

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

## Comput. Methods Appl. Mech. Engrg.

journal homepage: [www.elsevier.com/locate/cma](http://www.elsevier.com/locate/cma)

# Quality measures for the evaluation of machine learning architectures on the quantification of epistemic and aleatoric uncertainties in complex dynamical systems

Stephen Guth, Alireza Mojahed, Themistoklis P. Sapsis\*

Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

## ARTICLE INFO

### Keywords:

Uncertainty quantification  
Quality measures  
Gaussian process  
Ensemble neural networks  
Reduced order modeling

## ABSTRACT

Machine learning methods for the construction of data-driven reduced order model models are used in an increasing variety of engineering domains, especially as a supplement to expensive computational fluid dynamics for design problems. An important check on the reliability of surrogate models is Uncertainty Quantification (UQ), a self assessed estimate of the model error. Accurate UQ allows for cost savings by reducing both the required size of training data sets and the required safety factors, while poor UQ prevents users from confidently relying on model predictions. We introduce two quality measures for the quantification of uncertainty that are suitable for high dimensional problems: the distribution of normalized residuals on validation data and the distribution of estimated uncertainties. While the distribution of estimated uncertainties evaluates how confident the model is in making predictions, the distribution of normalized residuals evaluates how accurate these uncertainty estimates are. We examine several machine learning techniques, including both Gaussian processes and a family UQ-augmented neural networks: Ensemble neural networks (ENN), Bayesian neural networks (BNN), Dropout neural networks (D-NN), and Gaussian neural networks (G-NN). We evaluate UQ accuracy (distinct from model accuracy) We apply these metrics to two model data sets representative of complex dynamical systems: an ocean engineering problem in which a ship traverses irregular wave episodes, and a dispersive wave turbulence system with extreme events, the Majda–McLaughlin–Tabak model. We present conclusions concerning model architecture and hyperparameter tuning for use in applications of UQ to dynamical systems, such as active learning schemes.

## 1. Introduction

For a wide variety of engineering problems, numerical predictions are only a first step—also necessary is an estimate of the result error. The economic costs of unconstrained errors are evidenced by the costs of conservative tolerances and safety factors. For design problems, uncertainty quantification (UQ) is the general problem of estimating the magnitude of the error between model results and the (unobserved) ground truth [1,2], and it has wide application across many modeling domains [3–6]. Traditional design methodologies based on theory or simulation allow for error propagation, sensitivity analysis, and convergence studies, however these approaches are difficult to extend to the black box techniques currently available from machine learning.

For UQ, the applied machine learning literature has primarily drawn from Gaussian Process (GP) techniques [7], which have built-in UQ in the form of a posterior distribution [8–10]. GP techniques, sometimes called kriging, are especially useful for active

\* Corresponding author.

E-mail address: [sapsis@mit.edu](mailto:sapsis@mit.edu) (T.P. Sapsis).

<https://doi.org/10.1016/j.cma.2024.116760>

Received 22 June 2023; Received in revised form 21 November 2023; Accepted 6 January 2024

0045-7825/© 2024 Elsevier B.V. All rights reserved.

learning (optimal experimental design) problems [11–14] and multi-fidelity problems [15–17]. However, GP techniques are limited in two significant ways: the  $O(n^3)$  scaling of computation time with the number of training samples, and the more subtle breakdown as the dimensionality of the problem increases past  $n = 2$  with the dimensionality of the problem input [18].

The other main approach in the machine learning literature is to augment neural networks with some kind of uncertainty quantification [2,19]. The first thread is to use ensemble techniques to implement the Bayes Model Average from Variational Inference [18,20]. The second thread is to combine GP and neural network technology to combine the scalability of neural networks with the posterior distribution of GP. Various works in the direction include the marginal likelihood loss function [21,22] the sparse inducing point framework [23–26], Deep Gaussian Process [27,28], and Deep Kernel Learning [29,30]. These techniques seek to avoid the  $O(n^3)$  scaling costs of exact GP inference by replacing exact calculations with approximations. We also refer to recent studies [19,31] for a comprehensive comparison for various scientific ML data sets.

While several machine learning architectures have been proposed to evaluate uncertainty, at the moment, there is an important gap on evaluating the performance of these UQ algorithms, especially in problems that ‘live’ in high dimensions, where visualization of errors is only possible through low-dimensional projections. Global metrics that compare the full probability distribution of the output is a possibility in some problems where a probability measure is available for the input space, e.g. Kullback–Leibler divergence, but they do not provide an understanding of the sources of uncertainty, i.e. aleatoric or epistemic, or whether low performance is due to inaccurate predictions of the model or inaccurate uncertainty estimation. Here we introduce two criteria that characterize the UQ performance in a holistic and informative way: (i) the distribution of normalized residuals on validation data, which measures how accurately the model estimates its own errors, and (ii) the distribution of epistemic uncertainties, which measures the typical magnitude of the model errors. The two measures are complementary as they characterize different aspects of the model and they need to be used concurrently.

We use the introduced measures to compare the UQ skill of selected data-driven methods as applied to two datasets associated with complex dynamical systems from ocean engineering and turbulent systems. While previous works have also considered comparison of different UQ methods using NN architectures [19,31], they primarily focus on low-complexity dynamical systems (e.g. KdV equation) or problems with very small uncertainty and without a strongly turbulent character. The chosen problems here are characterized by high complexity in the sense that the uncertainty of the input variables is large (i.e. not of a perturbation type), while for the first data set there is also important aleatoric uncertainty, and for the second data set there are intermittent instabilities leading to highly non-linear extreme events. In this sense the data sets have a mix of uncertainty types (measurement noise vs noise free) and a range of dimensions (from  $n = 2$  to  $n = 20$ ). We compare GP against three deep neural network-based techniques for estimating epistemic uncertainty: ensemble networks, Bayesian networks, and dropout networks; as well as one technique for estimating aleatoric uncertainty: Gaussian networks.

The structure of this paper is as follows. First, in Section 2, we present the typology of uncertainty we follow in this work, as well as the two datasets we use for evaluation. Next, in Section 3, we describe each of the machine learning architectures we examine below. In Section 4, we briefly describe the metrics we plot to examine the UQ calibration. Finally, in Section 5, we present our results. We additionally include an appendix with alternate visualizations of the uncertainty of each surrogate model in 2D restrictions of each dataset.

## 2. Uncertainty quantification

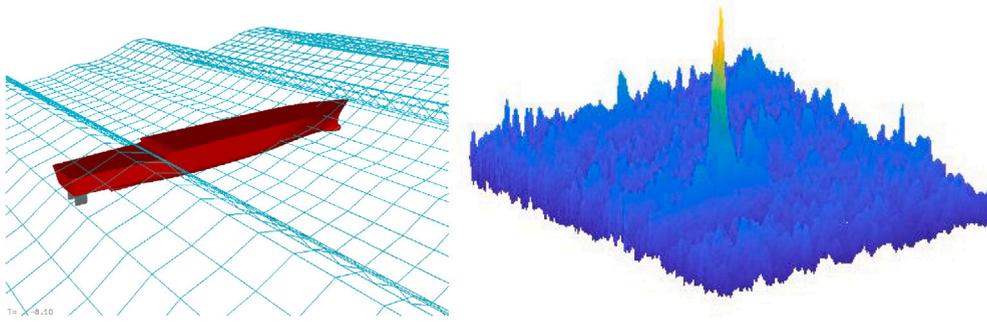
### 2.1. Typology of uncertainty

In a wide variety of modeling scenarios, it is not only important that the model be accurate (that is, generate predictions with low errors), but that the model provide accurate estimates for the its own errors. While UQ, and accurate UQ, is important for many reasons, we choose three to highlight. First, for applications that require high reliability, accurate UQ can help engineers to decide if the model reliability is sufficient for the task. Second, for some statistical tasks, measuring the uncertainty of the model is as important or more important than measuring the mean prediction. Third, an important strand of optimal experimental design involves choosing new experimental designs in regions of the parameter space where the model is most uncertain. While each of these tasks require UQ, they each examine a subtly different flavor of uncertainty. In summarizing a long literature of describing uncertainty [32], we distinguish between two types of uncertainty: aleatoric uncertainty and epistemic uncertainty.

Aleatoric uncertainty, sometimes called data uncertainty, is the irreducible uncertainty associated with unmeasurable randomness. For instance, repeated coin flip will not reduce the uncertainty of subsequent flips. The classic example is an experiment in coin flipping. No matter how much data you collect, the uncertainty in your estimate of which face the coin will show after its next flip will never decrease. In this paper, aleatoric uncertainty is the spread in outcomes when we repeat the experiment we might expect to measure if we performed the same experiment more than once.

Epistemic uncertainty, sometimes called model uncertainty, is the uncertainty associated with lack of knowledge or sparse data. For many types experiments, we expect that experiments nearby to existing samples in parameter space will yield similar results. Epistemic uncertainty is exhibited when many possible models explain the data we already have, but offer different predictions for unseen data. In general, we reduce epistemic uncertainty by collecting more data—particularly, data from novel experimental designs.

These two measures of uncertainty have different applications. For active sampling procedures, the epistemic uncertainty is the natural quantity of interest—choose the experimental design that will teach us the most. For many steady state statistical problems, the aleatoric uncertainty itself is the quantity we want to measure. Finally, for reliability questions, we are interested in the total



**Fig. 1.** Left: Sample visualization of the LAMP simulation. Here, the vessel passes through extreme waves throwing green water on deck fore and aft. Right: Sample realization of the MMT model, near an extreme event.

uncertainty—the sum of aleatoric and epistemic. The careful typology of uncertainty is important [1], because different techniques may measure epistemic and aleatoric uncertainties differently. In this paper, we refer to the standard deviation of the aleatoric uncertainty as  $\sigma_n$  and the corresponding epistemic standard deviation as  $\sigma_e$ .

The UQ challenge for model design with machine learning techniques is to produce reasonable estimates of the epistemic and aleatoric uncertainty. The standard neural network provides *no* estimate of its own uncertainty; this stands in contrast with GP models, which provide extensive estimates of their own uncertainties. In this paper, we summarize the UQ from different machine learning techniques both on a reference ocean engineering system problem, one where both epistemic and aleatoric uncertainty are significant, and a reference fluid dynamics system problem, where aleatoric uncertainty is known to be negligible.

## 2.2. Ship loads in irregular seas data set (LAMP data set)

For our first benchmark dataset, we use an ocean engineering problem setup described in [9] and used previously for benchmarking active learning ideas for extremes in [18]. Briefly, the Large Amplitude Motions Program (LAMP) v4.0.9 (May 2019) is used to simulate the structural bending moments of a marine vessel passing through a prescribed wave episode [33–36]. The primary bending moment of interest, the Vertical Bending Moment (VBM), is primary the result of hydrostatic Froude Krylov forces and slamming forces when the vessel bow first lifts clear of the water’s surface and then crashes down into the water. For both mechanisms, steeper waves lead to greater VBM, but there are important nonlinear dependencies on wave shape [9]. These nonlinearities lead to non-Gaussian VBM statistics, which can be seen in the left subfigure of Fig. 2.

The wave episodes on a finite time interval are defined by the coefficient vector  $\vec{\alpha}$  and a reduced order model (ROM) derived from a Karhunen–Loève (KL) expansion. Likewise, the VBM time series across that interval is projected onto a low dimensional subspace derived from principal component analysis. Together, this allows treating a causal system as a simple nonlinear mapping from one vector space to another.

The ML technique is applied to construct a surrogate model mapping from  $\vec{\alpha}$  (the wave episode) to  $\vec{q}$  (the resulting VBM). For practical purposes, we can reconstruct the VBM time series from the  $\vec{q}$  using the relationship

$$M_y(t, \vec{\alpha}) = \sum_{i=1}^{12} q_i(\vec{\alpha}) V_i(t), \quad t \in [0, T] \quad (1)$$

where  $V_i(t)$  are the principal components of the VBM. In [9], the authors found that the first twelve principle components are sufficient to capture the output statistics because the VBM energy spectrum decays faster than the wave elevation spectrum (shown below in Fig. 3) regardless of the number of wave components retained. The behavior is related to the nonlinear physics which prevent wave components shorter than the vessel hull length from contributing to VBM—no more than 12 of these are sufficient to capture the dominant part of the energy. In this work we limit ourselves to examining the  $\vec{\alpha}$  to  $\vec{q}$  step. Specifically, we restrict our attention to the component  $q_1$ , which, along with  $q_2$ , is the most significant term in the VBM reconstruction.

The statistics of the input coefficients are controlled by the ocean sea state, which follows a statistically stationary Gaussian distribution. For this work, we choose a steady sea state corresponding to the JONSWAP spectrum with significant wave height  $H_s = 13.2m$  and modal wave period  $\omega_m = 10s$  [37]. This is a quite extreme sea state, which leads to many extreme waves, and result in extreme internal forces, a visualization of which is shown in the left subfigure of Fig. 1. While the choice of spectrum does not *directly* affect the surrogate model, it affects the statistics of the output. In this case the map  $q_i(\vec{\alpha})$  is nonlinear and the resultant statistics highly non-Gaussian.

For each dataset, summarized in Table 1, we first choose a pair  $(n, T)$  that describes the KL representation of the wave episodes, where  $n$  is the truncation order (dimension of input) and  $T$  is the duration of the wave episode. Then, we perform  $n_s$  separate numerical experiments, where the  $\vec{\alpha}$  are drawn using Latin Hypercube Sampling (LHS) from an  $n$ -dimensional hyperbox with interval  $[-4.5, 4.5]$ . LHS is a variation of uniform sampling which avoids the density fluctuations common to true (pseudo-) random sampling.

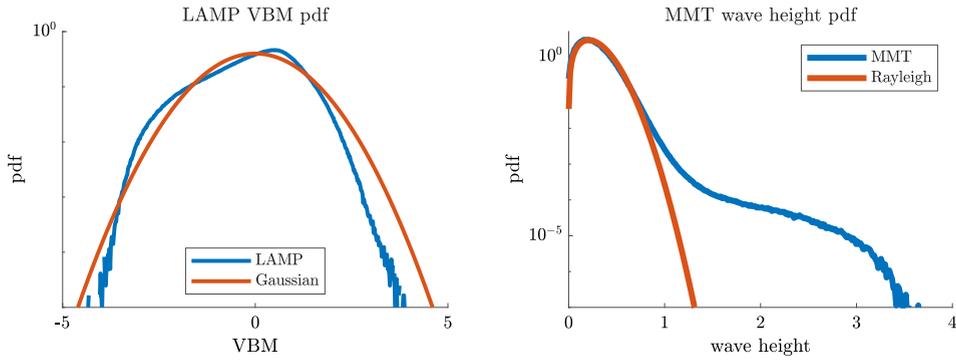


Fig. 2. Left: Total steady state statistics of the VBM for the LAMP problem. In this work, only the recovery of component  $q_1$  is considered (Eq. (2)). Right: Total steady state statistics of the wave height for the MMT problem. In this work, only the peak wave heights are considered (Eq. (6)).

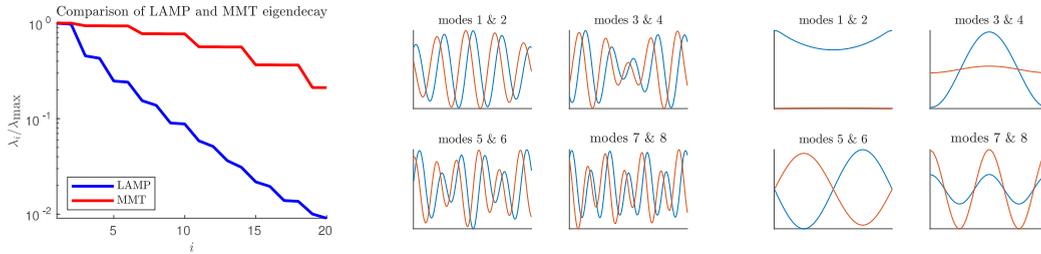


Fig. 3. Left: Comparison of the eigenspectral decay for LAMP and MMT reduced order models. Center: First few LAMP wave episode eigenmodes. Right: First few MMT initial condition eigenmodes, real parts. Eigenmodes are presented in pairs to emphasize approximate symmetries.

**Table 1**  
Selection of LAMP data sets used.

Input dimension	Interval Length	Training set	Validation set
2	40	25	675
4	40	462	238
10	40	2000	1000

The hyperbox radius,  $z^* = 4.5$ , represents  $4.5\sigma$  deviations in the choice of the components of  $\vec{\alpha}$ . Compared to representative sampling from the steady state, given by

$$\vec{\alpha} \sim \mathcal{N}(0, \mathbf{I}_n), \tag{2}$$

LHS sampling preferentially samples extreme inputs, up to the cutoff  $z^* = 4.5$ . For this reason, the LAMP data set includes more extreme wave episodes, and more extreme internal forces, than a data set that would have been sampled from the steady state distribution (Eq. (2)).

Finally, each wave episode is equipped with a stochastic prelude, [9]. The stochastic prelude is designed to extrapolate the sea surface elevation outside of the wave episode region, in order to set up the encounter conditions for the vessel. The stochastic prelude is constructed in a spectrum consistent manner so that the random encounter conditions agree with the steady state conditions while also smoothly matching the wave episode. Because the stochastic prelude is randomly chosen, it is an important source of aleatoric uncertainty in the relationship between  $\vec{\alpha}$  and  $\vec{q}$ .

We note that the KL expansion provides both a natural ordering of wave episode coefficients and the spectrum of eigenvalues representing the contribution to the wave episode energy of each dimension. Thus the  $n = 2$  data set keeps the two most significant KL modes, while the  $n = 10$  data set keeps the ten most significant modes.

### 2.3. Dispersive wave turbulence data set (MMT data set)

The Majda–McLaughlin–Tabak (MMT) model, first introduced in [38], is 1D nonlinear model of deep water wave dispersion. It has been studied in the context of weak turbulence and intermittent extreme events [18,39–44]. The MMT model is “a prototype system that generates rare events,” which are caused by the “intermittent formation and collapse of extreme events” through a Benjamin Fier type side-band instability [43,45].

Other nonlinear wave systems such as the Nonlinear Schrodinger (NLS) equation [46] display more Gaussian statistics, though heavy tailed statistics have been reported [47] and higher order approximations such as the Modified Nonlinear Schrodinger (MNLS)

**Table 2**  
Selection of MMT data sets used.

Input dimension	Training set	Validation set
2	25	975
4	462	99 538
20	10 000	90 000

equation [48] capture Benjamin Feir type side-band instabilities. Another nonlinear wave equation, the Korteweg De Vries (KdV) equation [49], has been studied in the presence of weak noise and weak damping. Unlike MMT solutions, KdV solutions do not exhibit heavy tailed statistics because there is no positive Lyapunov exponent [50], which contrasts against the unstable dimensions of the MMT attractor.

A sample visualization of an extreme event in the MMT model is shown in the right subfigure of Fig. 1, and a representation of the heavy tailed steady state statistics is shown in the right subfigure of Fig. 2. The governing equation is given by

$$iu_t = |\partial_x|^{\alpha_m} u + \lambda |\partial_x|^{\frac{\beta}{4}} \left( |\partial_x|^{\frac{\beta}{4}} u \right)^2 |\partial_x|^{\frac{\beta}{4}} u + iDu, \quad (3)$$

where  $u$  is a complex scalar, the pseudodifferential operator  $|\partial_x|^{\alpha_m}$  is defined through the Fourier transform as follows:

$$|\partial_x|^{\alpha} \widehat{u}(k) = |k|^{\alpha} \widehat{u}(k),$$

exponents  $\alpha_m$  and  $\beta$  are chosen model parameters, and  $D$  is a selective Laplacian. Our parameter selection is taken from [18], and our implementation follows [43]. Briefly, the spatial domain is periodic on  $[0, 1]$  discretized into 512 points. The parameters are chosen so that  $\lambda = -4$  (focusing case),  $\alpha = \frac{1}{2}$  (deep water waves case), and  $\beta = 0$ .

Our data set is constructed by specifying a (complex) initial condition, according to

$$u(x, t = 0) = \sum_{j=1}^m \alpha_j \sqrt{\lambda_j} \phi_j(x) \quad (4)$$

where  $\vec{\alpha} \in \mathbb{C}^m$  is a vector of complex coefficients with zero mean and unit covariance, the basis vectors  $\phi_j(x)$  are the first  $m$  eigenvectors of a KL expansion, and the normalization coefficients  $\lambda_j$  are the corresponding eigenvalues. This KL expansion is taken from a Gaussian stochastic process over  $x \in [0, 1]$ , whose autocorrelation is given by the complex periodic kernel

$$k(x, x') = \sigma_u^2 \exp(i2 \sin^2(\pi(x - x'))) \exp\left(\frac{-2 \sin^2(\pi(x - x'))}{l_u^2}\right) \quad (5)$$

with  $\sigma_u^2 = 1$  and  $l_u = 0.35$ . We present brief remarks about this process below in Section 2.4. We further note that we split the  $m$ -dimensional complex  $\vec{\alpha}$  into a  $2m$ -dimensional real  $\vec{\alpha}$  corresponding to the real and imaginary parts. We further restrict the  $\vec{\alpha}$  to lie within a hyperbox of radius  $z^* = 6$  (c.f.  $z^* = 4.5$  above in Section 2.2). Finally, we define the output map as

$$y(\vec{\alpha}) = \|\text{Re}(u(x, T = 50; \vec{\alpha}))\|_{\infty}, \quad (6)$$

that is, the output of the map is the most extreme (real)  $u$  in the domain after  $T = 50$  time units have elapsed. A table of the data sets considered in this paper is given in Table 2.

#### 2.4. Comparison of data sets

We briefly note a few features of the two data sets. In Fig. 2, we display the steady state statistics of each problem, compared to appropriately scaled Gaussian statistics. In the LAMP problem, we note the sharp left cutoff that corresponds to the bow of the ship getting submerged under the water. For the MMT problem, we emphasize the heavy right tail of extreme events corresponding to nonlinear wave focusing. For each problem, our data sets do not directly represent these steady state statistics; however, the major features we note are represented in our data sets in this work.

In the left subplot of Fig. 3, we compare the eigenspectrum decay of the reduced order models used in the LAMP dataset and the MMT dataset. This eigenspectrum describes the energy associated with each mode in the input space. For the LAMP dataset, this corresponds the KL decomposition of a JONSWAP spectrum with  $H_s = 13.2m$  and  $w_m = 10s$  [37,51]. For the MMT dataset, this corresponds to the KL decomposition of the autocorrelation kernel given in Eq. (5). As a preliminary matter, we note that the even/odd symmetry leads to paired eigenvalues for the LAMP input spectrum. For the MMT spectrum, there is both an even/odd symmetry and an (approximate) real/imaginary symmetry, leading to fourfold repeated eigenvalues.

When we compare the two eigenspectra from Fig. 3, it is clear that the eigenspectrum decays faster for the LAMP dataset than the MMT dataset. The faster decay implies that higher order modes are less important in the LAMP data set compared with the MMT data set. This implies that the even with 10D data, the LAMP dynamics are fundamentally lower dimensional than the MMT dynamics. In the right two subplots of Fig. 3, we show (the real parts of) the first eight eigenmode shapes for the LAMP and MMT datasets. The most important feature to note is that the characteristic wavelength of the low order modes for the LAMP problem are shorter than the interval length—approximately four wave periods fit into the interval. This is related to the parameters (peak

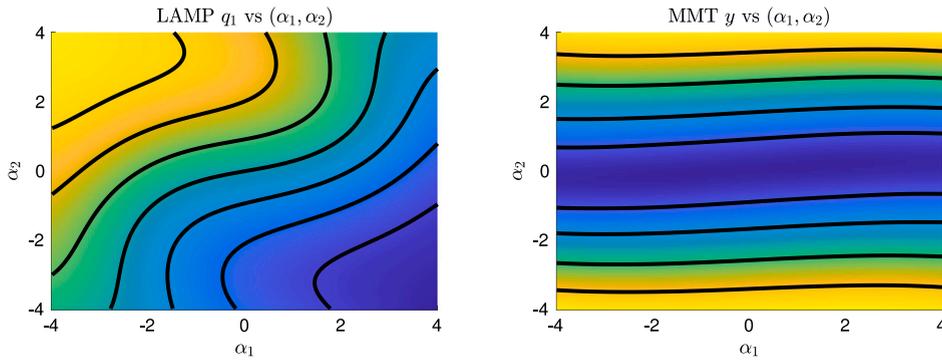


Fig. 4. Left: Plot of relationship between  $(\alpha_1, \alpha_2)$  and  $\mathbb{E}[q_1]$  for the LAMP data set. Right: Plot of the relationship between  $(\alpha_1, \alpha_2)$  and  $y$  for the MMT data set. Both plots are slices, i.e., all other inputs  $(\alpha_3, \dots)$  are zero.

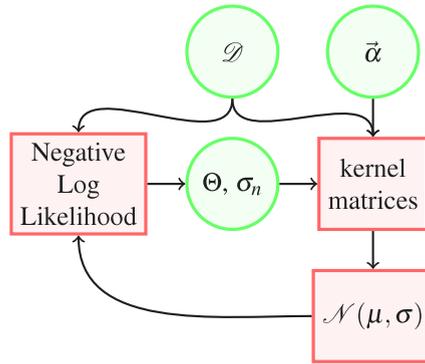


Fig. 5. Flowchart model of Gaussian Process.

frequency) of the chosen sea spectrum. For the MMT problem, the characteristic wavelength of the low order modes are longer than the interval. Also, since the autocorrelation function is defined on a periodic domain, the resulted eigenmodes are Fourier modes.

Next, the two data sets differ in the presences of aleatoric uncertainty, and intrinsic measurement error in the data. The MMT data is the product of a deterministic 1D PDE, whose only noise source is due to the limits of machine precision. The LAMP data, however, has stochastic encounter conditions associated with the stochastic prelude technique and a particular wave spectrum [9]. The stochastic encounter conditions manifest as an approximately homoskedastic aleatoric noise term in the  $q_1$  data.

Besides the energy spectrum of the input side of the data and the aleatoric error, these two problems also differ in the shape of the output side of the data. In Fig. 4, we present a slice representation of each data set, where the first two input components  $(\alpha_1, \alpha_2)$  are varied and the resulting output ( $\mathbb{E}[q_1]$  for LAMP,  $y$  for MMT) are plotted. We see at once that the spatial features in the model are very different. In particular, the LAMP data has very complicated joint dependence on the different components  $\alpha_1$  and  $\alpha_2$ . In contrast, the relationship for the MMT data is somewhat more varied, and lacks the ‘leveling off’ seen in the LAMP slice. This behavior carries through the output statistics— $p_Q(q_1)$  has short tails, while  $p_Y(y)$  has a very long right tail.

### 3. Machine learning architectures

#### 3.1. Gaussian process regression

The gold standard for uncertainty quantification in non-parametric regression is Gaussian Processes (GP). In this method, we use the training data to learn the hyperparameters of a kernel function  $k(\vec{\alpha}_1, \vec{\alpha}_2; \theta)$ , which describes the statistical relationship between model outputs for two different model inputs. The choice of kernel function, and especially its smoothness, is an important component of any GP model that encodes its inductive bias [52]. For this work, we use the squared exponential kernel with Automatic Relevance Determination (ARD), given by

$$k(\vec{\alpha}_1, \vec{\alpha}_2; \theta) = \theta_\sigma \exp\left(-\frac{1}{2} \sum_{i=1}^d \frac{(\alpha_{1,i} - \alpha_{2,i})^2}{\theta_i^2}\right). \tag{7}$$

The hyperparameters for GP include both the kernel parameters from Eq. (7) (a variance parameter  $\theta_\sigma$  and length scales  $\theta_i$ ) as well as the intrinsic noise parameter  $\sigma_n$ . The hyperparameter  $\sigma_n$  can be understood as a measure of aleatoric measurement error, or

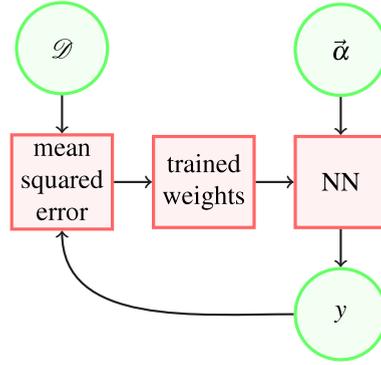


Fig. 6. Flowchart model of a traditional feed-forward Neural Network.

alternately as a numerical regularizer as from ridge regression. These hyperparameters are optimized by minimizing the Negative Log Likelihood (NLL), sometimes called negative marginal log likelihood, given by

$$-\log p(\mathbf{y}|\mathbf{A}) = \frac{1}{2} \mathbf{y}^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} + \frac{1}{2} \log |\mathbf{K} + \sigma_n^2 \mathbf{I}| + \frac{n}{2} \log 2\pi \quad (8)$$

Once we have optimized the hyperparameters, we calculate the posterior distribution with the relation

$$p(y | \bar{\alpha}) \sim \mathcal{N}(\mu(\bar{\alpha}), \sigma^2(\bar{\alpha})) \quad (9)$$

where,

$$\mu(\bar{\alpha}) = \mathbf{K}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{Y} \quad (10)$$

$$\sigma^2(\bar{\alpha}) = \mathbf{K}_{**} - \mathbf{K}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{K}_* \quad (11)$$

$\mu(\bar{\alpha})$  is the posterior mean,  $\sigma^2(\bar{\alpha})$  is the posterior variance,  $\mathbf{K}$ ,  $\mathbf{K}_*$ , and  $\mathbf{K}_{**}$  are kernel matrices, and  $\mathbf{Y}$  is the observed output data vector (for LAMP, the  $q_i$ , for MMT, the  $y$ ). Details of these calculations can be found in standard references, such as Rasmussen and Williams [7]. A schematic representation of this architecture is shown in figure Fig. 5.

The advantages of GP for surrogate modeling are twofold. First, because GP calculates a posterior distribution, it has built in UQ. Second, the internal structure of the variance calculation provides a simple decomposition of the uncertainty into an aleatoric term ( $\sigma_n$ ), and an epistemic term (the remainder).

On the other hand, GP has two major disadvantages. First, the matrix math of exact GP calculations involves inverting a matrix whose size scales like the number of training samples. This  $O(n_s^3)$  computational cost scaling prohibits the use of GP for large data sets ( $n_s \gtrsim 1000$ ), at least without significant approximations. Second, exploration of the kernel function space is difficult for feature vectors with more than a few components high dimensional spaces. As a result, GP models struggle with even modest high dimensional input spaces ( $n \gtrsim 5$ ). Because GP is restricted away from higher high dimensional input, we do not consider GP in conjunction with the functional input described below in Section 3.5.

### 3.2. Traditional neural networks

In this section we briefly recapitulate some of the features of a traditional feed-forward neural network. During the inference step, computation passes from the input layer (here,  $\bar{\alpha}$ ) through a number of hidden layers. At each hidden layer, the values from the previous layer are multiplied by a weight matrix, and then fed into a nonlinear activation function such as sigmoid or hyperbolic tangent. Common activation functions include the sigmoid and the hyperbolic tangent. In this paper, we use the rectified linear unit (ReLU). At the final (output) layer, the values from the last hidden layer are combined in a weighted sum, and a single output is reported.

The values of the weights parameters are determined in a separate training step during which. During training, the model iterates through a training set, and compares model predicted output to the training data according to some loss function. Model errors are used to adjust the network weights through the backpropagation algorithm. The literature on neural network training is vast, and we direct the interested reader to a reference such as [53] or [54]. We note that, unlike the deterministic inference (feed-forward) step, network training is generally a non-deterministic stochastic gradient descent (SGD) process.

In Fig. 6 we display a schematic flowchart of the relationship between the network input ( $\bar{\alpha}$ ), output ( $y$ ), and weights, which in turn depend on the training data ( $\mathcal{D}$ ). The architecture hyperparameters for a feed-forward neural network are the number of hidden layers, the size of each hidden layer, and the choice of activation function. Additionally, important training hyperparameters include the training time (epochs), the training rate (automatically adjusted by the Adams gradient descent), and the batch size (division of data during training). Additionally, regularization techniques, such as dropout or  $L_2$  Tikhonov regularization, may be used to combat overfitting. We emphasize that, unlike the GP case, the simple neural network produces a single (scalar) output prediction instead of a distribution. This form of prediction is insufficient for UQ, so we next turn to a series of techniques to modify neural networks in order to allow for UQ.

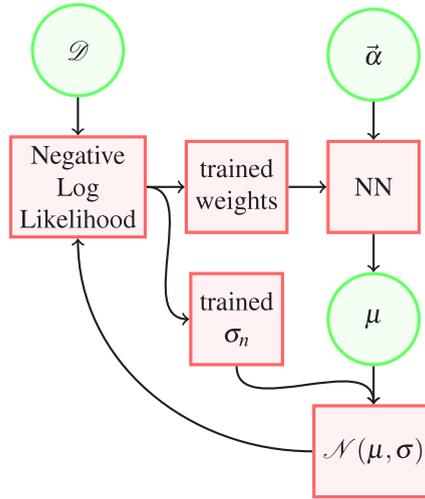


Fig. 7. Flowchart model of Gaussian Neural Network.

### 3.3. Gaussian neural networks

The first method of introducing uncertainty quantification to neural networks takes direct inspiration from GP, while attempting to avoid the iron  $O(n_s^3)$  cost of matrix inversion. Wilson et al. (2015) [24] developed KISS-GP with an approximate kernel matrix for fast inversion and determinant calculations. Calandra et al. (2016) [55] developed Manifold GP as a technique to learn an encoder and latent space representation instead of kernel function hyperparameters. Burt et al. (2020) [25] investigated using variational inference instead of an exact kernel to speed matrix inversions.

The common point of these approaches is Inducing Point GP or Deep Kernel Learning, also investigated by [30] and [29,56]. The common idea is to take a final Gaussian Process ‘layer,’ but replace the kernel matrix calculations with a different approximate learning scheme. By retaining the posterior distribution, we are able to keep the built-in UQ from GP based techniques.

In Fig. 7, we show a schematic representation of our implementation, which we call Gaussian Neural Network (G-NN). Starting from Eq. (9) for the posterior distribution, we train a neural network to learn the function  $\mu(\vec{\alpha})$ . Simultaneously, we train a parameter for the homoskedastic variance term  $\sigma_n$ . Because the G-NN outputs a distribution instead of a single valuepoint, we use the same NLL loss function for training as GP, given in Eq. (8). The G-NN model fits a homoskedastic noise parameter  $\sigma_n$  to the data set, in a close match to how GP estimates the intrinsic noise parameter  $\sigma_n$ . However, unlike, GP, G-NN does not estimate the epistemic portion of the posterior variance., which could vary with  $\vec{\alpha}$ .

We note that, like the  $\sigma_n$  term in GP, it is possible to learn a more complicated function form for the intrinsic noise term in G-NN. However, overfitting  $\sigma_n$  is a distinct danger and. Additionally, for the ocean engineering dataset we examine, a homoskedastic aleatoric uncertainty model is a reasonable approximation. We note also that, in our experience, training a G-NN on a noiseless dataset introduces numerical convergence complications and is not recommended. For those cases, such as the MMT dataset here, we recommend using a traditional deterministic output layer.

### 3.4. Prediction ensembles

One of the most used frameworks for introducing UQ into NN models is Variational Inference (VI) [57]. Under this framework, the latent variables of the model (network weights) have some conditional distribution given the observed training data. Forward inference from these exact conditional distributions would produce a predicted output distribution, whose spread would include the model epistemic uncertainty.

As computing the exact conditional distributions is computationally intractable, the literature has proposed methods to approximately sample model realizations to produce the Bayesian Model Average [19]. In this paradigm, we construct an ensemble of neural networks, each of which produces a single prediction. We then approximate the true posterior distribution with the ensemble of predictions, using the relation (written for convenience withfor a scalar output)

$$p(y | \vec{\alpha}, \mathcal{D}) \approx \frac{1}{n_e} \sum_{i=1}^{n_e} p(y | \vec{\alpha}, \hat{\theta}_i), \tag{12}$$

where  $\mathcal{D}$  is the observed training data,  $n_e$  is the ensemble size, and  $\hat{\theta}_i$  is the parametrization (weights) of the  $i^{\text{th}}$  network. In the simple case where each network produces a single prediction  $\{y_i\}$ , we compute the ensemble variance, given by

$$\sigma_{\text{ens}}^2 = \frac{1}{n_e - 1} \sum_{j=1}^{n_e} (\bar{y} - y_j)^2, \tag{13}$$

and use it as a proxy for the epistemic variance:  $\sigma_{\text{ens}}^2 \approx \sigma_{\epsilon}^2$ .

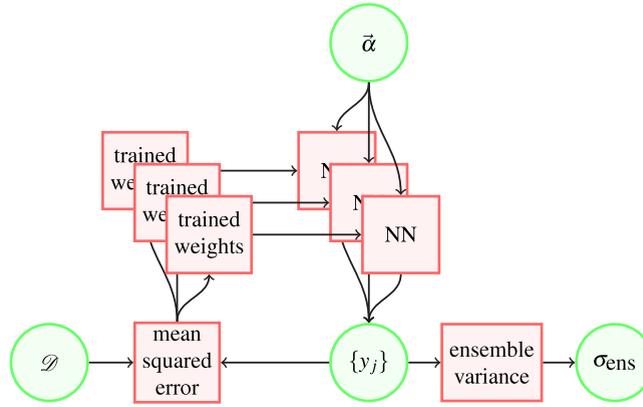


Fig. 8. Flowchart model of Ensemble Neural Network (ENN).

In the following sections, we present three techniques to produce an ensemble of predictions: Ensemble Neural Networks (ENN), Bayesian Neural Networks (BNN), and Dropout Neural Networks (D-NN). For each technique, we explain how to generate an ensemble of predictions, and how to choose/identify the parameter  $n_e$  which controls the ensemble size. Finally, we emphasize that the ensemble technique is fully compatible with the G-NN technique. That is, we can use both an ensemble technique and the Gaussian process layer to estimate aleatoric and epistemic uncertainties simultaneously. When we combine these techniques, we produce an ensemble of Gaussian distributions  $\mathcal{N}(\mu_i, \sigma_{n,i})$ . From this ensemble, we estimate

$$\sigma_n^2 = \frac{1}{n_e} \sum_{i=1}^{n_e} \sigma_{n,i}^2, \quad \sigma_\epsilon^2 = \text{Var}(\{\mu_i\}), \quad (14)$$

in a decomposition based on the law of total variance.

#### 3.4.1. Ensemble neural networks

The first technique for creating a prediction ensemble is the Ensemble Neural Network (ENN), the implementation of which we borrow from an operator network investigation by Pickering et al. (2022) [18]. The underlying idea of ENN is to train multiple neural networks with different (random) initial weights on the same training data. Then, during inference, each network is sampled separately, producing an ensemble of predictions. We simply use the ensemble mean and ensemble variance of these predictions as parameters of a Gaussian posterior distribution  $p(y_i | \vec{\alpha})$ . This architecture is shown schematically in Fig. 8.

Intuitively, this approach has promise. Near training data, the loss function will enforce that each network in the ensemble agree with one another, because they must agree with the training data. In regions far from the training data, however, there is no such pressure, and the ensemble predictions are free to diverge. Unlike the other approaches we describe, the ENN architecture has no built-in regularization term. For our results, we include an  $L_2$  weight regularizer.

**Ensemble Size.** Because each neural network must be trained separately, training time is proportional to the ensemble size,  $n_e$ . Pickering et al. [18] has suggested that just two networks ( $n_e = 2$ ) was sufficient to estimate relative uncertainties for active sampling. We found that a larger ensemble – up to  $n_e = 8$  – gave better estimates for more robust uncertainty quantification. In particular, the hyperparameter  $n_e$  trades training time and memory costs against statistical power for inferring the true ensemble-distribution variance from the ensemble-sample variance. As  $n_e$  increases, the accuracy of the ensemble variance improves. However, training time, and memory cost, is directly proportional to  $n_e$ . This cost is particularly acute in active learning applications, where many surrogates must be trained iteratively. Conversely, the computational cost is somewhat ameliorated by the ease of parallelization—each member of the ensemble may be trained independently.

In Fig. 9, we show how the distribution of the model uncertainty,  $p_s(\sigma_\epsilon)$ , i.e. the histogram all values of the estimated epistemic uncertainty over the input parameter space, vary with the ensemble size,  $n_e$ . We use this quantity as a collective measure for the overall shape of  $\sigma_\epsilon$  for high dimensional problems. For the 10D LAMP dataset, we run a parameter convergence study for the steady state statistics (left subplot, log pdf difference) and the pointwise epistemic errors (right subplot, mse). We note that, after an initial period of rapid convergence, the improvement for increasing  $n_e$  rapidly levels off (the fast drop near  $n_e = 100$  is an artifact of using  $n_e = 100$  as the convergence target). We conclude that, for ENN, there is a significant improvement from  $n_e = 2$  to  $n_e = 3$ , and noticeable improvements perhaps through  $n_e = 10$ , though the benefit is paid for by a linear increase in training (and inference) cost. For all ENN results below, we fix  $n_e = 8$ .

#### 3.4.2. Bayesian neural networks

Plain feedforward neural networks are prone to overfitting. When applied to supervised learning problems these networks are also often incapable of correctly assessing the uncertainty in the training data and so make overly confident decisions about the

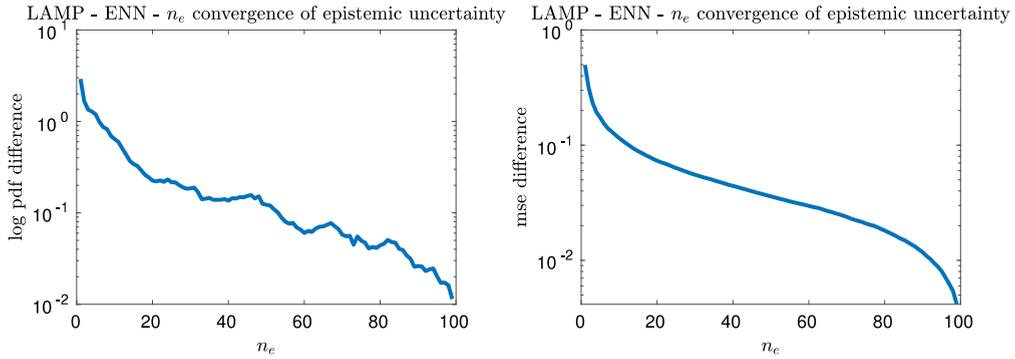


Fig. 9. Convergence curves of  $\sigma_\epsilon$  as functions of ensemble size,  $n_e$ , for ENN surrogate,  $N = 10, n_s = 2000$ . **Left:** log pdf difference:  $\int |\log p_S(\sigma_\epsilon^{n_e}) - \log p_S(\sigma_\epsilon^{100})| d\sigma_\epsilon$ . **Right:** Mean of squared differences:  $\sum (\sigma_\epsilon^{n_e} - \sigma_\epsilon^{100})^2 / n_s$ .

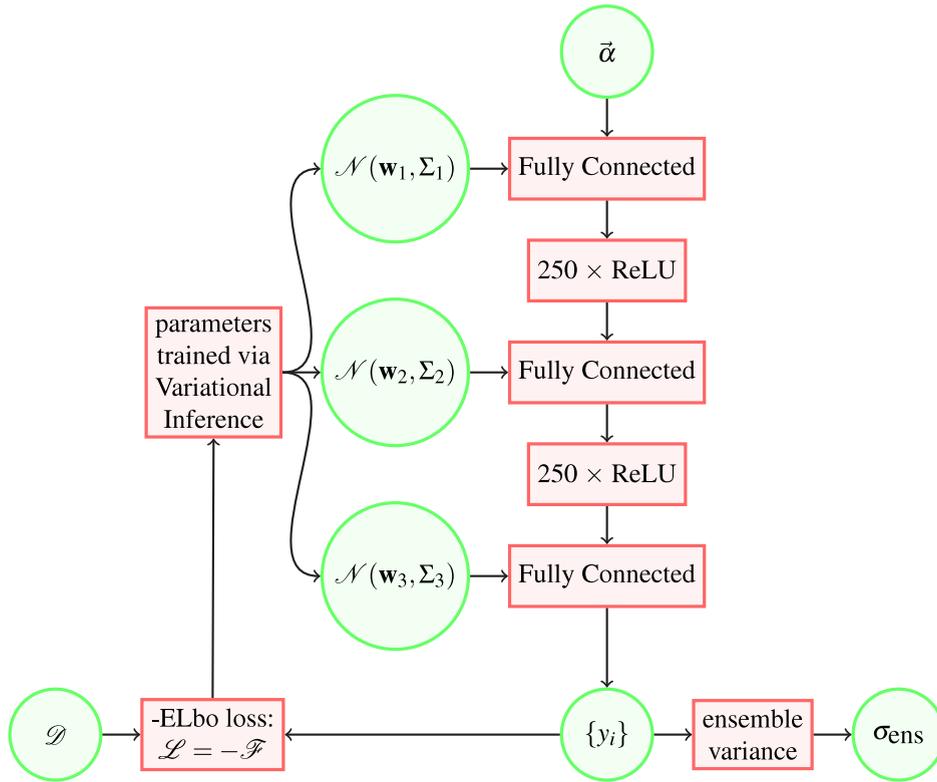


Fig. 10. Flowchart model of Bayesian Neural Network (BNN) with 2 hidden layers.

correct class, prediction or action. These issues are addressed by using variational Bayesian learning to introduce uncertainty in the weights of the network. In this approach, all the weights in the neural networks are represented by probability distributions over possible values, rather than having a single fixed value as in traditional neural networks [58,59]. In Fig. 10, we show the relationship of the Bayesian network probability distributions to the traditional neural network weights.

Learnt representations and computations must therefore be robust under perturbation of the weights, but the amount of perturbation each weight exhibits is also learnt in a way that coherently explains variability in the training data. Thus instead of training a single network, essentially an ensemble of networks are trained, where each network has its weights drawn from a shared,

learned probability distribution. Unlike other ensemble methods, this approach typically only doubles the number of parameters yet trains an infinite ensemble using unbiased Monte Carlo estimates of the gradients. In general, exact Bayesian inference on the weights of a neural network is intractable as the number of parameters is very large and the functional form of a neural network does not lend itself to exact integration. Instead we take a variational approximation to exact Bayesian updates.

**Training a Bayesian Neural Network.** Bayesian inference for neural networks calculates the posterior distribution of the weights given the training data,  $P(\mathbf{w}|D)$ . This distribution answers predictive queries about unseen data by taking expectations: the predictive distribution of an unknown label  $\hat{y}$  of a test data item  $\hat{\mathbf{x}}$ , is given by  $P(\hat{y}|\hat{\mathbf{x}}) = \mathbb{E}_{P(\mathbf{w}|D)}[P(\hat{y}|\hat{\mathbf{x}}, \mathbf{w})]$ . Each possible configuration of the weights, weighted according to the posterior distribution, makes a prediction about the unknown label given the test data item  $\hat{\mathbf{x}}$ . Thus taking an expectation under the posterior distribution on weights is equivalent to using an ensemble of an uncountably infinite number of neural networks. Unfortunately, this is intractable for neural networks of any practical size. To work around this problem, we employ variational learning, i.e., variational learning finds the parameters  $\theta$  of a distribution on the weights  $q(\mathbf{w}|\theta)$  that minimizes the Kullback–Leibler (KL) Weight Uncertainty in Neural Networks divergence with the true Bayesian posterior on the weights. The resulting cost function is known as the variational free energy or the expected lower bound (ELBo) and is expressed as:

$$F(D, \theta) = \mathbf{KL}[q(\mathbf{w}|\theta) \parallel P(\mathbf{w})] - \mathbb{E}_{q(\mathbf{w}|\theta)}[\log P(D|\mathbf{w})] \quad (15)$$

The cost function of Eq. (15) is a sum of a data-dependent part, which is referred to as the likelihood cost, and a prior-dependent part, which is referred to as the complexity cost. The cost function embodies a trade-off between satisfying the complexity of the data  $D$  and satisfying the simplicity prior  $P(\mathbf{w})$ . Eq. (15) is also readily given an information theoretic interpretation as a minimum description length cost. Exactly minimizing this cost naively is computationally prohibitive. Instead gradient descent and various approximations are used. Due to complexity of the approach and the stochastic nature of the method, the modeling and the optimization is done through the so called “Bayes By Backprop” method. For more details on the implementation please see the original article by [60]. After the training is done, for every forward-pass (prediction), weights of the BNN are randomly drawn from  $y(\mathbf{w})$ . Assuming that  $y(\mathbf{w})$  is a normal distribution, the weights at every forward-pass are:

$$w_i = w_0^i + \varepsilon \sigma_i, \quad \varepsilon \sim \mathcal{N}(0, 1). \quad (16)$$

Consequently, if an input is passed through the BNN several times, the outputs will be slightly different each time and form their own distribution. The mean of these distributions are considered the outputs (predictions) and their standard deviation is the epistemic uncertainty. For the results presented here, we pass each input through the BNN 100 times ( $n_e = 100$  ensemble size of 100).

### 3.4.3. Dropout neural networks

Bayesian probability theory offers mathematically grounded tools to reason about model uncertainty, but these usually come with a prohibitive computational cost. It is perhaps surprising then that it is possible to cast recent deep learning tools as Bayesian models without changing either the models or the optimization. It can be shown that the use of dropout (and its variants) in NNs can be interpreted as a Bayesian approximation of a well known probabilistic model: the Gaussian process (GP) [27]. Dropout is used in many models in deep learning as a way to avoid over-fitting and this interpretation suggests that dropout approximately integrates over the model weights [61,62]. In other words, it is shown that a neural network with arbitrary depth and non-linearities, with dropout applied before every layer, is mathematically equivalent to an approximation to the probabilistic deep Gaussian process (marginalized over its covariance function parameters). It can also be shown that the dropout objective, in effect, minimizes the Kullback–Leibler divergence between an approximate distribution and the posterior of a deep Gaussian process (marginalized over its finite rank covariance function parameters) [63].

The approach we take in this article is that, before each layer of our  $8 \times [250]$  neural network model (except the first hidden layer), we add a dropout layer with dropout rate of  $r_D$ . The dropout layer sets neurons (input to the next layer) to zero randomly with probability of  $r_D$  at every *training step*, while scaling up the rest (non-zero neuron or inputs) by a factor of  $1/(1 - r_D)$  so that the sum over all neurons (inputs) remains unchanged. Typically, after the training is done, the dropout layers are switched off which makes the neural network to have deterministic output, however, here, we keep them on during prediction, as well. This means that during *every forward-pass* (prediction) some of the neurons in each hidden layer are randomly *switched off*, i.e., at every forward-pass the prediction is made by a slightly changed NN. Inference on a D-NN will predict different values when the same input is run multiple times. This means that the D-NN predicts a different output associated with each input every time it is passed through the D-NN. This leads to D-NN having a probabilistic output.

We show a schematic representation of the D-NN architecture in Fig. 11, and in particular emphasize the differences between the D-NN and BNN architectures: the BNN probability distribution controls the *weights*, whereas the Dropout probabilities control the *nodes* in each hidden layer.

In Fig. 12, we display some of the effects of  $r_D$  and  $n_e$  on the normalized residuals and epistemic uncertainty distributions. In the left subfigure, we show how changing  $r_D$  and  $n_e$  affects the distribution of uncertainties. On the right, we see how changing  $r_D$  and  $n_e$  affects the distribution of normalized residuals. We note that increasing  $r_D$  has complicated effects on the distribution of  $\sigma_e$ , but it generally increases  $\sigma_n$  and the range of  $\sigma_e$ . Additionally, we note that increasing  $r_D$  causes the distribution of normalized residuals to more closely approximate the standard Gaussian. Fig. 12 implies that increasing  $n_e$  from 5 to 50 does not have any significant

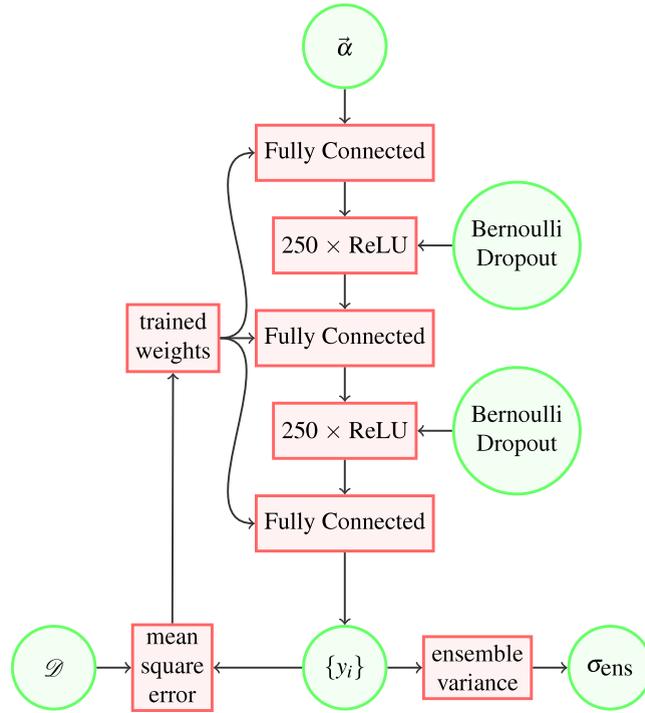


Fig. 11. Flowchart model of Dropout Neural Network (D-NN) with 2 hidden layers.

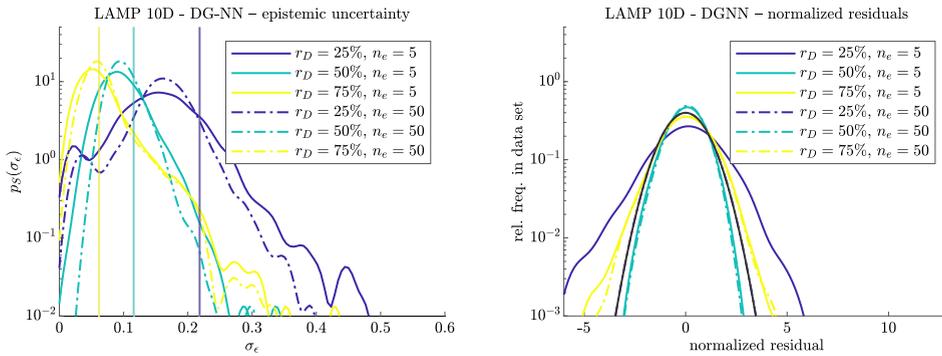


Fig. 12. Comparison of Gaussian D-NN (DG-NN) surrogate models for different dropout parameters,  $r_D$ , and ensemble size,  $n_e$ .  $N = 10, n_s = 2000$ . Additionally, the configuration of the neural network is as described before. **Left:** Distribution of epistemic uncertainties and aleatoric uncertainties (vertical lines). **Right:** Distribution of normalized residuals compared against a normal distribution (black solid curve).

effects on normalized residuals and epistemic uncertainty distribution. Additionally, Fig. 13 depicts the effects of ensemble size,  $n_e$ , and  $r_D$  on estimating  $\sigma_\epsilon$  both in value and distribution.

**Parameter Study.** In the previous section, we hinted at one of the parameters that directly affect the performance of Dropout Neural Networks (D-NNs) and Dropout Gaussian Neural Networks (DG-NNs), i.e., ensemble size ( $n_e$ ). We showed the convergence of  $\log p_S(\sigma_\epsilon^{n_e})$  and the mean squared error in  $\sigma_\epsilon$  vs. ensemble size ( $n_e$ ) - cf. Fig. 13. While it shows orders of magnitudes improvement as  $n_e$  increases, Fig. 12 shows that qualitatively, even an ensemble size of  $n_e = 5$  can accurately predict unseen data and estimate epistemic uncertainties compared to  $n_e = 50$ .

Additionally, figures 12 and 13 shows the effects of another parameter, dropout rate ( $r_D$ ), on DG-NN prediction of unseen data (implied in distribution of normalized residuals) and estimation of epistemic and aleatoric uncertainties. We can make the following conclusions, in terms of bias–variance trade-off, by examining Fig. 12:

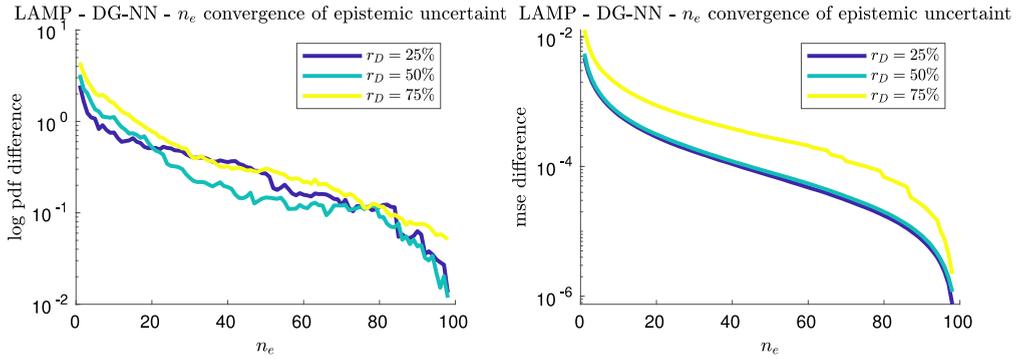


Fig. 13. Convergence curves of  $\sigma_\epsilon$  for different  $r_D$  as functions of ensemble size,  $n_e$ . **Left:** log pdf difference:  $\int |\log p_S(\sigma_\epsilon^{n_e}) - \log p_S(\sigma_\epsilon^{100})| d\sigma_\epsilon$ . **Right:** Mean of squared differences:  $\sum (\sigma_\epsilon^{n_e} - \sigma_\epsilon^{100})^2 / N_{samples}$ .

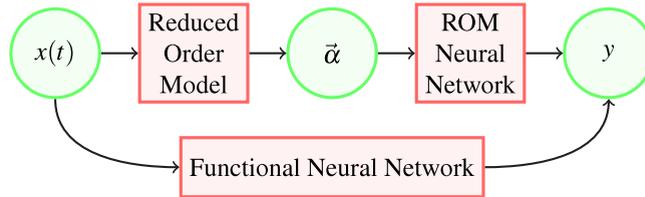


Fig. 14. Flowchart model for functional input.

- **Small dropout rate ( $r_D = 25\%$ ):** For small  $r_D$ , the model estimates low uncertainty values. In addition, the large tails of the normalized residuals distribution (compared to a normal distribution) corresponding to  $r_D = 25\%$  implies that the model is unable to generalize to previously unseen data. In summary, with small dropout rate the model tends to overfit, i.e., **low bias–high variance**.
- **Large dropout rate ( $r_D = 75\%$ ):** For large  $r_D$ , the model estimates much larger uncertainty values. Also, the narrow nature of the associated normalized residuals distribution implies that the model generalizes *too well* to unseen data, that is, **high bias–low variance**. The conclusion here is that large dropout rates create models that tend to underfit, i.e., **high bias–low variance**.
- **50% dropout rate ( $r_D = 50\%$ ):** We have observed that the model with  $r_D = 50\%$  not only estimates the uncertainties well but also generalizes well to unseen data as evident by its corresponding normalized residuals distribution. Our conclusion is that a DG-NN model with 50% dropout rate properly trains on the training data and generalizes well to unseen data, i.e., **medium bias–medium variance**.

### 3.5. Functional inputs

So far, we have considered the input to our machine learning models to be coefficient vectors for a reduced order model (ROM): wave time series for the LAMP example and initial condition for the MMT example. This choice is appropriate for GP, which struggle with higher dimensional input. However, neural network architectures are not so limited.

We may consider instead feeding a time series or spatial series input into a first fully connected layer of the neural network, after which the network architecture would be the same as before. In the LAMP case, this first layer would be a 1024 dimensional vector, representing the time series sampled uniformly on the interval  $[0, T]$ , where  $T = 40$  is the duration of the wave episodes considered in the LAMP simulation. In the MMT case, this first layer would also be a 1024 dimensional vector, the first 512 components of which correspond to the real part of the PDE initial conditions on the interval  $[0, 1]$ , and the second 512 components of which correspond to the imaginary parts.

In Fig. 14, we show a cartoon relationship between the neural network described previously, the ROM network, and the functional network. In the ROM network, we include a preprocessing step to transform the input into a low dimensional vector representation. Conversely, in the functional network, we feed the input directly into a fully connected layer of the network. In each case, the

**Table 3**  
Comparison of UQ capabilities for various ML techniques.

Architecture	Epistemic UQ?	Aleatoric UQ?
GP	Yes	Yes
NN	No	No
G-NN	No	Yes
E-NN	Yes	No
D-NN	Yes	No
B-NN	Yes	Yes
EG-NN	Yes	Yes
DG-NN	Yes	Yes

**Table 4**  
Selection of architecture details used in neural network models for LAMP presented below.

Architecture	Depth	Activation function	loss function	explicit regularization
E-NN	$8 \times [250]$	ReLU	mse	Yes
EG-NN	$8 \times [250]$	ReLU	NLL	Yes
B-NN	$8 \times [250]$	ReLU	-ELBo	No
DG-NN	$8 \times [250]$	ReLU	NLL	No

**Table 5**  
Selection of architecture details used in neural network models for MMT presented below.

Architecture	Depth	Activation function	loss function	explicit regularization
E-NN	$8 \times [250]$	ReLU	mse	Yes
D-NN	$8 \times [250]$	ReLU	mse	No

network returns a scalar output. This may be compared to the literature on neural operators such as DeepONet and Fourier Neural Operator [64–66], where both the input and the output side of the network are functions.

We make a special note that we keep the same network architecture—we have not introduced recurrent networks or convolutional networks. The functional input mode is compatible with all of the UQ techniques discussed above, but takes little advantage of the most recent developments in neural operator architecture. In Section 5.3, we present some comparisons of surrogate models using ROM input and functional input.

### 3.6. Combinations and summary

In Table 3, we summarize the uncertainty quantification ability of each machine learning technique discussed. In general, the G-NN final layer (included in B-NN) is required to estimate aleatoric uncertainty. Further, some kind of ensemble technique is required to estimate epistemic uncertainty. However, both B-NN and DG-NN provide an alternative ensemble technique, and DG-NN generally overperformed compared to E-NN.

Each of the large scale architectures we trialed also has a number of adjustable hyperparameters. We summarize our final selections in Tables 4 and 5. We note that the width and depth of our neural networks,  $8 \times [250]$ , is far into the overparameterization regime. Our choice of activation function, ReLU, is also a standard activation function. Some techniques, such as B-NN and EG-NN were found to be sensitive to choice of activation function, while others, like DG-NN, were not.

Additionally, we note that Dropout, and to a lesser extent the Variational Inference technique behind B-NN, are implicit regularization strategies. For E-NN and EG-NN, which cannot take advantage of this built in regularization, we used  $L_2$  regularization of weights.

## 4. Quality measures

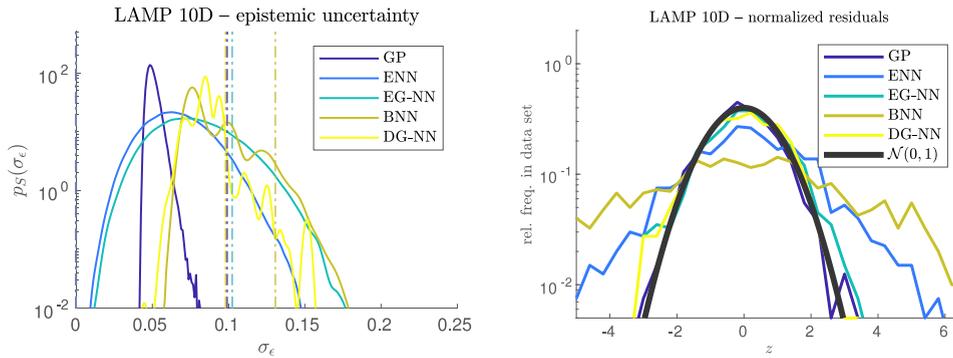
To measure the quality of each regression method we cannot simply rely on quantifying their regression accuracy, as this does not contain information for the uncertainty quantification part. To this end we adopt two separate measures.

### 4.1. Normalized residual distribution

The first quality metric we examine is the distribution of *normalized residuals* (NR), sometimes called z-scores. This is given by the histogram of the normalized error *computed on the validation set*:

$$z_i = \frac{y_i - \mu(\alpha_i)}{\sigma(\alpha_i)}. \quad (17)$$

The NR histograms measures the self-consistency of the regression scheme, i.e. how honest it is on accurately predicting the regions of the input space where the prediction uncertainty is high. It is not directly related to the quality of the regression, i.e. how



**Fig. 15.** Left: Distributions of epistemic uncertainty  $\sigma_\epsilon$ , for LAMP data and different ML techniques.  $N = 10, n_s = 2000$ . Right: Distributions (pdf) of normalized residuals  $z_i = \frac{y_i - \mu(\alpha_i)}{\sigma(\alpha_i)}$ .

accurate it is, but rather to whether the errors on the validation data are consistent in magnitude and location with the a priori estimated uncertainty. For accurately estimated uncertainties, the NR histogram should match a distribution with zero mean and unit variance.

When we present NR plots, we include an overlaid standard Gaussian distribution ( $\mathcal{N}(0, 1)$ ) to help with evaluation. We emphasize that a well calibrated UQ scheme will not, in general, match the *shape* of the standard Gaussian distribution, but will match the *mean* (i.e., 0) and the *variance* (i.e., 1). Generally, a NR histogram with variance greater than one, or generally heavier tails, is an indicator of over-fitting for the corresponding regression scheme, i.e. larger errors than what is expected by the estimated uncertainty. On the other hand, if the NR variance is less than one, we have higher estimated uncertainty than the actual errors indicate. This could be the case, for example, when the regression schemes operates as a smoother over the data, a behavior typically combined by a large uncertainty (aleatoric or epistemic). Alternately, this can occur from unconverged model fitting.

For each of our test cases, we generate the distribution of normalized residuals by separating the data into a training set and a validation set. For each point in the validation set, none of which have been seen by the model during training, we compute the residual (model error) and divide by the total model uncertainty at that point. We emphasize that the set of normalized residuals computed this way are **not** a distribution, because the samples from the validation set are not independent of the samples from the training set—in particular, the validation data are precisely the samples not selected for the training set! This proviso is more significant for the LAMP data, where computational costs limit the total number of simulations, than for the MMT data, where only a small fraction of the total data is used for training.

## 4.2. Uncertainty distribution

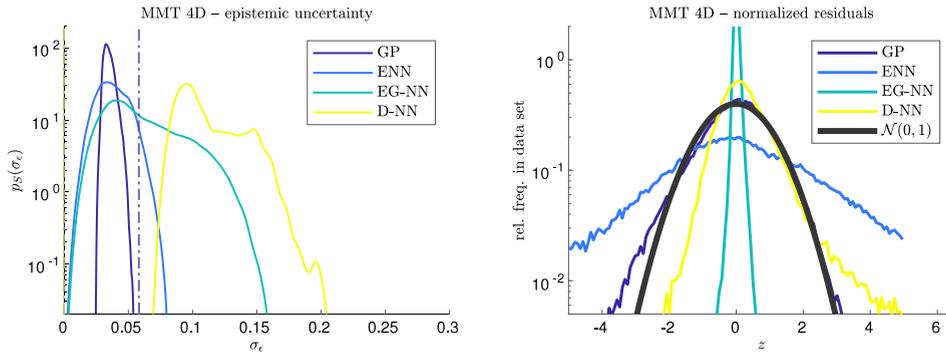
While NR quantifies the consistency of the estimated uncertainty with the model errors, it does not provide an overall measure of the accuracy of the regression method. To cover this aspect, we also adopt the uncertainty histogram, i.e. a histogram of all the estimated epistemic uncertainties,  $\sigma(\alpha_i)$ , over the validation set. For models that include an estimate of aleatoric uncertainty, we additionally include the estimated aleatoric uncertainty—for this work, assumed to be uniform over the input space. It is important to emphasize that the distribution of epistemic uncertainties and the aleatoric uncertainty measure how accurate each regression scheme ‘believes’ it is, through its own uncertainty quantification scheme.

Whether this uncertainty is consistent with the accuracy of the regression, it is measured by the NR distribution. To this end, these two measures should be used in combination to quantify the quality of any regression schemes that aims to quantify uncertainty. A high quality regression scheme is associated with small values of the uncertainty, and an NR distribution with unit variance. The distribution of uncertainties is especially useful for active sampling schemes, where (large) outlier uncertainties represent important gaps in experimental data.

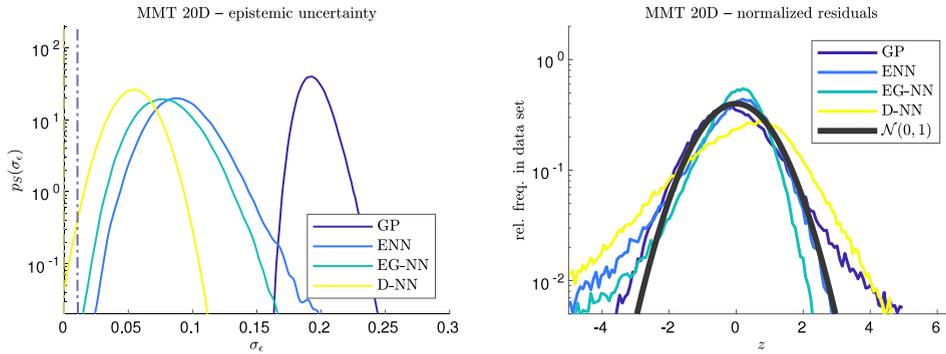
## 5. Results

### 5.1. LAMP dataset

First, we consider UQ for the LAMP dataset. In Fig. 15, we display the distribution of uncertainties, as well as the normalized residuals for  $N = 10$  dimensional LAMP data. In the left column, we display the pdf of epistemic uncertainties for each ML technique, along with a vertical line for the estimate aleatoric uncertainty,  $\sigma_n$ . On the right column, we plot the normalized residuals, given by Eq. (17). We note that GP, EG-NN, and D-GNN all estimate a very similar value for  $\sigma_n$ . It is important to emphasize that despite the relatively high dimensionality of the data set its inherently low dimensional structure allows for GP to work very well in terms of minimizing epistemic uncertainty. The NN based methods have comparable performance in terms of epistemic uncertainty. In terms of UQ properties, i.e. capability to model own errors, besides ENN and BNN both of which underestimate model error, the UQ for GP, EG-NN, and DG-NN are all well calibrated.



**Fig. 16.** Left: Distributions of epistemic uncertainty  $\sigma_\epsilon$ , for MMT data and different ML techniques.  $N = 4, n_s = 2000$ . Right: Distributions of normalized residuals  $z_i = \frac{y_i - \mu(\alpha_i)}{\sigma(\alpha_i)}$ .



**Fig. 17.** Left: Distributions of epistemic uncertainty  $\sigma_\epsilon$ , for MMT data and different ML techniques.  $N = 20, n_s = 10000$ . Right: Distributions of normalized residuals  $z_i = \frac{y_i - \mu(\alpha_i)}{\sigma(\alpha_i)}$ .

## 5.2. MMT dataset

In Figs. 16 and 17, we compare the UQ performance for selected models on MMT datasets of dimensionality  $N = 4$  and  $N = 20$ . We chose not to include BNN in this comparison, because it performed poorly in the LAMP trial and preliminary investigation for MMT confirmed this performance. We further note that we include D-NN instead of DG-NN for this dataset. This choice was made to capitalize on our a priori knowledge that the MMT dataset does not include aleatoric uncertainty, combined with our experience that the G-NN final layer complicates training when the true  $\sigma_n \approx 0$ .

We observe that GP performs fairly well compared to NN architectures for  $N = 4$  dimensions in terms of estimated uncertainty, but its performance deteriorates for  $N = 20$ . This is in line with the intuition that GP scales poorly to higher dimensional input space. Note that while the GP model uncertainty/errors are large absolutely in  $N = 20$  dimensions, the UQ is still well calibrated according to the normalized residual plot. On the other hand, for the higher dimensional case, the NN based architectures perform very well in minimizing epistemic uncertainties with the D-NN having the best performance. Moreover, they have good UQ-error consistency properties with the D-NN doing slightly worse compared with the other methods, i.e. its UQ underestimates its large errors.

## 5.3. Functional inputs

In addition to comparisons between different surrogate models, we also compared the ENN and EG-NN techniques using ROM inputs and functional inputs as described above in Section 3.5. In Figs. 18 and 19, we present comparisons for both LAMP and MMT data sets. In Fig. 18, we see that functional inputs lead to significantly better predictions in the LAMP problem with a 10 dimensional space of wave episodes. This is evidenced by a distribution of estimated model errors (left plot) much smaller for functional inputs (red), without concomitant overfitting (right plot). In fact, the observed *underfitting* suggests that network expressivity is a limiting factor.

In Fig. 19, we see that for the MMT data, the functional inputs lead to smaller improvements. While for 4D input data, the different input methods has no clear difference, but 20D input data, the functional input mode shows a clear improvement, if not as dramatic as displayed by the LAMP data. Why does functional input improve the results from LAMP data more than for MMT data? It seems that functional neural networks have a stronger advantage when the underlying structure of the data set has lower

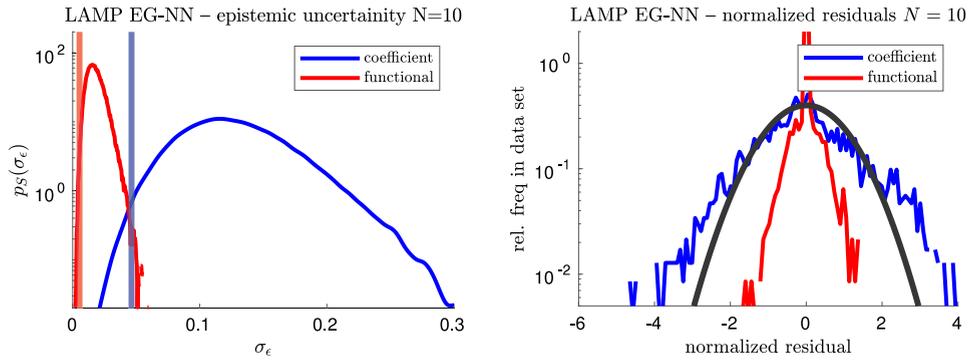


Fig. 18. Comparison of EG-NN with ROM inputs and with functional inputs for 10D LAMP data,  $n_s = 2000$ . **Left:** Comparison of uncertainties. **Right:** Comparison of normalized residuals.

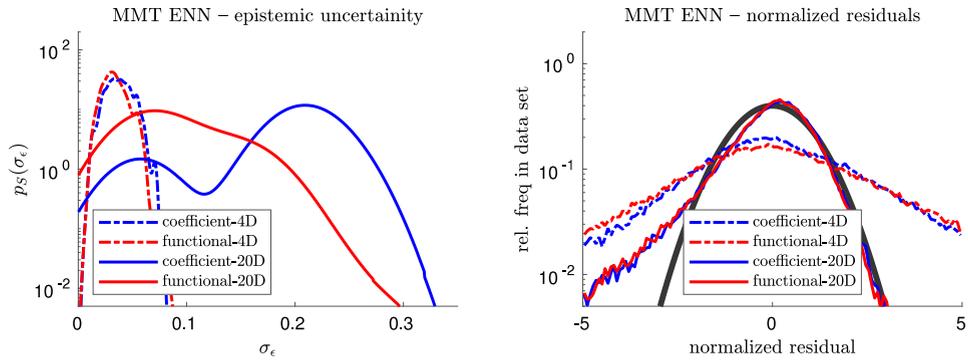


Fig. 19. Comparison of ENN with ROM inputs and with functional inputs for 4D and 20D MMT data ( $n_s = 462$  and  $n_s = 10000$ ). **Left:** Comparison of uncertainties. **Right:** Comparison of normalized residuals.

dimensionality (Section 2.4). Thus, the LAMP wave episodes, which inherit more structure from their KL basis, are able to better take advantage of the functional input improvement, compared with the MMT data set.

### 6. Conclusions

In this paper, we have reviewed a number of existing machine learning methods for combining neural network surrogate models with UQ for estimating both aleatoric and epistemic uncertainties in data sets associated with complex dynamical systems. First, we presented the G-NN model, which allows any neural network architecture, including traditional deterministic feed-forward networks, to estimate aleatoric uncertainty associated with a training data set. Next, we presented three implementations of the ensemble technique for estimating epistemic uncertainty: ENN, BNN and D-NN. Each of these ensemble techniques may be combined with the G-NN final layer, allowing for a neural network surrogate model to estimate both aleatoric and epistemic uncertainties. Finally, motivated by recent work on neural operators, we described how to simply change the network architecture to handle functional inputs.

We compared each of these techniques, along with a GP model for reference, on two different real datasets obtained by complex dynamical systems: one with an important intrinsic aleatoric component (LAMP) and the other without (MMT). We compared performance with metrics tailored specifically to the question of UQ: normalized residuals and predicted uncertainties. The distribution of normalized residuals from a hold-out test set provide an estimate of how well calibrated the model UQ is and may indicate presence of underfitting or overfitting. The distribution of the model uncertainties indicated how uncertain the model is generally.

In the data set with intrinsically low dimensional structure (LAMP) as well as the low-dimensional version of the MMT data set, GP always performs the best: it has one of the lowest epistemic uncertainties, as it is also self consistent, i.e. errors in the validation test are correlated well with the estimated uncertainty in terms of magnitude and location. On the other hand, for the higher dimensional MMT data set, GP has important limitations, as expected. In this case, ensemble methods, in particular D-NN or ENN perform consistently better. These prediction ensemble methods also have excellent performance in the normalized residual test. We found that BNN models underperformed in our investigation, especially compared to D-NN, their most similar competitor. Further, there were also significant difficulties in training BNN models. Regarding the estimation of the aleatoric uncertainty, we

found that the addition of adding a Gaussian layer in ensemble methods (EG-ENN, DG-NN) results in accurate estimation of the aleatoric uncertainty, i.e. comparable with GP.

Finally, we examined the use of ROM (coefficient) inputs compared to functional inputs. In the case of the LAMP data, functional inputs performed significantly better than coefficient inputs with minimal tuning. However, in the case of MMT data functional inputs provided only marginal improvements after significant optimization. We conjecture that this difference is due to differences in the inherent dimensionality of the two data sets.

Overall, the GP is the recommended method for data sets that ‘live’ in low dimensional spaces or have a low-dimensional structure. For intrinsically higher-dimensional data sets we recommend the dropout technique for constructing ensemble variance estimates, with a balanced dropout fraction of  $r_D = 0.5$ . However, we find that the ensemble networks may also be useful, especially in situations where complicated model architecture may make the dropout technique challenging to implement. The D-NN architecture is especially appropriate for active learning applications whose dimensionality is too great for effective GP modeling.

### CRedit authorship contribution statement

**Stephen Guth:** Conceptualization, Formal analysis, Investigation, Methodology, Validation, Visualization, Writing – original draft, Writing – review & editing. **Alireza Mojahed:** Data curation, Formal analysis, Investigation, Writing – original draft. **Themistoklis P. Sapsis:** Conceptualization, Formal analysis, Funding acquisition, Resources, 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.

### Acknowledgments

We acknowledge support from the AFOSR, United States MURI grant no. FA9550-21-1-0058, the DARPA, United States grant no. HR00112290029, and the ONR, United States grant no. N00014-21-1-2357. We are grateful to Dr. Ethan Pickering for useful discussions.

### Appendix

#### A.1. LAMP data $N=2$ — visualization of surrogate and epistemic uncertainty

A reasonable question for comparison across these techniques is, where is the epistemic uncertainty maximized? This is important for uncertainty quantification generally, but especially for active learning with uncertainty sampling based techniques.

In Fig. 20, left column, we show the surrogate mean  $\mu(\vec{\alpha})$  in the restricted  $N = 2$  dimensional case, with  $n_s = 25$  training points (c.f. Fig. 4). We show this image for background before we display the spatial distribution of epistemic uncertainties in the right column of Fig. 20. In particular, note the belt of maximum gradient passing along an approximately 45 degree angle near the origin. In Fig. 20, right column, we show epistemic uncertainty as a function of  $(\alpha_1, \alpha_2)$  for a 2D dataset with  $n_s = 25$ . The low dimension and small training set allows one to compare the location of the training data to the location of regions of high epistemic uncertainty.

We note that, while generally we hold up GP as the gold standard of uncertainty quantification, for the question of the spatial distribution of epistemic uncertainty we give GP no special deference. The kernel matrix math for  $\sigma(\vec{\alpha})$  is independent of the  $q_1$  component of the training data. As a result, GP uncertainty estimation largely ignores the existence of regions with large gradient, except to the extent they are captured during the hyperparameter optimization.

Additionally, the spatial distribution of the epistemic uncertainties (and not the absolute values) is of particular important to active sampling. In uncertainty sampling based active sampling methods, the activation function is a product of the epistemic variance, possibly along with some weight factor:

$$u_{\text{US}}(\vec{\alpha}) = w(\alpha)\sigma_\epsilon^2(\vec{\alpha}). \quad (18)$$

In this application, scaling all the uncertainties by a constant factor leaves the relative ordering of the acquisition function unchanged.

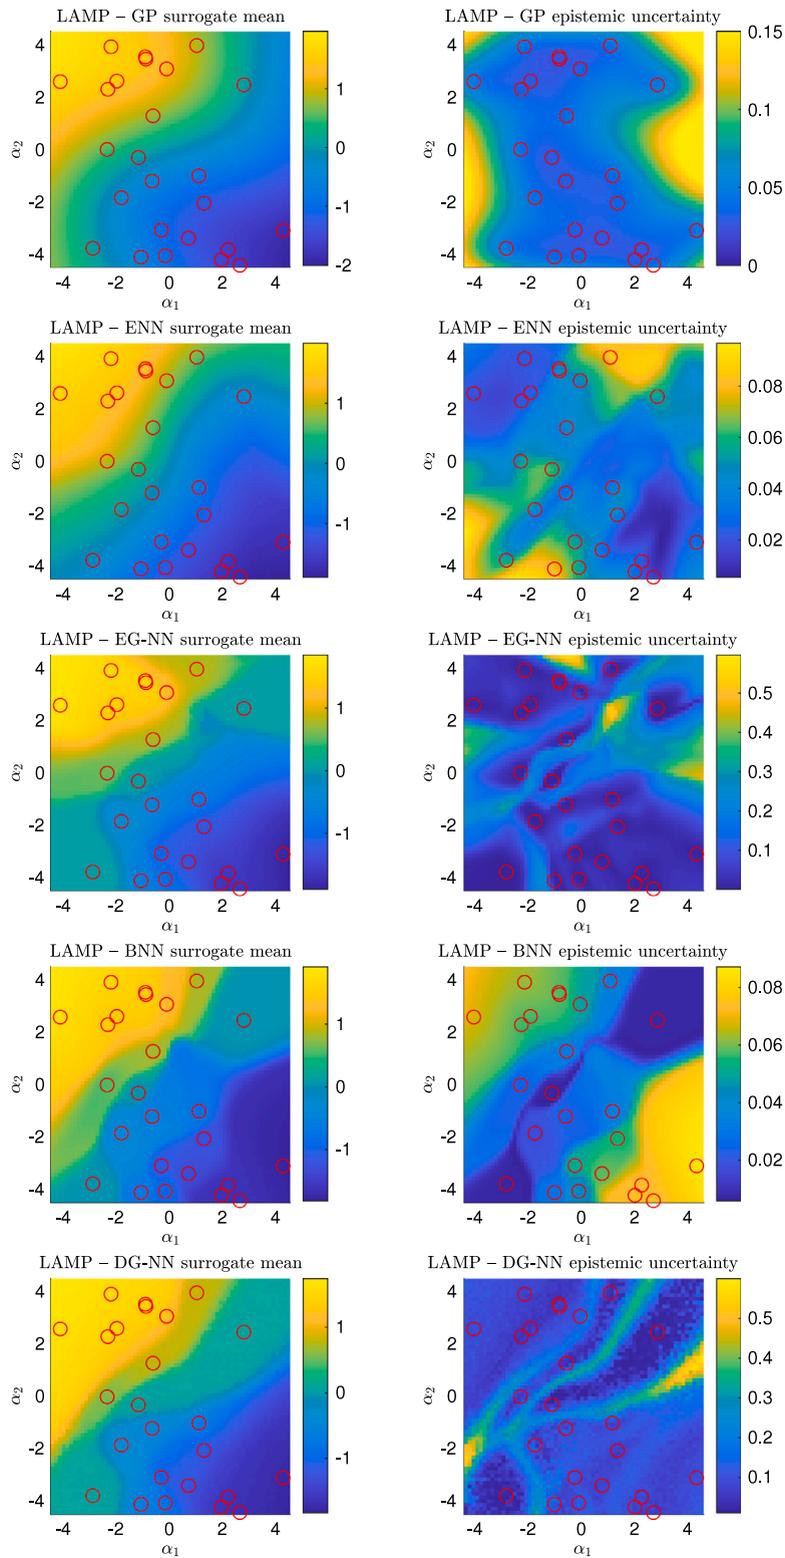


Fig. 20. Spatial distributions of the (Left) surrogate mean  $\mu(\bar{\alpha})$  and (Right) epistemic uncertainty  $\sigma_e(\bar{\alpha})$ , for LAMP data and different ML techniques.  $N = 2, n_s = 25$ .

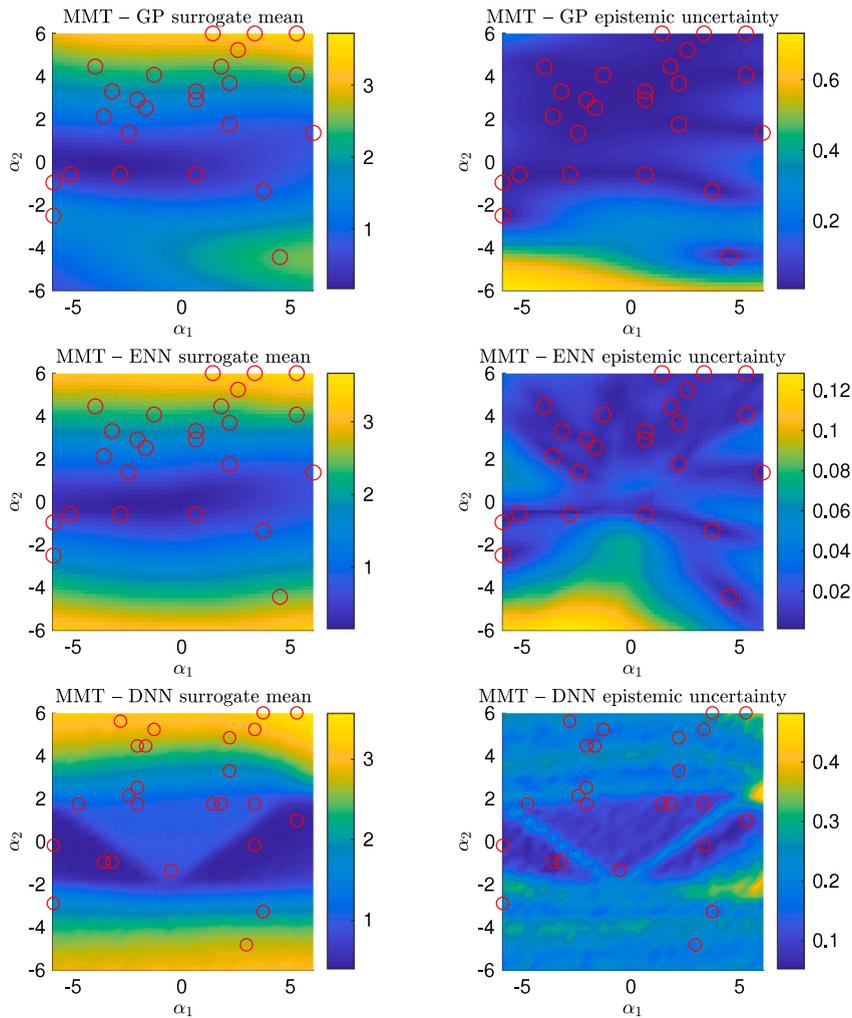


Fig. 21. Spatial distributions of the (Left) surrogate mean  $\mu(\bar{\alpha})$  and (Right) epistemic uncertainty  $\sigma_e(\bar{\alpha})$ , for MMT data and different ML techniques.  $N = 2$ ,  $n_s = 25$ .

#### A.2. MMT data $N=2$ – visualization of surrogate and epistemic uncertainty

For completeness, in Fig. 21 we present the corresponding epistemic plots for the MMT dataset for the GP, ENN and D-NN surrogate models. The MMT system displays a simpler character than the LAMP system—in the 2D input restriction, the output is only weakly dependent on  $\alpha_1$ .

#### References

- [1] J.O. Berger, L.A. Smith, On the statistical formalism of uncertainty quantification, *Annu. Rev. Stat. Appl.* 6 (1) (2019) 433–460.
- [2] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U.R. Acharya, V. Makarenkov, S. Nahavandi, A review of uncertainty quantification in deep learning: Techniques, applications and challenges, *Inf. Fusion* 76 (2021) 243–297.
- [3] N. Kennamer, A. Ihler, D. Kirkby, Empirical study of MC-dropout in various astronomical observing conditions, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019.
- [4] G. Wang, W. Li, M. Aertsen, J. Deprest, S. Ourselin, T. Vercauteren, Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks, *Neurocomputing* 338 (2019) 34–45.
- [5] A. Akbari, R. Jafari, Personalizing activity recognition models through quantifying different types of uncertainty using wearable sensors, *IEEE Trans. Biomed. Eng.* 67 (9) (2020) 2530–2541.
- [6] A. Charalampopoulos, T. Sapsis, Uncertainty quantification of turbulent systems via physically consistent and data-informed reduced-order models, *Phys. Fluids* 34 (7) (2022) 075120.
- [7] C.E. Rasmussen, C.K.I. Williams, *Gaussian Processes for Machine Learning*, in: *Adaptive Computation and Machine Learning*, MIT Press, Max-Planck-Gesellschaft, Biologische Kybernetik, Cambridge, MA, USA, 2006, p. 248.
- [8] X. Gong, Y. Pan, Sequential Bayesian experimental design for estimation of extreme-event probability in stochastic input-to-response systems, *Comput. Methods Appl. Mech. Engrg.* 395 (2022) 114979.

- [9] S. Guth, T.P. Sapsis, Wave episode based Gaussian process regression for extreme event statistics in ship dynamics: Between the Scylla of Karhunen–Loève convergence and the charybdis of transient features, *Ocean Eng.* 266 (2022) 112633.
- [10] S. Guth, E. Katsidomiotaki, T. Sapsis, Statistical modeling of fully nonlinear hydrodynamic loads on offshore wind turbine foundations using wave episodes and targeted CFD simulations through active sampling, 2023, <http://dx.doi.org/10.22541/au.168537006.67360210/v1>, Authorea Preprints.
- [11] M.A. Mohamad, T.P. Sapsis, Sequential sampling strategy for extreme event statistics in nonlinear dynamical systems, *Proc. Natl. Acad. Sci.* 115 (44) (2018) 11138–11143.
- [12] A. Blanchard, T. Sapsis, Bayesian optimization with output-weighted optimal sampling, *J. Comput. Phys.* 425 (2020) 109901.
- [13] A. Blanchard, T. Sapsis, Output-weighted optimal sampling for Bayesian experimental design and uncertainty quantification, *SIAM/ASA J. Uncertain. Quantif.* 9 (2) (2021) 564–592.
- [14] Y. Yang, A. Blanchard, T. Sapsis, P. Perdikaris, Output-weighted sampling for multi-armed bandits with extreme payoffs, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* 478 (2260) (2022) 20210781.
- [15] P. Perdikaris, D. Venturi, J.O. Royset, G.E. Karniadakis, Multi-fidelity modelling via recursive co-kriging and Gaussian–Markov random fields, *Proc. Math. Phys. Eng. Sci.* (2015).
- [16] H. Babaei, C. Bastidas, M. Defilippo, C. Chrysostomidis, G. Karniadakis, A multi-fidelity framework and uncertainty quantification for sea surface temperature in the Massachusetts and cape cod bays, *Earth Space Sci.* 7 (2020).
- [17] B. Champenois, T. Sapsis, Machine learning framework for the real-time reconstruction of regional 4D ocean temperature fields from historical reanalysis data and real-time satellite and buoy surface measurements, *Physica D: Nonlinear Phenomena* 459 (2024) 134026.
- [18] E. Pickering, S. Guth, G.E. Karniadakis, T. Sapsis, Discovering and forecasting extreme events via active learning in neural operators, *Nat. Comput. Sci.* 2 (2022) 823–833.
- [19] A.F. Psaros, X. Meng, Z. Zou, L. Guo, G.E. Karniadakis, Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons, *J. Comput. Phys.* 477 (2023) 111902.
- [20] E. Pickering, T. Sapsis, Information FOMO: The unhealthy fear of missing out on information. a method for removing misleading data for healthier models, 2022, arXiv preprint arXiv:2208.13080.
- [21] M.B. Tomczak, R.E. Turner, Marginal likelihood gradient for Bayesian neural networks, in: *Third Symposium on Advances in Approximate Bayesian Inference*, 2021.
- [22] P. Schwöbel, M. Jørgensen, S.W. Ober, M. Van Der Wilk, Last layer marginal likelihood for invariance learning, in: G. Camps-Valls, F.J.R. Ruiz, I. Valera (Eds.), *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, in: *Proceedings of Machine Learning Research*, 151, PMLR, 2022, pp. 3542–3555.
- [23] J. Hensman, A. Matthews, Z. Ghahramani, Scalable Variational Gaussian Process Classification, in: G. Lebanon, S.V.N. Vishwanathan (Eds.), *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, in: *Proceedings of Machine Learning Research*, vol.38, PMLR, San Diego, California, USA, 2015, pp. 351–360.
- [24] A. Wilson, H. Nickisch, Kernel interpolation for scalable structured Gaussian processes (KISS-GP), in: F. Bach, D. Blei (Eds.), *Proceedings of the 32nd International Conference on Machine Learning*, in: *Proceedings of Machine Learning Research*, vol.37, PMLR, Lille, France, 2015, pp. 1775–1784.
- [25] D.R. Burt, C.E. Rasmussen, M. van der Wilk, Convergence of sparse variational inference in Gaussian processes regression, *J. Mach. Learn. Res.* 21 (131) (2020) 1–63.
- [26] S. Ober, L. Aitchison, Global inducing point variational posteriors for Bayesian neural networks and deep Gaussian processes, in: *ICML*, 2021.
- [27] A. Damianou, N.D. Lawrence, Deep Gaussian processes, in: C.M. Carvalho, P. Ravikumar (Eds.), *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, in: *Proceedings of Machine Learning Research*, vol.31, PMLR, Scottsdale, Arizona, USA, 2013, pp. 207–215.
- [28] J. Bradshaw, A. Matthews, Z. Ghahramani, Adversarial examples, uncertainty, and transfer testing robustness in Gaussian process hybrid deep networks, 2017.
- [29] J. van Amersfoort, L. Smith, A. Jesson, O. Key, Y. Gal, Improving deterministic uncertainty estimation in deep learning for classification and regression, 2021, *CoRR* abs/2102.11409, arXiv:2102.11409.
- [30] S.W. Ober, C.E. Rasmussen, M. van der Wilk, The promises and pitfalls of deep kernel learning, in: C. de Campos, M.H. Maathuis (Eds.), *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, in: *Proceedings of Machine Learning Research*, vol.161, PMLR, 2021, pp. 1206–1216.
- [31] Z. Zou, X. Meng, A.F. Psaros, G.E. Karniadakis, NeuralUQ: A comprehensive library for uncertainty quantification in neural differential equations and operators, 2022, arXiv preprint arXiv:2208.11866.
- [32] E. Hüllermeier, W. Waegeman, Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods, *Mach. Learn.* 110 (2021).
- [33] Y. Shin, V. Belenky, W. Lin, K. Weems, A. Engle, K. McTaggart, J. Falzarano, B. Hutchison, M. Gerigk, S. Grochowalski, Nonlinear time domain simulation technology for seakeeping and wave-load analysis for modern ship design, *Trans. - Soc. Naval Archit. Mar. Eng.* 111 (2003) 557–583.
- [34] W.-M. Lin, S. Zhang, K. Weems, P. Jones, M. Meinhold, B. Metcalf, A.M. Powers, Numerical simulation and validation study of wetdeck slamming on high speed catamaran, in: *International Conference on Numerical Ship Hydrodynamics [9th]*, 2007.
- [35] W.-M. Lin, M. Collette, D.R. Lavis, S.D. Jessup, J. Kuhn, Recent hydrodynamic tool development and validation for motions and slam loads on ocean-going high-speed vessels, 2007.
- [36] W.-M. Lin, S. Zhang, K.M. Weems, Numerical simulations of surface effect ship in waves, in: *Proceedings of the 2010 Conference on Grand Challenges in Modeling & Simulation*, GCMS '10, Society for Modeling & Simulation International, Vista, CA, 2010, pp. 414–421.
- [37] K. Hasselmann, T. Barnett, E. Bouws, H. Carlson, D. Cartwright, K. Enke, J. Ewing, H. Gienapp, D. Hasselmann, P. Kruseman, A. Meerburg, P. Muller, D. Olbers, K. Richter, W. Sell, H. Walden, Measurements of wind-wave growth and swell decay during the joint north sea wave project (JONSWAP), *Deut. Hydrogr. Z.* 8 (1973) 1–95.
- [38] A.J. Majda, D.W. McLaughlin, E.G. Tabak, A one-dimensional model for dispersive wave turbulence, *J. Nonlinear Sci.* 7 (1997) 9–44.
- [39] D. Cai, A. Majda, D. McLaughlin, E. Tabak, Dispersive wave turbulence in one dimension, *Physica D* 152–153 (2001) 551–572.
- [40] V. Zakharov, P. Guyenne, A. Pushkarev, F. Dias, Wave turbulence in one-dimensional models, *Physica D* 152–153 (2001) 573–619, *Advances in Nonlinear Mathematics and Science: A Special Issue to Honor Vladimir Zakharov*.
- [41] V. Zakharov, F. Dias, A. Pushkarev, One-dimensional wave turbulence, *Phys. Rep.* 398 (1) (2004) 1–65.
- [42] B. Rumpf, L. Biven, Weak turbulence and collapses in the Majda–McLaughlin–Tabak equation: Fluxes in wavenumber and in amplitude space, *Physica D* (204) (2005) 188–203.
- [43] T.P.S. Will Cousins, Quantification and prediction of extreme events in a one-dimensional nonlinear dispersive wave model, *Physica D* (280–281) (2014) 48–58.
- [44] S. Guth, T. Sapsis, Machine learning predictors of extreme events occurring in complex dynamical systems, *Entropy* 21 (925) (2019).
- [45] T.B. Benjamin, J.E. Feir, The disintegration of wave trains of deep water, *J. Fluid Mech.* 27 (1967) 417–430.
- [46] V.E. Zakharov, Stability of periodic waves of finite amplitude on the surface of a deep fluid, *J. Appl. Mech. Tech. Phys.* 9 (2) (1972) 190–194.
- [47] M. Onorato, A.R. Osborne, M. Serio, S. Bertone, Freak waves in random oceanic sea states, *Phys. Rev. Lett.* 86 (2001) 5831–5834.
- [48] K.B. Dysthe, Note on a Modification to the Nonlinear Schrödinger Equation for Application to Deep Water Waves, *Proc. R. Soc. Lond. Ser. A* 369 (1736) (1979) 105–114.

- [49] D.J. Korteweg, G. de Vries, XLI. On the change of form of long waves advancing in a rectangular canal, and on a new type of long stationary waves, *Lond. Edinburgh Dublin Philos. Mag. J. Sci.* 39 (240) (1895) 422–443.
- [50] A. Sergeeva, E. Pelinovsky, T. Talipova, Nonlinear random wave field in shallow water: Variable Korteweg-de Vries framework, *Nat. Hazards Earth Syst. Sci.* 11 (2) (2011) 323–330.
- [51] P.D. Scavounos, Karhunen–Loève representation of stochastic ocean waves, *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* 468 (2145) (2012) 2574–2594.
- [52] M.G. Genton, Classes of kernels for machine learning: A statistics perspective, *J. Mach. Learn. Res.* 2 (2002) 299–312.
- [53] S. Shalev-Shwartz, S. Ben-David, *Understanding Machine Learning - from Theory to Algorithms*, Cambridge University Press, 2014, pp. I–XVI, 1–397.
- [54] K.P. Murphy, *Probabilistic Machine Learning: An Introduction*, MIT Press, 2022.
- [55] R. Calandra, J. Peters, C.E. Rasmussen, M.P. Deisenroth, Manifold Gaussian processes for regression, in: 2016 International Joint Conference on Neural Networks, IJCNN, 2016, pp. 3338–3345.
- [56] J. Van Amersfoort, L. Smith, Y.W. Teh, Y. Gal, Uncertainty estimation using a single deep deterministic neural network, in: H.D. III, A. Singh (Eds.), *Proceedings of the 37th International Conference on Machine Learning*, in: *Proceedings of Machine Learning Research*, vol.119, PMLR, 2020, pp. 9690–9700.
- [57] D.M. Blei, A. Kucukelbir, J.D. McAuliffe, Variational inference: A review for statisticians, *J. Amer. Statist. Assoc.* 112 (518) (2017) 859–877.
- [58] L.V. Jospin, H. Laga, F. Boussaid, W. Buntine, M. Bennamoun, Hands-on Bayesian neural networks—A tutorial for deep learning users, *IEEE Comput. Intell. Mag.* 17 (2) (2022) 29–48.
- [59] E. Goan, C. Fookes, Bayesian neural networks: An introduction and survey, in: K.L. Mengersen, P. Pudlo, C.P. Robert (Eds.), *Case Studies in Applied Bayesian Data Science*, CIRM Jean-Morlet Chair, Fall 2018, Springer International Publishing, Cham, 2020, pp. 45–87.
- [60] C. Blundell, J. Cornebise, K. Kavukcuoglu, D. Wierstra, Weight uncertainty in neural network, in: *International Conference on Machine Learning*, PMLR, 2015, pp. 1613–1622.
- [61] G.E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, R.R. Salakhutdinov, Improving neural networks by preventing co-adaptation of feature detectors, 2012, *arXiv preprint arXiv:1207.0580*.
- [62] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* 15 (56) (2014) 1929–1958.
- [63] Y. Gal, Z. Ghahramani, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning, in: *International Conference on Machine Learning*, PMLR, 2016, pp. 1050–1059.
- [64] K. Bhattacharya, B. Hosseini, N.B. Kovachki, A.M. Stuart, Model Reduction And Neural Networks For Parametric PDEs, *SMAI J. Comput. Math.* 7 (2021) 121–157.
- [65] N. Kovachki, S. Lanthaler, S. Mishra, On universal approximation and error bounds for Fourier neural operators, *J. Mach. Learn. Res.* 22 (2021) Art–No.
- [66] L. Lu, X. Meng, S. Cai, Z. Mao, S. Goswami, Z. Zhang, G.E. Karniadakis, A comprehensive and fair comparison of two neural operators (with practical extensions) based on FAIR data, *Comput. Methods Appl. Mech. Engrg.* 393 (2022) 114778.