highlighted in Fig. 10. This was done to ensure that the underlying results that the PINNs surrogate is trained upon were valid.

The focus of the Delft3D simulator, and subsequent PINNs surrogate model, is to accurately model water surface elevation in the English

Channel modelling domain. As can be seen in Fig. 10, the modelling domain exhibits significant spatial heterogeneity in elevation, with certain areas of the Channel reaching depths of 150 m while other extensive sections are 20 m deep or shallower. To ensure robustness, a surrogate model should accurately and efficiently replicate the spatio-temporal variability in the water surface elevations in the modelling domain, representing a complex dynamic regression problem.

To illustrate the significant variability in water surface elevations across the domain, Fig. 11 A displays time series showing the water surface elevations at ten independent boundary segments for the first simulation. As expected, the structure is highly similar due to the cyclical nature of tidal dynamics. However, due to the variation in elevation, the mean values are widely different. To improve the neural network training, all the data, including input boundary conditions and outputs (water surface elevation matrices), were centered to have zero





Fig. 9. Illustration of the spatial difference in performance (RMSE) between the CNN model and the PINNs model. Positive values indicate areas where the CNN loss is higher and negative where the PINNs loss is higher.

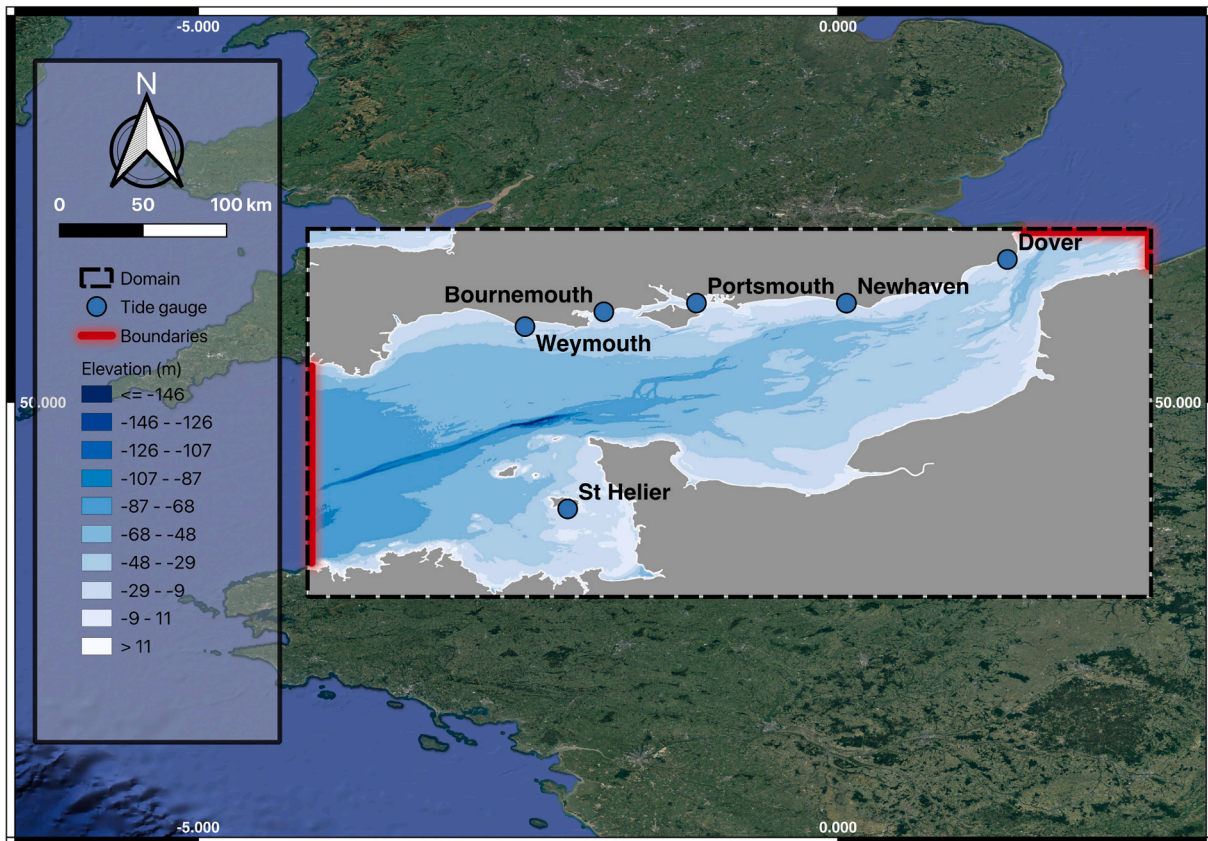


Fig. 10. Delft3D modelling domain of the English Channel with bathymetry, the boundaries at which water levels are applied from, and tide gauge validation locations illustrated.

mean, with the impact on surface elevations at the boundaries illustrated in Fig. 11 B. During the cross-validation process, the training and test sets were centered by subtracting the means of the input and output data

from the training set. This was done to ensure that the test-set mean was approximated using the training-set mean, while also avoiding any sharing of data between the two sets. This method helped to improve the

**Table 2**  
Simulation periods for the Delft3D tidal model.

Simulation	Start	End
1	01-03-2005	01-04-2005
2	01-05-2006	01-06-2006
3	01-07-2007	01-08-2007
4	01-09-2008	01-10-2008

accuracy and reliability of the results obtained from the neural network model. An important clarification is that centering the data in this manner is a linear transformation and will not affect the inequality outlined in Eq. (13).

The proposed tidal model has 40 distinct boundaries, each varying at every timestep. The PINNs surrogate still takes the form outlined in Eq. (21), where boundary conditions are aggregated, along with the 8 preceding timesteps to form inputs to the surrogate of the form  $\mathbf{x}^t \in \mathbb{R}^{40 \times 9} = \mathbb{R}^{360}$ . The outputs are still represented by vectors (matrices representing the modelling domain flattened into vectors),  $\mathbf{y}^t \in \mathbb{R}^{287 \times 330} = \mathbb{R}^{94380}$ . To reflect the change in boundary parameterisation, the CNN for the LFP surrogate had a kernel stride of 1, whereas the Delft3D CNN will have a kernel stride of 9 to reflect the change in boundaries but maintains the same temporal basis to the parameterisation. To minimise overfitting and reduce the sensitivity of the outputs to translations in the input space, after each convolutional layer, max pooling layers will be added with a kernel size of 2 to reduce the parameterisation of the model (Goodfellow et al., 2016).

**3.2.2. Results**

The assessment for the Delft3D PINNs surrogate is the same as outlined previously in §3.1.2. Cross-validation is performed by partitioning data according to the discrete simulations the model was used for, resulting in 5-fold CV for model assessment. Similarly, as described in Section 2.5, the evaluation of the surrogate model is conducted primarily using the ordinary version of RMSE instead of the modified RMSE metric utilised in the LFP surrogate model. The modified RMSE was used in LFP simulations as most cells remained dry, i.e., had a fluid depth of 0, across the simulations, and hence were excluded from the assessment. However, in the case of the Delft3D model, most cells contained high levels of fluid depth across the simulations, and only a minority of cells were dry.

The results for the developed surrogate models are outlined in Table 3, using the same criteria as previously used for the case of the LFP

surrogate models. The PINNs surrogate model outperformed the data-driven CNN model, resulting in a notable increase in performance. The average reduction in the error rate for the modified RMSE metric was approximately 11 %, highlighting the superiority of the physics-informed approach. Although this improvement was smaller compared to the previous case study, the PINNs model still demonstrated a considerable advantage over the purely data-driven surrogate model. However, although the expected reduction in error is not substantial, it can be concluded that an average reduction in error of over 10 % is not a result of random noise, indicating that the proposed PINNs-based approach outperforms the purely data-driven surrogate model. Furthermore, the PINNs model achieves a reduction of approximately 20 % in the volume discrepancy, further highlighting its superior performance.

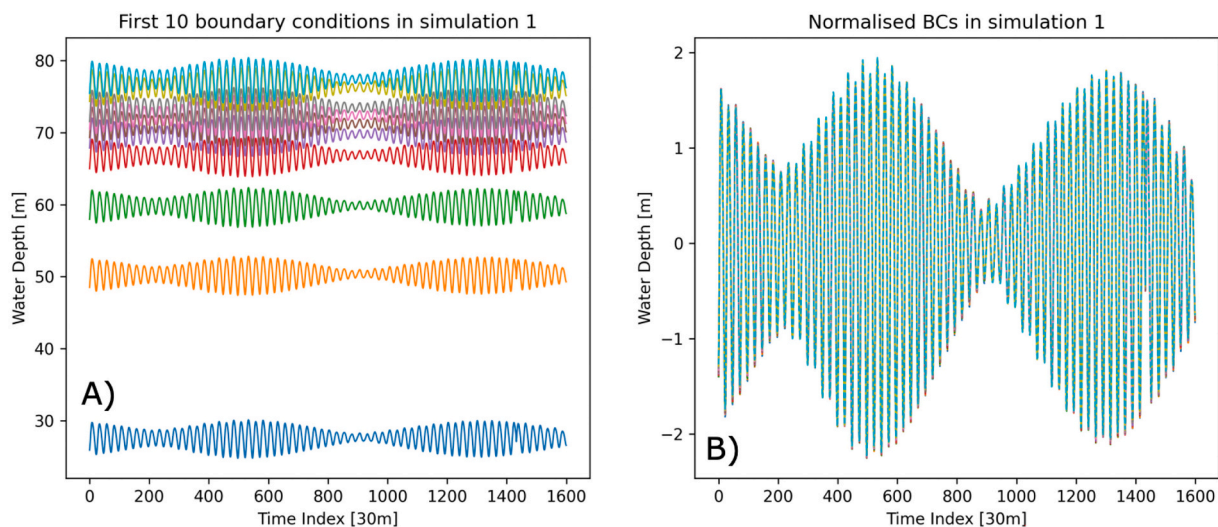
To illustrate predictive performance, Figs. 12 and 13 show the PINNs model's predictions at the six validation tide gauge locations outlined in Fig. 10. Fig. 12 shows a snapshot of the predicted values, highlighting a random period of 500 timesteps (approximately 10 days) for each of the locations, and Fig. 13 outlines the macro view, showing the aggregated predictions across all CV partitions (which can be seen through negligible discontinuities in the figures). These figures demonstrate the remarkable accuracy achieved by the PINNs surrogate model and its ability to handle the significant variability in the mean water surface elevation across cells, as evidenced by the diverse range of average depths at different locations.

Further, the overall spatial distribution of errors for the PINNs model can be observed in Fig. 14. The figure displays the aggregated validation RMSE in every cell in the domain across all timesteps. The errors are generally low in the majority of the domain but moderately higher in small sub-domains. These subdomains correspond to locations where the coastal geometry and topography are complex, resulting in tidal dynamics that deviate from those of the rest of the domain, leading to lower model performance.

To show how the model performance varies temporally, Fig. 15a shows the average error of the model over a 4-hour period of high tide at the beginning of March and Fig. 15b shows the average error of the

**Table 3**  
Mean CV scores for the CNN and PINNs models on the Delft3D data.

Model	mRMSE	RMSE	Vol. Discrepancy
CNN (MSE)	0.154	0.150	$2.26 \times 10^6$
PINNs	0.139	0.137	$1.84 \times 10^6$



**Fig. 11.** Boundary condition values for 10 of the total 40 boundary locations for the first simulation period. A) Original untransformed data. B) Centered data after subtracting the mean.



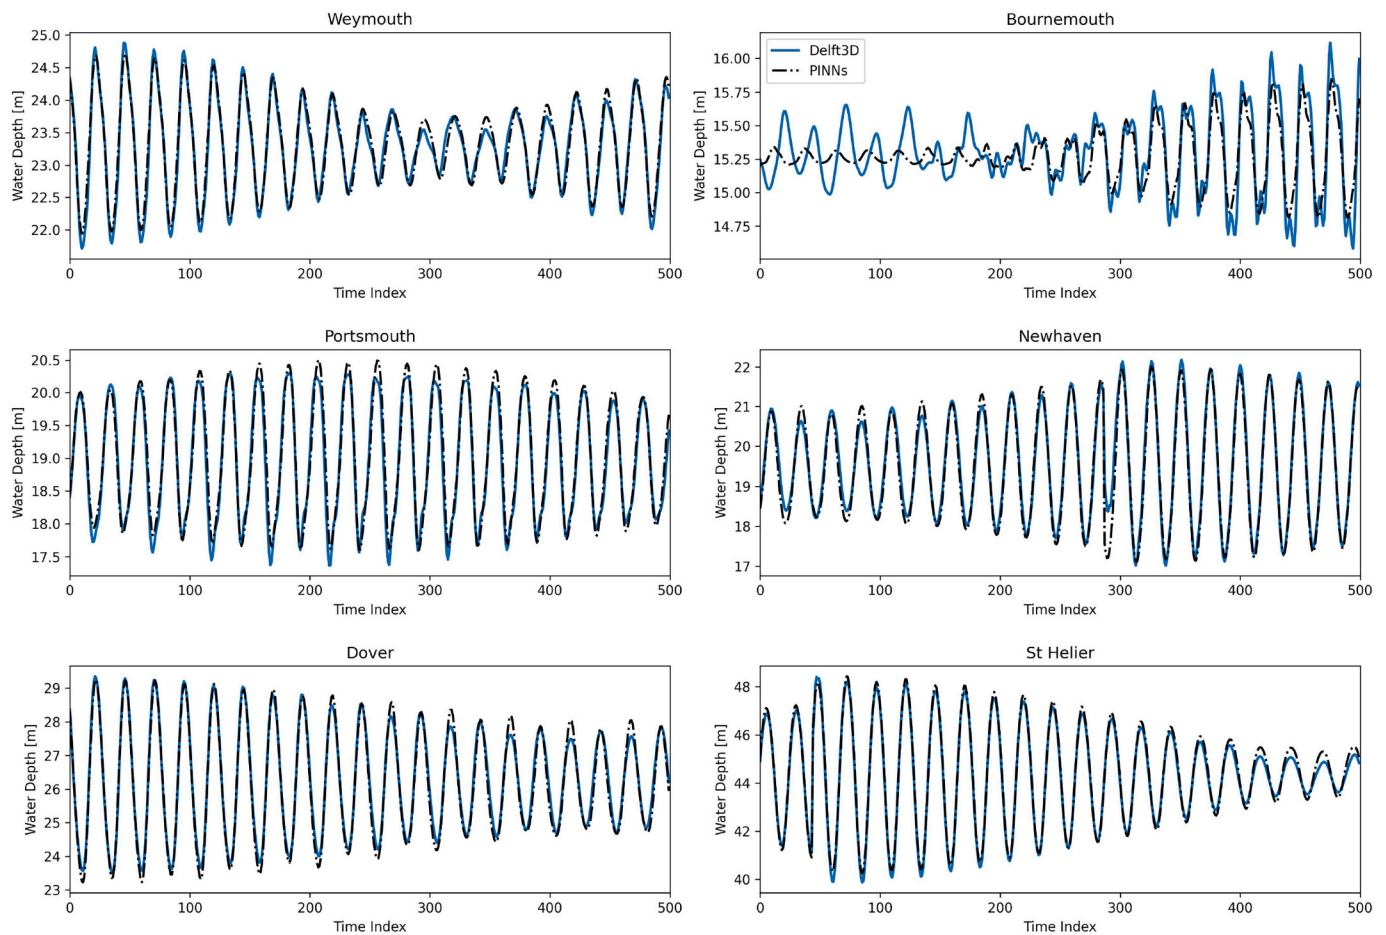


Fig. 12. PINNs predictive performance over random 500-timestep periods for each validation location in Fig. 10.

model over another 4-hour period of high-tide at the beginning of September. In Fig. 15a, the overall tidal level is considerably higher than it is in Fig. 15b, likely leading to sub-domains of higher losses, i.e., areas where the true water levels are higher relative to the averages observed in the training data. The performance of the model will vary with respect to how similar the true values are to the training data. The overall extrapolation issue, however, by training the surrogate on a representative sample of data and incorporating prior information through PINNs these issues can be mitigated.

#### 4. Discussions

The utilisation of neural network-based surrogate models in flood modelling has presented several advantages, including significantly reduced computational and time requirements for inferring new scenarios. By explicitly introducing prior knowledge of the system being modelled into the structure of the surrogate, i.e., such as outlined here with physics-based regularisation, a more robust ML-based model can be produced. Physics-informed approaches to ML look to have the potential to partially resolve existing issues in supervised ML, such as data efficiency, explainability of predictions, and generalisation performance. Moreover, by leveraging the capabilities of neural networks, such models can be constructed with no additional computational overhead at inference time and can be highly scalable; adapting well to large datasets of high-dimensional feature spaces.

A primary limitation of a PINNs-based approach, in contrast to purely data-driven neural network, is the increased development skills and knowledge needed to develop these models, as demonstrated in Section 2.3. For instance, the design of the custom loss function used in

PINNs is more demanding, especially when compared to standard loss functions available in widely-used neural network frameworks. In addition, the structure of the PINNs models, which incorporates recurrent connections to past and future target values, necessitates a modified training routine, further adding to the complexity. However, many PINNs models, like the surrogate outlined in this study, will involve more computation during model training but once the model is trained, the cost of inference will be the same as it would for a data-driven surrogate.

Surrogate models generally have the same inherent limitations as supervised ML models do. Their ability to generalise well and extrapolate to new, unseen data that was not observed in the training data can be highly varied. Considering this, and since the surrogate models are usually trained on limited data points, achieving highly accurate predictive performance on unseen data that differs from the training data is not guaranteed. As a result, changes in underlying data distributions (i.e., drift), such as new spatial locations, can necessitate model re-training or adjustment to ensure achieving the desired accuracy. This poses challenges in ensuring consistent reliability across varying scenarios, highlighting the importance of ongoing performance monitoring and validation.

Furthermore, the process of creating surrogate models can be laborious and resource-intensive. The development of a surrogate model not only requires a well-calibrated numerical model, with all the attendant calibration and validation steps, but also demands careful experimental design, data selection, pre-processing, and training. Although the benefits of having a trained surrogate can be significant in terms of reduced computational requirements for simulations, the initial investment in terms of time and effort is substantial. Given the inherent limitations, it

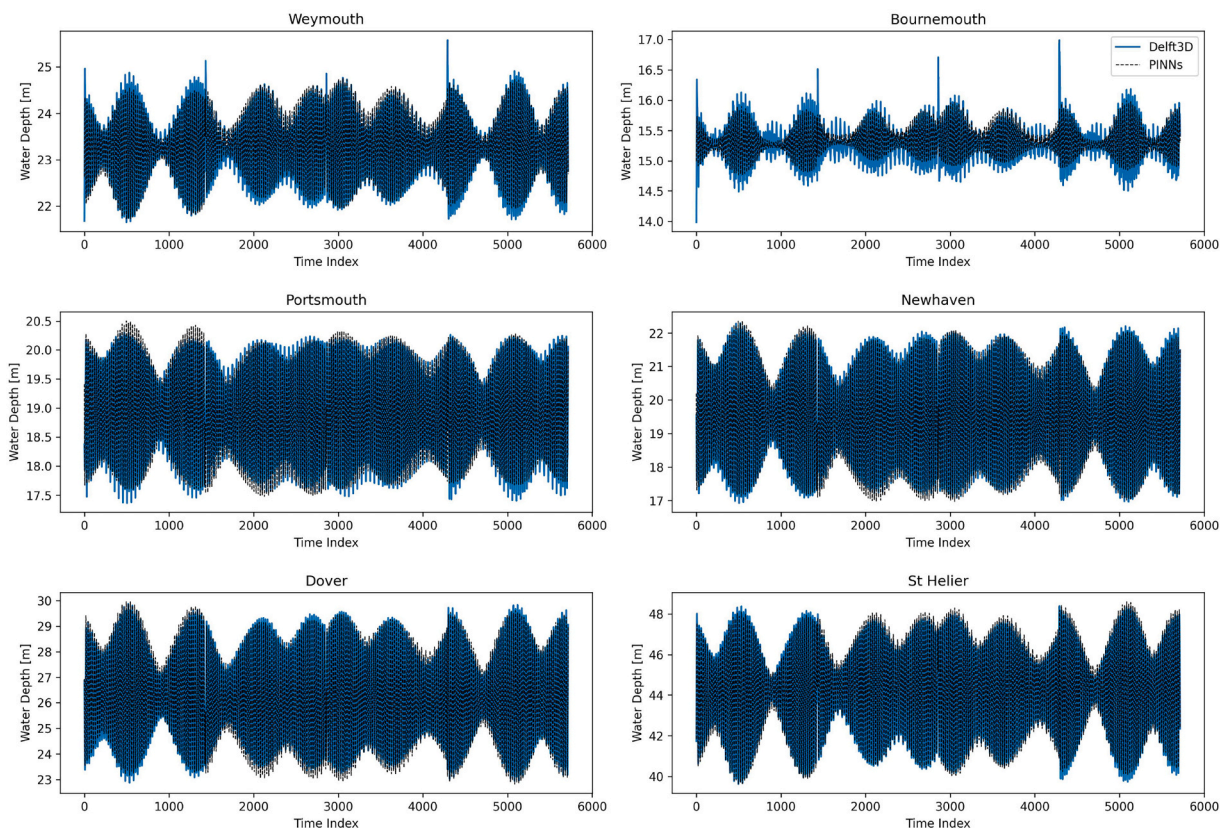


Fig. 13. Full predictive performance for the PINNs model with predictions aggregated across CV partitions.




Fig. 14. Spatial distribution of errors for the PINNs model showing aggregated validation RMSE in every cell across all timesteps.


is vital for researchers and practitioners to carefully weigh the pros and cons of developing a surrogate model.

5. Conclusions


This study introduced methodological innovations by outlining a

novel PINN model based on discrete conservation laws and applying it to two numerical hydrodynamic simulators, LISFLOOD-FP and Delft3D. The hydrodynamic simulators and both the PINNs and data-driven CNN-based surrogates developed in this study have proven effective in both inland flood and tidal modelling, showcasing the versatility and robustness of the proposed approach. The findings of this study





(a) March



(b) September

**Fig. 15.** Average error (RMSE) for the 4-hour high-tide periods of different simulation periods (high-tide predictive performance in different months and seasonal patterns).

demonstrate that the proposed PINNs-based surrogate model can be applied universally to accurately approximate any hydrodynamic simulator with high efficiency and accuracy. The novelty of the proposed method involves aggregating predicted quantities of interest and complex boundary conditions over time, which enables the direct quantification of violations of the conservation laws governing the system. By including additional terms to the loss function that penalize predictions violating these laws to a greater extent, a more robust surrogate model for hydrodynamic simulators has been developed. Furthermore, the surrogate developed could be extended to predict other conserved quantities in different applications.

By extending the existing PINNs approaches, no compromises on the functionality of the surrogate model were required. Unlike existing

PINNs approaches, which typically require spatiotemporal coordinates as inputs to the network, generating latent solutions under the same boundary conditions as the training outputs, the proposed approach allows for greater flexibility in the specification of boundary conditions and the prediction of quantities of interest. Here, the ability to flexibly parameterise boundary conditions in the network is maintained such that the surrogate model can be directly applied to entirely new boundary conditions. The proposed model is able to accurately predict quantities of interest at all spatial locations simultaneously. Additionally, this study demonstrates that by using a convolutional neural network architecture, the resulting PINNs approach can be easily applied to architectures beyond dense, fully-connected networks.

The results showed that in both experiments, for LISFLOOD-FP and

Delft3D, the PINNs model outperformed a benchmark CNN model in which the same network architecture was trained using an ordinary MSE loss function. The examples presented in this study demonstrate the high degree of accuracy provided by the PINNs model, surpassing the data-driven CNN by approximately 25 % and 11 % for the two respective applications. Furthermore, the proposed PINNs model was found to exhibit more stable optimisation behaviour, especially around the minima, although convergence took more weight updates, which is due to including the physical laws governing the system and boundary conditions. The impact of random weight initializations in the networks has been taken into account to condition for their effects. Therefore, any performance differences between the models are most likely due to the optimization criteria of the networks, which adds a high degree of certainty to the results obtained.

When using neural network-based models generally there is no straightforward approach to directly quantify the uncertainty associated with the predictions. In complex settings such as environmental forecasting, where predictions can have far-reaching implications through planning and policy, the ability to robustly consider predictive uncertainty allows for a more sophisticated interpretation of predictions. Other ML-based surrogate models such as Gaussian Processes do allow for this type of easy uncertainty quantification. However, they lack the same scalability as neural networks with respect to large, high-dimensional datasets. GPs also allow for the incorporation of prior information through a Bayesian framework, although this framework may require assumptions, which limit the complexity of prior information (Owhadi et al., 2015) and so may not be as appropriate for implementing physics-informed models. Future research directions should focus on utilising the scalability and flexibility of PINNs models with the uncertainty quantification capabilities of Bayesian methods.

It should be noted that the PIML approaches represent a unique combination of existing physics-based numerical modelling approaches and purely data-driven ML-based approaches. This study recommends that, in cases where a comprehensive knowledge of the governing physics is available, physics-informed approaches should be more widely adopted in surrogate modelling of scientific and engineering applications. These approaches could provide higher accuracy, more explainable forecasts and better data efficiency. In an era of proliferate supercomputer usage and a rapidly changing climate, there is a more urgent need than ever to develop low computational cost, rapid approaches to forecasting important environmental and climate variables.

#### CRedit authorship contribution statement

**James Donnelly:** Conceptualization, Methodology, Software, Validation, Writing – original draft, Writing – review & editing. **Alireza Daneshkhah:** Supervision, Writing – review & editing. **Soroush Abolfathi:** Supervision, Writing – review & editing.

#### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Data availability

Data will be made available on request.

#### Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.scitotenv.2023.168814>.

#### References

- Al-Jamimi, H.A., BinMakhashen, G.M., Saleh, T.A., 2022. Artificial intelligence approach for modeling petroleum refinery catalytic desulfurization process. *Neural Comput. & Applic.* 34 (20), 17809–17820.
- Almajid, M.M., Abu-Al-Saud, M.O., 2022. Prediction of porous media fluid flow using physics informed neural networks. *J. Pet. Sci. Eng.* 208, 109205.
- Bates, P.D., 2022. Flood inundation prediction. *Annu. Rev. Fluid Mech.* 54, 287–315.
- Bates, P.D., De Roo, A., 2000. A simple raster-based model for flood inundation simulation. *J. Hydrol.* 236 (1–2), 54–77.
- Brunton, S.L., Noack, B.R., Koumoutsakos, P., 2020. Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* 52, 477–508.
- Cai, S., Wang, Z., Wang, S., Perdikaris, P., Karniadakis, G.E., 2021. Physics-informed neural networks for heat transfer problems. *J. Heat Transf.* 143 (6).
- Chang, E.T., Strong, M., Clayton, R.H., 2015. Bayesian sensitivity analysis of a cardiac cell model using a gaussian process emulator. *PLoS One* 10 (6), e0130252.
- Chen, Y., Lu, L., Karniadakis, G.E., Dal Negro, L., 2020. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Opt. Express* 28 (8), 11618–11633.
- Conti, S., Gosling, J.P., Oakley, J.E., O'Hagan, A., 2009. Gaussian process emulation of dynamic computer codes. *Biometrika* 96 (3), 663–676.
- Cuomo, S., Di Cola, V.S., Giampaolo, F., Rozza, G., Raissi, M., Piccialli, F., 2022. Scientific Machine Learning Through Physics-informed Neural Networks: Where We Are and What's Next. *arXiv preprint arXiv:2201.05624*.
- Deltares, 2022. Delft3D-FLOW, Simulator of Multi-dimensional Hydrodynamic Flows and Transport, User Manual.
- Depina, I., Jain, S., Mar Valsson, S., Gotovac, H., 2022. Application of physics-informed neural networks to inverse problems in unsaturated groundwater flow. In: *Georisk: Assessment and Management of Risk for Engineered Systems and Geohazards*, 16 (1), pp. 21–36.
- Donnelly, J., Abolfathi, S., Pearson, J., Chatrabgoun, O., Daneshkhah, A., 2022. Gaussian process emulation of spatio-temporal outputs of a 2d inland flood model. *Water Res.* 225, 119100.
- Donnelly, J., Abolfathi, S., Daneshkhah, A., 2023. A Physics-informed Neural Network surrogate model for tidal simulations. *ECCOMAS Procedia* 836–844.
- Donnelly, J., Daneshkhah, A., Abolfathi, S., 2024. Forecasting global climate drivers using Gaussian processes and convolutional autoencoders. *Eng. Appl. Artif. Intell.* 128, 107536.
- Fanou, M., Daneshkhah, A., Eden, J.M., Remesan, R., Palade, V., 2023a. Hydro-morphodynamic modelling of mangroves imposed by tidal waves using finite element discontinuous Galerkin method. *Coast. Eng.* 182, 104303.
- Fanou, M., Daneshkhah, A., Eden, J.M., Yang, J., See, S., Palade, V., 2023b. Physics informed neural networks to model the hydro-morphodynamics of mangrove environments. *ECCOMAS Procedia* 822–835.
- Feng, D., Tan, Z., He, Q., 2023. Physics-informed neural networks of the saint-venant equations for downscaling a large-scale river model. *Water Resour. Res.* 59 (2), e2022WR033168.
- Garca-Alba, J., B arcena, J.F., Ugarteburu, C., Garc a, A., 2019. Artificial neural networks as emulators of process-based models to analyse bathing water quality in estuaries. *Water Res.* 150, 283–295.
- Goodfellow, I., Bengio, Y., Courville, A., 2016. *Deep Learning*. MIT press.
- Horstman, E., Dohmen-Janssen, M., Hulscher, S., 2013. Modeling tidal dynamics in a mangrove creek catchment in delft3d. In: *Coastal Dynamics*, Vol. 2013, pp. 833–844.
- Jagtap, A.D., Mao, Z., Adams, N., Karniadakis, G.E., 2022. Physics-informed neural networks for inverse problems in supersonic flows. *J. Comput. Phys.* 466, 111402.
- Kabir, S., Patidar, S., Xia, X., Liang, Q., Neal, J., Pender, G., 2020. A deep convolutional neural network model for rapid prediction of fluvial flood inundation. *J. Hydrol.* 590, 125481.
- Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., Yang, L., 2021. Physics-informed machine learning. *Nat. Rev. Phys.* 3 (6), 422–440.
- Karpatne, A., Watkins, W., Read, J., Kumar, V., 2017. Physics-guided Neural Networks (pgnn): An Application in Lake Temperature Modeling. *arXiv preprint arXiv:1710.11431*, 2.
- Kashinath, K., Mustafa, M., Albert, A., Wu, J., Jiang, C., Esmailzadeh, S., Azizzadenesheli, K., Wang, R., Chattopadhyay, A., Singh, A., et al., 2021. Physics-informed machine learning: case studies for weather and climate modelling. *Phil. Trans. R. Soc. A* 379 (2194), 20200093.
- Kochkov, D., Smith, J.A., Alieva, A., Wang, Q., Brenner, M.P., Hoyer, S., 2021. Machine learning-accelerated computational fluid dynamics. *Proc. Natl. Acad. Sci.* 118 (21), e2101784118.
- Kumbier, K., Carvalho, R.C., Vafeidis, A.T., Woodroffe, C.D., 2018. Investigating compound flooding in an estuary using hydrodynamic modelling: a case study from the Shoalhaven River, Australia. *Nat. Hazards Earth Syst. Sci.* 18 (2), 463–477.
- Kutz, J.N., 2017. Deep learning in fluid dynamics. *J. Fluid Mech.* 814, 1–4.
- Lateef, S.A., Oyeohan, I.A., Oyeohan, T.A., Saleh, T.A., 2022. Intelligent modeling of dye removal by aluminumized activated carbon. *Environ. Sci. Pollut. Res.* 29 (39), 58950–58962.
- Lin, Q., Leandro, J., Wu, W., Bhola, P., Disse, M., 2020. Prediction of maximum flood inundation extents with resilient backpropagation neural network: case study of Kulmbach. *Front. Earth Sci.* 8, 332.
- Liu, B., Tang, Q., Zhao, G., Gao, L., Shen, C., Pan, B., 2022. Physics-guided long short-term memory network for streamflow and flood simulations in the lancang-mekong river basin. *Water* 14 (9), 1429.
- Longobardi, S., Lewalle, A., Coveney, S., Sjaastad, I., Espe, E.K., Louch, W.E., Musante, C. J., Sher, A., Niederer, S.A., 2020. Predicting left ventricular contractile function via



- gaussian process emulation in aortic-banded rats. *Phil. Trans. R. Soc. A* 378 (2173), 20190334.
- Lütjens, B., Leshchinskiy, B., Requena-Mesa, C., Chishtie, F., Díaz-Rodríguez, N., Boulais, O., Piña, A., Newman, D., Lavin, A., Gal, Y., et al., 2020. Physics-informed Gans for Coastal Flood Visualization. *arXiv preprint arXiv:2010.08103*.
- Lütjens, B., Leshchinskiy, B., Requena-Mesa, C., Chishtie, F., Díaz-Rodríguez, N., Boulais, O., Sankaranarayanan, A., Pina, A., Gal, Y., Rassi, C., et al., 2021. Physically-consistent Generative Adversarial Networks for Coastal Flood Visualization. *arXiv preprint arXiv:2104.04785*.
- Mao, Z., Jagtap, A.D., Karniadakis, G.E., 2020. Physics-informed neural networks for high-speed flows. *Comput. Methods Appl. Mech. Eng.* 360, 112789.
- Massoud, E.C., 2019. Emulation of environmental models using polynomial chaos expansion. *Environ. Model Softw.* 111, 421–431.
- Moreno-Rodenas, A.M., Bellos, V., Langeveld, J.G., Clemens, F.H., 2018. A dynamic emulator for physically based flow simulators under varying rainfall and parametric conditions. *Water Res.* 142, 512–527.
- Mosavi, A., Ozturk, P., Chau, K.-W., 2018. Flood prediction using machine learning models: literature review. *Water* 10 (11), 1536.
- Neal, J., Schumann, G., Fewtrell, T., Budimir, M., Bates, P., Mason, D., 2011. Evaluating a new lisflood-fp formulation with data from the summer 2007 floods in Tewkesbury, UK. *J. Flood Risk Manag.* 4 (2), 88–95.
- O'Hagan, A., 2006. Bayesian analysis of computer code outputs: a tutorial. *Reliab. Eng. Syst. Saf.* 91 (10–11), 1290–1300.
- O'Loughlin, F.E., Neal, J., Schumann, G., Beighley, E., Bates, P.D., 2020. A lisflood-fp hydraulic model of the middle reach of the Congo. *J. Hydrol.* 580, 124203.
- Owhadi, H., Scovel, C., Sullivan, T., 2015. Brittleness of bayesian inference under finite information in a continuous world. *Electron. J. Stat.* 9 (1), 1–79.
- Qian, K., Mohamed, A., Claudel, C., 2019. Physics Informed Data Driven Model for Flood Prediction: Application of Deep Learning in Prediction of Urban Flood Development. *arXiv preprint arXiv:1908.10312*.
- Raissi, M., Perdikaris, P., Karniadakis, G.E., 2017. Physics Informed Deep Learning (part i): Data-driven Solutions of Nonlinear Partial Differential Equations. *arXiv preprint arXiv:1711.10561*.
- Raissi, M., Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378, 686–707.
- Raissi, M., Yazdani, A., Karniadakis, G.E., 2020. Hidden fluid mechanics: learning velocity and pressure fields from flow visualizations. *Science* 367 (6481), 1026–1030.
- Roelvink, J., Van Banning, G., 1995. Design and development of delft3d and application to coastal morphodynamics. *Oceanogr. Lit. Rev.* 11 (42), 925.
- Sahli Costabal, F., Yang, Y., Perdikaris, P., Hurtado, D.E., Kuhl, E., 2020. Physics-informed neural networks for cardiac activation mapping. *Front. Phys.* 8, 42.
- Shi, B., Wang, P., Jiang, J., Liu, R., 2018. Applying high-frequency surrogate measurements and a wavelet-ann model to provide early warnings of rapid surface water quality anomalies. *Sci. Total Environ.* 610, 1390–1399.
- Shustikova, I., Domeneghetti, A., Neal, J.C., Bates, P., Castellarin, A., 2019. Comparing 2d capabilities of hec-ras and lisflood-fp on complex topography. *Hydrol. Sci. J.* 64 (14), 1769–1782.
- Skinner, C.J., Coulthard, T.J., Parsons, D.R., Ramirez, J.A., Mullen, L., Manson, S., 2015. Simulating tidal and storm surge hydraulics with a simple 2d inertia based model, in the Humber Estuary, UK. *Estuar. Coast. Shelf Sci.* 155, 126–136.
- Sudret, B., Marelli, S., Wiart, J., 2017. Surrogate models for uncertainty quantification: an overview. In: 2017 11th European Conference on Antennas and Propagation (EUCAP). IEEE, pp. 793–797.
- Vatvani, D., Zweers, N., Van Ormondt, M., Smale, A., De Vries, H., Makin, V., 2012. Storm surge and wave simulations in the gulf of mexico using a consistent drag relation for atmospheric and storm surge models. *Nat. Hazards Earth Syst. Sci.* 12 (7), 2399–2410.
- Vousdoukas, M.I., Voukouvalas, E., Mentaschi, L., Dottori, F., Giardino, A., Bouziotas, D., Bianchi, A., Salamon, P., Feyen, L., 2016a. Developments in large-scale coastal flood hazard mapping. *Nat. Hazards Earth Syst. Sci.* 16 (8), 1841–1853.
- Vousdoukas, M.I., Voukouvalas, E., Annunziato, A., Giardino, A., Feyen, L., 2016b. Projections of extreme storm surge levels along europe. *Clim. Dyn.* 47 (9), 3171–3190.
- Waldman, S., Baston, S., Nimalidinne, R., Chatzirodou, A., Venugopal, V., Side, J., 2017. Implementation of tidal turbines in mike 3 and delft3d models of pentland firth & orkney waters. *Ocean Coast. Manag.* 147, 21–36.
- Wang, J., Balaprakash, P., Kotamarthi, R., 2019. Fast domain-aware neural network emulation of a planetary boundary layer parameterization in a numerical weather forecast model. *Geosci. Model Dev.* 12 (10), 4261–4274.
- Yan, Y., Li, X., Sun, W., Fang, X., He, F., Tu, J., 2023. Semi-surrogate modelling of droplets evaporation process via XGBoost integrated CFD simulations. *Sci. Total Environ.* 895, 164968 <https://doi.org/10.1016/j.scitotenv.2023.164968>.
- Yang, J., Jakeman, A., Fang, G., Chen, X., 2018. Uncertainty analysis of a semi-distributed hydrologic model based on a gaussian process emulator. *Environ. Model Softw.* 101, 289–300.
- Yang, Q., Yuan, Q., Gao, M., Li, T., 2023. A new perspective to satellite-based retrieval of ground-level air pollution: simultaneous estimation of multiple pollutants based on physics-informed multi-task learning. *Sci. Total Environ.* 857, 159542.
- Yang, S.-N., Chang, L.-C., 2020. Regional inundation forecasting using machine learning techniques with the internet of things. *Water* 12 (6), 1578.
- 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. *J. Comput. Phys.* 397, 108850.