Contents lists available at ScienceDirect

Theoretical & Applied Mechanics Letters

journal homepage: www.elsevier.com/locate/taml

Letter

Physics-constrained bayesian neural network for fluid flow reconstruction with sparse and noisy data

Luning Sun, Jian-Xun Wang*

ᵃ *Department of Aerospace and Mechanical Engineering, University of Notre Dame, Notre Dame, IN 46556, USA*
ᵇ *Center for Informatics and Computational Science, University of Notre Dame, Notre Dame, IN 46556, USA*

A B S T R A C T

In many applications, flow measurements are usually sparse and possibly noisy. The reconstruction of a high-resolution flow field from limited and imperfect flow information is significant yet challenging. In this work, we propose an innovative physics-constrained Bayesian deep learning approach to reconstruct flow fields from sparse, noisy velocity data, where equation-based constraints are imposed through the likelihood function and uncertainty of the reconstructed flow can be estimated. Specifically, a Bayesian deep neural network is trained on sparse measurement data to capture the flow field. In the meantime, the violation of physical laws will be penalized on a large number of spatiotemporal points where measurements are not available. A non-parametric variational inference approach is applied to enable efficient physics-constrained Bayesian learning. Several test cases on idealized vascular flows with synthetic measurement data are studied to demonstrate the merit of the proposed method.

Reconstruction of a flow field from limited and noisy measurements is of great significance yet challenging in many engineering applications. For example, the rapid development in flow magnetic resonance (MR) imaging techniques enables noninvasive assessment of hemodynamic information for cardiovascular research and healthcare [1]. However, the resolution and signal-to-noise ratio (SNR) of MR images still remain the limiting factors for clinical applications [2]. Similar scenarios can also be found in monitoring wind farms or other aerodynamic systems, where measurement sensors (e.g., lidar) are usually placed at sparse locations and thus the collected data are also sparse and noisy [3].

Because of its wide range of applications, full-field reconstruction of sparse, noisy flow data has become an active research area and received a great deal of attention. In order to compensate for the incompleteness and sparsity of the gappy data, additional information is required, which can be obtained either from an offline flow database or a physics-based model. Based on what type of information is incorporated, the existing flow reconstruction (i.e., superresolution) methods can be organized into two groups. (i) When large offline full-field flow data sets are available, the coherent structures and correlation features of the fluid flow can be extracted, which will be utilized to reconstruct the high-resolution flow fields from sparse online data. Proper orthogonal decomposition (POD) [4] or dynamic mode decomposition (DMD) [5, 6] are commonly used for flow feature extraction. For instance, gappy POD has been applied for steady and unsteady flow-field reconstruction in various applications [7–12]. To overcome the linearity limitations of POD and DMD, deep learning based approaches (e.g., autoencoder neural networks) have been recently developed to extract nonlinear latent representations of the flow field from massive offline data [13]. As an alternative, sparsity-promoting representation techniques, e.g., compressed sensing, have also been demonstrated to be able to achieve the same goal more robustly when data is noisy [14, 15]. All the algorithms described above rely on a large number of flow datasets for offline "training", which might not be available in many cases. (ii) Instead of learning from the offline database, the other type of flow reconstruction methods takes advantage of a physics-based model, e.g., computational

---

* Corresponding author.
E-mail address: jwang33@nd.edu (J.X. Wang).

fluid dynamics (CFD) model, which is able to provide full-field flow predictions. The sparse measurement data can be fused into the model-based predictions using data assimilation (DA) techniques, e.g., ensemble Kalman filter, particle filters, or variational based DA techniques [16–21]. Nonetheless, physics-based simulations are time-consuming in general while the DA process usually involves numerous model evaluations, which could be computationally prohibitive.

The recent advances of deep learning techniques for image superresolution [22, 23] open up new avenues for developing efficient algorithms of flow reconstruction from limited measurements. For example, neural networks (NN) have been used to learn POD coefficients [24] or directly capture an end-to-end mapping between the sparse measurements and the high-resolution flow field [25]. However, the success of these deep learning models is mostly dependent on a sufficient amount of offline training data, which, as mentioned above, are inaccessible in many applications, e.g., superresolution for flow MR imaging. To alleviate data sparsity, a physics-constrained deep learning strategy has been proposed [26–29], where physical laws of a system (e.g., Navier–Stokes equations in fluid mechanics) are leveraged to constrain the training process. Recently, this idea has attracted increasing attention and its merits have been demonstrated in solving a number of forward and inverse problems governed by classic partial differential equations (PDEs). Notably, the physics-informed neural networks (PINN) proposed by Rassi et al. [26] were applied to reconstruct a flow field by assimilating scalar concentration data of a flow field [30]. Sun et al. [28] developed a PINN-based fluid surrogate model with encoded boundary conditions and demonstrated that the flow solutions of parametric Navier–Stokes equations can be learned without using any labeled training data. Although the physics-constrained deep learning shows great promise for flow reconstruction of limited data, the measurement noise associated with the data and model-form uncertainties due to model inadequacy cannot be considered since the classic deep learning models are usually formulated in a deterministic way. Researchers have recently started to explore the uncertainty quantification (UQ) analysis of physics-constrained deep learning by using arbitrary polynomial chaos [31] and variational inference [27, 32, 33].

In this work, a physics-constrained Bayesian neural network (PC-BNN) is proposed for flow field reconstruction from sparse and noisy measurements. In contrast to previous works, the equation-constrained training is formulated in a Bayesian manner, where the posterior distribution of the NN weights will be obtained based on the likelihood function, which is defined by the uncertainty from both measurement noise and model inadequacy. Specifically, the confidence of the physical/physiological constraints is modeled in a probabilistic way, being combined with data uncertainty to form the likelihood function [34]. A nonparametric variation inference algorithm, Stein variation gradient decent (SVGD) [35], is adopted to efficiently perform the Bayesian learning with limited training overhead compared to its deterministic version. The merit of the proposed method is demonstrated on the reconstruction of idealized vascular flows with sparse and noisy velocity data. The rest of the paper is organized as follows. The proposed physics-constrained Bayesian neural network for flowfield reconstruction is introduced first. Then numerical studies on test flows with two idealized vascular geometries are presented. The roles of data and physical constraints in deep learning will be discussed. Finally, we will conclude the paper in the last paragraph.

The general idea of this work is to reconstruct a high-resolution flow field from low-resolution (sparse or possibly noisy) measurement data based on deep neural networks (DNN). Instead of training the DNN on extra offline databases of high-resolution flow fields, physical/physiological principles are leveraged to constrain the learning process and provide additional information for super-resolution. Namely, a pointwise DNN model will be trained on sparse velocity data to capture the flow field. In the meantime, the physical laws are imposed on a large number of spatiotemporal collocation points where measurements are not available. Therefore, the trained DNN is a smooth function in spatiotemporal space and can reconstruct the flow field with arbitrarily high resolution. The physics-constrained deep learning is usually formulated as a *deterministic* optimization problem, where a loss function is defined by combining both the data mismatches and the residuals of governing equations of a physical model, e.g., incompressible Navier–Stokes equations for Newtonian flows [28, 30]. By minimizing the physics-informed loss, the solution is expected to satisfy the physical model as well as match the training data. This formulation here is referred to as the *deterministic physics-constrained deep learning*. However, when the physical model is not perfect and noisy data are used, the prediction uncertainty regarding model inadequacy and measurement noise cannot be considered in such a deterministic learning process. To address this issue, we developed a *probabilistic physics-constrained Bayesian learning* framework, where the physics-constrained training is formulated in a Bayesian way. Instead of defining the loss, a physics-informed *likelihood* function is constructed, where the measurement noise and equation residuals are modeled as random variables with specified distributions. Given the physics-informed likelihood and specified prior information (DNN initialization), the posterior distribution of the DNN weights can be computed based on the Bayes's theorem. Considering the high dimensionality of DNN, variational inference (VI) is employed to enable feasible Bayesian deep learning. All these components are described further below.

As mentioned above, a DNN approximator $\mathbf{f}^\theta(t, \mathbf{x}) = [\mathbf{u}^\theta, P^\theta]$ is constructed to capture the true pointwise flow solution $\tilde{\mathbf{f}}(t, \mathbf{x}) = [\tilde{\mathbf{u}}, \tilde{P}]$, where $\mathbf{u}, P$ represent velocity and pressure, and θ represents DNN parameters (e.g., weights and bias). The training of this neural network relies on two pieces of information: sparse (noisy) velocity data $\mathbf{u}^d$ and a physical model of the fluid system. The data-based loss component can be defined straightforwardly as the data mismatch, $\|\mathbf{u}^\theta - \mathbf{u}^d\|$, while the physics-based loss component is built upon the fluid governing equations. Here, we model the fluid dynamics by a set of incompressible Navier–Stokes equations with the Newtonian assumption,

$$\mathcal{R}(\mathbf{u}, P) = 0 := \begin{cases} \underbrace{\nabla \cdot \mathbf{u} = 0}_{\text{Mass conservation}}, & \mathbf{x}, t \in \Omega_f \times [0, T], \\ \underbrace{\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} + \frac{1}{\rho}\nabla P - \nu\nabla^2\mathbf{u} + \mathbf{b}_f = 0}_{\text{Momentum conservation}}, \\ & \mathbf{x}, t \in \Omega_f \times [0, T], \end{cases}$$

(1)

where $\mathcal{R}(\mathbf{u}, P)$ represents the residual function of the Navier–Stokes equations; t and $\mathbf{x}$ are temporal and spatial coordinates, respectively; ρ and ν are density and viscosity of the fluid, respectively; $\mathbf{b}_f$ is the body force; To determine unique flow solutions, proper initial $(\mathcal{I}(\mathbf{u}, P) = 0)$ and boundary conditions $(\mathcal{B}(\mathbf{u}, P) = 0)$ are required. The DNN-approximated solutions $[\mathbf{u}^\theta, P^\theta]$ are also expected to comply with the physical model, and thus the violation of the Eq. (1) will be penalized as well. Hence, the physics-regularized loss function can be defined as,

$$\mathcal{L}(\theta) = \left\| \mathcal{R}(\mathbf{u}^\theta, P^\theta) \right\| + \lambda_d \left\| \mathbf{u}^\theta - \mathbf{u}^d \right\|, \tag{2}$$

where all the derivative terms in $\mathcal{R}$ are computed using automatic differentiation and λ_d is a trainable penalty coefficient. The physics-constrained training is defined as a constrained optimization problem,

$$\theta^* = \underset{\theta}{\arg\min}\ \mathcal{L}(\theta),$$
$$\text{s.t.} \begin{cases} \mathcal{I}(\mathbf{x}, P^\theta, \mathbf{u}^\theta) = 0, & t = 0, \ \text{in}\ \Omega_f, \\ \mathcal{B}(t, \mathbf{x}, P^\theta, \mathbf{u}^\theta) = 0, & \text{on}\ \partial\Omega_f. \end{cases} \tag{3}$$

To impose the initial and boundary conditions (IC&BC), two strategies can be used: (i) IC&BC are formulated as additional penalty terms into the loss function and imposed in a soft manner, or (ii) they can be encoded into the DNN structure in a hard manner as shown in Ref. [28]. In this work, the pressure inlet/outlet boundary conditions will be enforced by construction while the no-slip wall boundary condition is imposed softly to avoid involving additional networks for complex geometries. In general, the data loss can only be computed on a handful of points due to data sparsity, but the residual of the physical model will be penalized on a large number of points randomly selected from the physical domain. The Adam stochastic gradient descent (SGD) algorithm [36] is used to solve this optimization problem.

The deterministic formulation of physics-constrained DNN has limitations when it comes to noisy data and imperfect physical models. To reflect uncertainties associated with the data and model, a probabilistic formulation should be considered, where the training is conducted in a Bayesian way. Namely, the DNN $\mathbf{f}^\theta(t, \mathbf{x})$ is initialized by specifying a prior distribution $p(\theta)$ for network parameters θ. By constructing the likelihood function $p(\mathcal{D}, \mathcal{R}|\theta)$ based on the sparse data $\mathcal{D} = \{\mathbf{u}^d\}$ and physical model $\mathcal{R} = 0$, the posterior distribution $p(\theta|\mathcal{D}, \mathcal{R})$ can be obtained using Bayes' rule,

$$p(\theta|\mathcal{D}, \mathcal{R}) \propto p(\theta)p(\mathcal{D}, \mathcal{R}|\theta). \tag{4}$$

By sampling the posterior, the trained DNN can provide a mean prediction as well as estimated uncertainties.

Although efficient Monte Carlo sampling approaches such as Markov chain Monte Carlo (MCMC) are standard for Bayesian inference and have been widely used to approximate the posterior distribution, they are usually infeasible for an extremely high-dimensional problem like DNN training, which may involve millions of parameters. VI, instead, recasts the Bayesian inference as a deterministic optimization problem by minimizing the Kullback–Leibler (KL) divergence between a proposed distribution $q(\theta)$ and the target distribution (i.e., posterior distribution) as,

$$\theta^* = \underset{\theta}{\arg\min}\ \mathbb{KL}\big(q(\theta)||p(\theta|\mathcal{D}, \mathcal{R})\big) = \underset{\theta}{\arg\min}\ \mathbb{E}_q\left[\log\left(\frac{q(\theta)}{p(\theta|\mathcal{D}, \mathcal{R})}\right)\right], \tag{5}$$

where $\mathbb{E}_q(\cdot)$ is the expectation with probability density q. The KL divergence is a measure of the discrepancy between two probability distributions. Most often, the proposed density is parameterized with a specified form of distributions. The performance of the parametric VI largely depends on the predefined family of distributions, which introduces deterministic biases [37]. In this work, a non-parametric VI method, SVGD [35], is adopted, which uses a set of n particles $\{\theta_i\}_{i=1}^n$ to directly minimize the KL divergence without the need of defining variational approximation family. The general idea is to iteratively move the set of particles towards the posterior distribution using the gradient ϕ of KL divergence gradient, which is proved to be proportional to the kernelized Stein operator within the unit ball of a reproducing kernel Hilbert space (RKHS) [37]. Accordingly, the SVGD update equations are given as,

$$\theta_i^{t+1} = \theta_i^t + \epsilon_t \phi(\theta_i^t), \tag{6}$$

where

$$\phi(\theta) = \frac{1}{n}\sum_{j=1}^n \left[k(\theta_j^t, \theta)\underbrace{\nabla_{\theta_j^t}\big(\log p(\theta_j^t) + \log p(\mathcal{D}, \mathcal{R}|\theta_j^t)\big)}_{\text{gradient}} + \underbrace{\nabla_{\theta_j^t} k(\theta_j^t, \theta)}_{\text{repulsive force}} \right], \tag{7}$$

where i represents particle index, ϵ_t is the step size at t iteration, and $k(x, \cdot)$ represents a positive definite kernel (e.g., radial basis function (RBF) is used in current work). As a result, an ensemble of DNNs corresponding to n parameter particles $\{\theta_i\}_{i=1}^n$ are trained by SVGD, where the "gradient" term moves the particles towards high-density regions of the posterior and the "repulsive force" term imposes diversity and avoids particle collapsing. Compared to parametric VI methods, the particle-based SVGD is able to capture multi-modal posteriors.

In realistic applications, a model only approximates reality and has model-form errors. Hence, it is natural to formulate the model constraints in a probabilistic way to reflect inadequacy of a model. Similar to the constrained Bayesian approach proposed by Wu et al. [34], the physical equations $\mathcal{R} = 0$ here are formulated as soft constraints, being a part of the likelihood function. We assume that the residual of governing equations obey a zero-mean Gaussian distribution,

$$p(\mathcal{R}(\mathbf{u}^\theta, P^\theta)|\theta) \sim \mathcal{N}(\mathbf{0}, \Sigma_R), \tag{8}$$

where covariance matrix Σ_R is is control parameter reflecting our confidence in the physical model. As the Navier–Stokes equations well describe the fluid dynamics in general, a small variance $(\sigma = 10^{-4})$ is specified in this work. Nonetheless, in practice, the physical model might be partially unknown or some of the model coefficients are uncertain [38]. This model-form uncertainty can also be characterized within the proposed Bayesian learning framework. Namely, the hyperparameters of the physical likelihood component (e.g., covariance matrix) are treated as trainable parameters, whose posterior can then be

learned during the training process. Without loss of generality, the sparse observation data errors can be assumed to follow zero-mean Gaussian distributions. Therefore, the log(*likelihood*) function can be explicitly written as the sum of log data likelihood and log equation likelihood,

$$\log(likelihood) = \log p(D|\theta, \Sigma_D) + \log p(\mathcal{R}|\theta, \Sigma_R), \tag{9}$$

where the data covariance matrix $\Sigma_D = \text{diag}(\boldsymbol{\sigma}_D)$ are learnable parameters, which can be learned from the data. The prior distribution of the data variance $\boldsymbol{\sigma}_D$ is modeled as an inverse Gamma distribution and the prior of DNN parameters are assumed as a student's *t*-distribution. The physics-constrained SVGD algorithm can be summarized in Algorithm 1.

Algorithm 1: Physics-constrained Stein variational gradient descent

Result: DNN parameters θ and learnable data variance Σ_D.

Sample prior distributions for θ and Σ_D with *n* particles;

for $i = 1 : n_t$ **do**

1. Calculate log posterior: $L(\theta) = \log p(\theta) + \log p(\Sigma_D) + \log p(\mathcal{D}|\theta, \Sigma_D) + \log p(\mathcal{R}|\theta, \Sigma_R)$;
2. Calculate the gradient ∇_θ by back-propagation;
3. Choose an appropriate kernel function $k(\theta, \cdot)$, and calculate kernel Stein operator ϕ, RBF kernel is used in current work;
4. Update θ and Σ_D by stochastic gradient descent (e.g., Adam);

end

After training, the physics-constrained Bayesian DNN can be used to reconstruct the flow field given sparse data and high-resolution coordinates by forward propagation. In the SVGD algorithm, an ensemble of trained DNNs will be obtained from the particle-based posterior approximation. Although the concrete form of the posterior is unknown, the statistics of the flowfield predictions can be estimated by the network ensemble using the Monte Carlo method. For example, the mean velocity field $\overline{\mathbf{u}}^\theta$ is computed by,

$$\overline{\mathbf{u}}^\theta = \mathbb{E}[\mathbf{u}(\mathbf{x}, t)] \approx \frac{1}{N} \sum_{i=1}^{n} \mathbf{u}^{\theta_i}(\mathbf{x}, t), \tag{10}$$

where *n* is the number of DNNs indexed by θ_i. The variance field (reflecting reconstruction uncertainty) is computed based on the law of total variance as shown in Ref. [39], where conditional covariance is defined as,

$$\begin{aligned}
\text{Cov}(\mathbf{u}^\theta|\mathbf{x}, t, D) &= \mathbb{E}_{\mathbf{w},\sigma_D}\left[\text{Cov}(\mathbf{u}^\theta|(\mathbf{x}, t; \mathbf{w}, \sigma_D))\right] \\
&+ \text{Cov}_{\mathbf{w},\sigma_D}\left\{(\mathbb{E}\left[\text{Cov}(\mathbf{u}^\theta|(\mathbf{x}, t; \mathbf{w}, \sigma_D))\right])\right\} \\
&= \mathbb{E}_{\sigma_D}[\Sigma_D] + \mathbb{E}_{\mathbf{w}}\left[\mathbf{u}^\theta(\mathbf{x}, t)^\top \mathbf{u}^\theta(\mathbf{x}, t)\right] \\
&- \mathbb{E}_{\mathbf{w}}\left[\mathbf{u}^\theta(\mathbf{x}, t)\right]\mathbb{E}_{\mathbf{w}}^\top\left[\mathbf{u}^\theta(\mathbf{x}, t)\right] \\
&\approx \frac{1}{N} \sum_{i=1}^{n}\left(\Sigma_{D_i} + \mathbf{u}_i^\theta(\mathbf{x}, t)\mathbf{u}_i^\theta(\mathbf{x}, t)^\top\right) \\
&- \left(\frac{1}{N}\sum_{i=1}^{n}\mathbf{u}_i^\theta(\mathbf{x}, t)\right)\left(\frac{1}{N}\sum_{i=1}^{n}\mathbf{u}_i^\theta(\mathbf{x}, t)\right)^\top, \tag{11}
\end{aligned}$$

where $\Sigma_D^i \sim p(\Sigma_D|D)$. With the defined mean and variance, a probabilistic flow reconstruction result can be obtained.

Several flow cases with two idealized vascular geometries (i.e., stenosis and aneurysm bifurcation) are investigated to demonstrate the performance of the proposed method for flow reconstruction from sparse data. In this study, data are generated by sampling the fully-resolved CFD solutions on sparse locations. We begin our numerical experiments by reconstructing the flow with noise-free data using deterministic physics-constrained (PC) deep learning (cases 1 & 2). Then we evaluate our proposed PC-BNN on the same flow reconstruction problems but with noisy data (cases 3 & 4). Both the reconstructed mean flow fields and uncertainties for different data noise levels are investigated.

A fully-connected network structure of 3 layers and 20 neurons per layer is built for all the flow cases. The Swish activation function [40] is specified in each layer except the output one, where a linear activation is applied. For both deterministic and probabilistic formulations, the Adam optimizer is used for training, where the batch size and initial learning rate are set as 50 and 1×10^{-3}, respectively. In the probabilistic formulation, the prior of NN parameters θ is given by a student's *t*-distribution $\theta \sim \text{St}(\theta|\mu, \lambda, \nu)$, where $\mu = 0, \lambda = 2a_0, \nu = a_0/b_0$. The shape and rate parameters a_0 and b_0 are specified as $a_0 = 1$ and $b_0 = 0.04$, respectively. Furthermore, data uncertainties (noise) are assumed homoscedastic, and thus the covariance matrix of data likelihood is a diagonal matrix where the prior distribution of the diagonal term is assumed to be an inverse Gamma IG($\beta|a_1, b_1$) with $a_1 = 2$ and $b_1 = 1 \times 10^{-6}$. The equation likelihood is assumed to follow a Gaussian distribution with variance $\sigma^2 = 1 \times 10^{-4}$. To perform SVGD, an ensemble of five NN samples are generated based on the prior. The Bayesian DNN and physics-constrained SVGD are implemented in the PyTorch platform [41]. The training of 6×10^4 SGD iterations is performed for deterministic cases, while 1.2×10^5 SGD iterations for probabilistic cases, on an NVIDIA GeForce RTX 2080 graphics processing unit (GPU) card. The code and dataset for this work will become available at https://github.com/Jianxun-Wang/Physics-constrained-Bayesian-deep-learning upon publication.

Case 1: In case 1, we aim to reconstruct the flow field in an idealized stenotic vessel from velocity data on very sparse locations (marked as "x" in Fig. 1a). As mentioned above the data are generated directly from the CFD benchmark (Fig. 1a) without adding any noise. Nonetheless, the data are too sparse to provide sufficient information for flow reconstruction. As shown in Fig. 1d, where the DNN is trained solely based on data, the reconstructed flow is not physical at all and flow features at the tapered region are distorted. If the training process is constrained by the divergence-free condition, i.e., continuity equation, the result can be significantly improved (e.g, flow speed decreasing due to increased radius can be observed in Fig. 1c), but notable discrepancies still exist compared to the CFD benchmark. When the training is constrained by both continuity and momentum equations (i.e., full Navier–Stokes equations) with boundary conditions, the velocity contour of the reconstructed flow field is almost identical to the CFD benchmark (see Fig. 1b). The relative reconstruction errors from the purely data-based learning, divergence-free constrained learning, and Navier–Stokes constrained learning are 24.8%, 11.3%, and 5.6%, respectively. The results showed here clearly demonstrate that proper physical constraints can provide additional information to compensate for data insufficiency and enable physical flow reconstruction using limited measurements.

Case 2: To further demonstrate effectiveness of the physical constraints for super-resolution, a more complex flow (i.e., flow in an idealized aneurysm bifurcation) is considered here. The model has a perfect "T" shape, where the flow starts from the bottom of the vertical tube and goes out through two 90° bifurcation arms, driven by a pressure drop $\Delta P = 0.1$. The dome at the
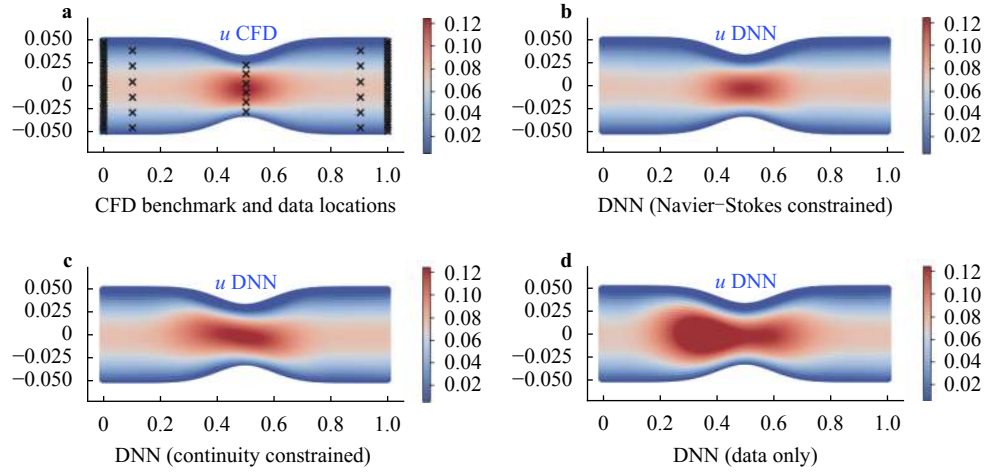

Fig. 1. Comparison of **a** the CFD benchmark (ground truth) with deterministic flow reconstruction results by **b** Navier–Stokes constrained DNN, **c** divergence-free constrained DNN, and **d** purely data-based DNN for a stenotic flow.
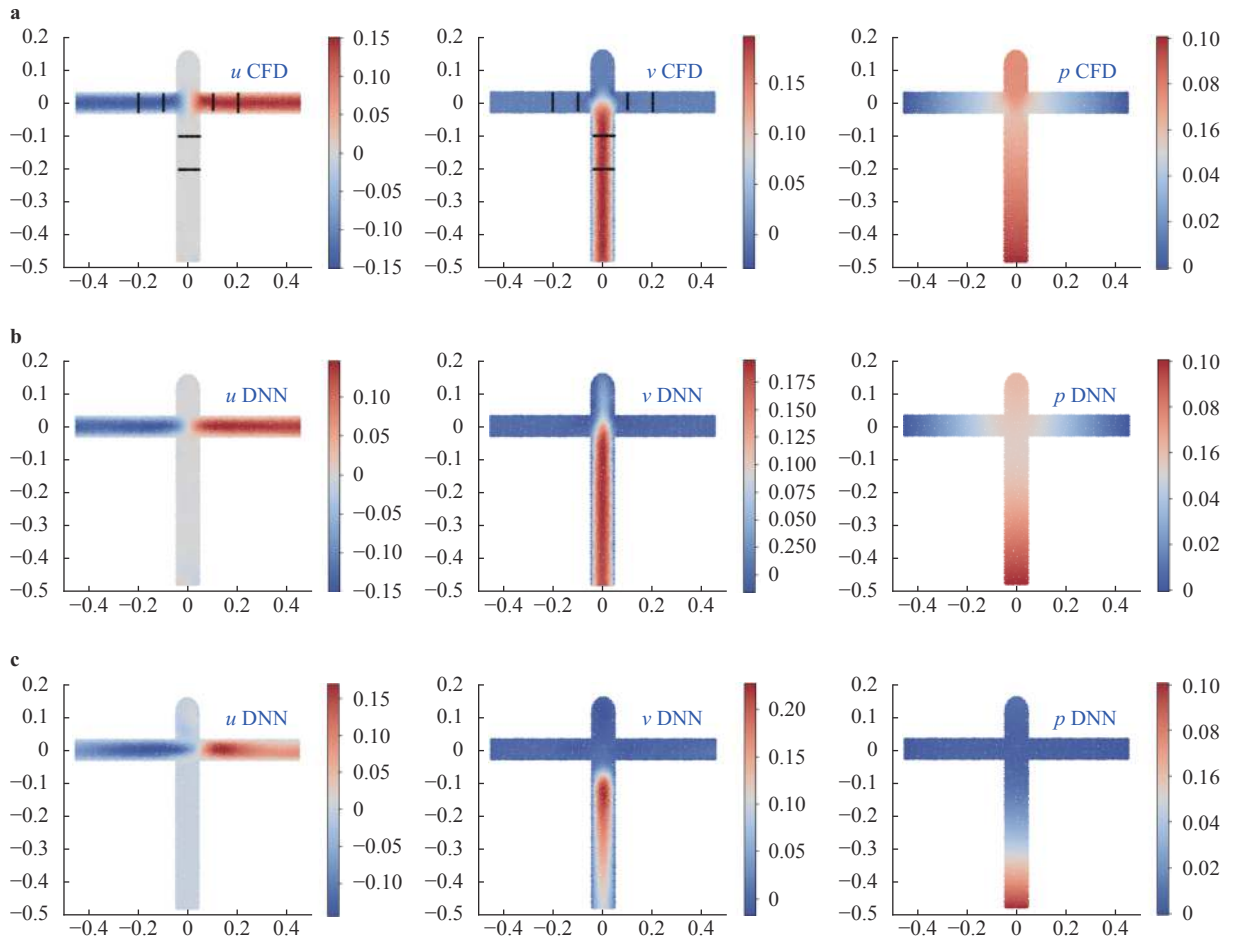

Fig. 2. Comparison of **a** CFD benchmark (with locations of labelled data) with deterministic flow reconstruction by **b** physics-constrained DNN and **c** purely data-based DNN for an aneurysm bifurcation flow.

end of the input tube represents an idealized terminal aneurysm. The data were obtained by probing the CFD velocity field on only six slices, which are very sparse in general (see Fig. 2a). Following the physics-constrained learning, where the wall boundary condition is enforced softly, the reconstructed velocity and pressure fields (see Fig. 2b) agree well with the CFD benchmark. For the sake of comparison, the purely data-based learning results are also presented in Fig. 2c, where the reconstructed velocity fields significantly differ from the CFD benchmarks. It is worthwhile to note that the purely data-based DNN fails to reconstruct the pressure field since no pressure data are used for training. However, the physics-constrained DNN can reasonably capture the general patterns of pressure field because of the constraints on the relation between pressure and velocity, imposed by the Navier–Stokes equations. Quantitatively, the relative reconstruction errors $(|f_{DNN} - f_{CFD}|/|f_{CFD}|)$ in u, v, and P from the purely data-based DNN are $35.1\%, 40.5\%$, and 69.9%, respective, while for Navier–Stokes constrained learning, the relative errors can be reduced to $13.7\%, 12.1\%$, and 12.8%. These comparisons show that the PC-NN remarkably improves the reconstruction accuracy for velocity, and it also can infer the pressure field with the same level of accuracy, where no data are used for training.

Case 3: In the two cases presented above, the flow fields are reconstructed from noise-free data based on deterministic physics-constrained learning. However, when the data are not only sparse but also noisy, the uncertainty due to measurement noises should be reflected in the reconstructed flow. Hence, PC-BNN is trained on the sparse, noisy data to enable robust flow reconstruction with quantified uncertainties. The same flow in Case 1 is reconstructed but with noisy data sampled from the CFD benchmark solutions that are corrupted by Gaussian noises of different levels. Similarly, the noisy flow is "observed" only at a few locations indicated by "x". Figure 3 shows the flow reconstruction results by PC-BNN, while the purely data-based solution is also plotted for comparison. We can see that the flow field corrupted by 10% Gaussian noise becomes unsmooth (Fig. 3a), and the purely data-based flow estimation (Fig. 3b) fails to capture any physical flow patterns. The relative reconstruction error increases to 83.3% (Fig. 1d). This is expected since the data are lack of both quantity and quality. In contrast, the mean-field of the reconstructed flow by PC-BNN (Fig. 3c) is in a good agreement with the CFD benchmark (Fig. 1a) and the noise can be notably reduced as well. The relative error of the mean reconstructed field is reduced from 83.3% to 6.9% by introducing the Navier–Stokes equation constraint. Moreover, the uncertainty of the reconstructed flow can be reasonably estimated as shown by the standard deviation (std) field in Fig. 3d. We have studied the reconstruction performance given different data noise levels $(5\%, 10\%, 15\%)$, and the prediction results and uncertainties are summarized in Table 1. The reconstruction error of the purely data-based DNN remarkably increases with increased data noise. Although the accuracy of the PC-BNN predictions also slightly decreases with the increased noise level, the performance is still satisfactory and flow physics can be captured reasonably well. It is important to note that the mean and maximum std of the reconstructed field also increases as the data noise becomes larger, demonstrating that the uncertainty of the reconstructed results can be well estimated by the PC-BNN, which reflects the effect of data noises.

Case 4: Lastly, the PC-BNN is applied to reconstruct the bifurcation flow from noisy, sparse data. Gaussian noises with different variances are added onto the CFD solution and six sections of the corrupted flow data (see Fig. 4) are used for training. Figure 4a shows the corrupted CFD solution and marks the locations of training data by "x". For the purely data-based learning, the reconstruction result by noisy data is much worse than that using noise-free data, which is expected (see Fig. 4b). In contrast, the PC-BNN still accurately captures the flow field and the mean velocity contour shown in Fig. 4c agrees with the CFD benchmark in Fig. 2a. Furthermore, the reconstruction uncertainty introduced by data noise can be reflected by the std field in Fig. 4d. The uncertainty is large at the left outlet region, indicating that the prediction at this area has low fidelity. Similarly, the performance of the PC-BNN on the data with different noise levels are studied, and the reconstruction accuracy and uncer-
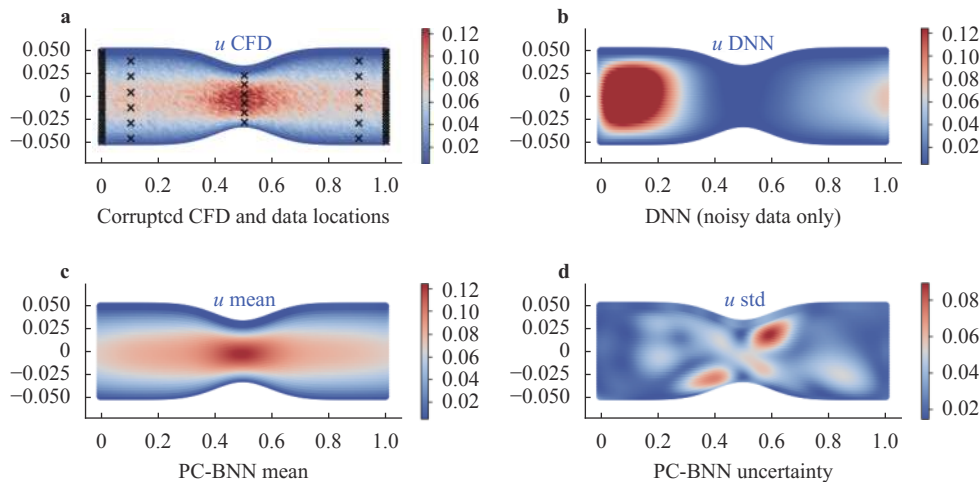

Fig. 3. **a** Comparison of reconstruction results of the stenotic flow with sparse, noisy data (10% noise) between **b** purely data-based DNN prediction and **c** mean velocity field reconstructed by PC-BNN. The uncertainty of PC-BNN based flow reconstruction (i.e., standard deviation field) is shown in panel **d**.
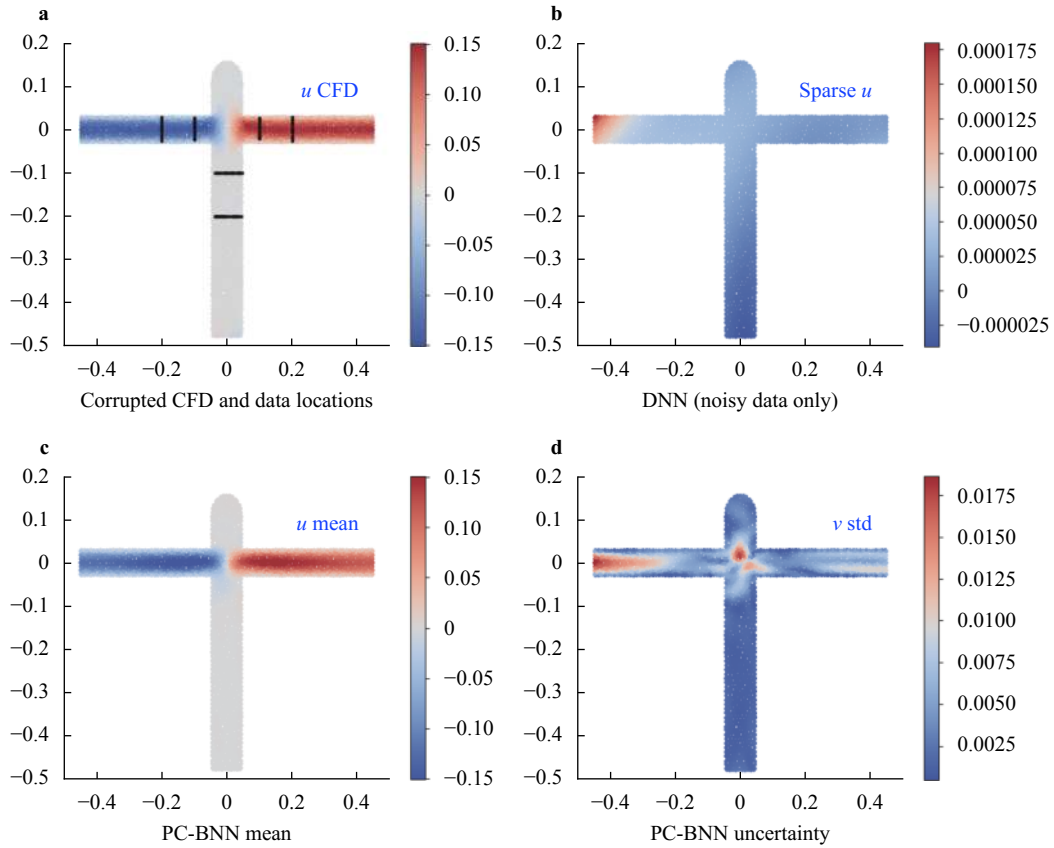
Fig. 4. **a** Comparison of reconstruction results of the bifurcation flow with sparse, noisy data (10% noise) between **b** purely data-based DNN prediction and **c** mean velocity field reconstructed by PC-BNN. The uncertainty of PC-BNN based flow reconstruction (i.e., standard deviation field) is shown in panel **d**.

Table 1 Mean errors and uncertainties of reconstructed u velocity fields from sparse, noisy data using PC-BNN. Note that both the error and std are normalized by the corresponding CFD benchmark solution.

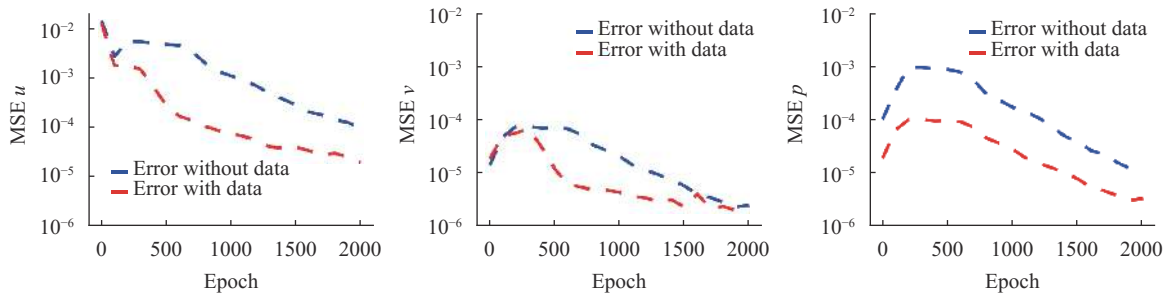| Reconstructed u velocity field of stenotic flow		
Level of noise	Relative mean error	Uncertainty (mean/max std)
5% noise Data only	0.623	Not applicable
PC-BNN	0.041	0.038/0.076
10% noise Data only	0.833	Not applicable
PC-BNN	0.069	0.050/0.083
15% noise Data only	0.941	Not applicable
PC-BNN	0.100	0.064/0.125
Reconstructed u velocity field of bifurcation flow		
Level of noise	Relative mean error	Uncertainty (mean/max std)
5% noise Data only	1.000	Not applicable
PC-BNN	0.132	0.046/0.210
10% noise Data only	1.000	Not applicable
PC-BNN	0.126	0.062/0.253
15% noise Data only	1.000	Not applicable
PC-BNN	0.173	0.067/0.325

Fig. 5. Test error histories of the physics-constrained learning with and without using training data. The mean square errors (MSE) of *u* (left), *v* (middel), and *p* (right) predictions are compared.

tainty are summarized in Table 1. The same trend as shown in Case 3 can be found: the mean reconstruction error and uncertainty will increase as the noise grows. The results from both Cases 3 & 4 show that the proposed PC-BNN can accurately reconstruct a high-resolution flow field from sparse and noisy data, and the prediction uncertainty can also be estimated.

In this work, the numerical results have demonstrated that a high-resolution flow field can be recovered following physics-constrained learning with sparse data and known physical constraints (i.e., Navier–Stokes equations). However, a previous work [28] in the context of surrogate modeling has shown that flow solutions of the Navier–Stokes equations can be obtained from physics-constrained deep learning even without any labeled data if the boundary conditions are imposed properly. Therefore, it is interesting to know what benefit can be gained by introducing additional sparse labeled data into PDE-constrained learning. Taking the stenotic flow as an example, we conducted a comparison study between the *data-free* PDE-constrained learning and *weakly data-based* (i.e., sparse data-based) PDE-constrained learning, where boundary conditions are both imposed softly with a penalty parameter $\lambda = 0.1$. Figure 5 shows the histories of test errors for velocity and pressure versus the number of training epochs. The test error of the weakly data-based learning (red dashed line) decreases much faster than that of the data-free learning (blue solid line). With the same number of training epochs, the prediction error from sparse data-based, physics-constrained learning is about one order of magnitude lower than purely physics-constrained learning without any labeled data. The comparison indicates that adding some labeled data would further improve the equation-constrained learning, which is consistent with the intuition since more information is used for the neural network training.

The objective of this work is to reconstruct a high-resolution flow field from sparse and possibly noisy data. To achieve this goal, we proposed a physics-constrained Bayesian deep learning framework, where the likelihood function is constructed based on the measurement uncertainty and model inadequacy. Stein variation gradient descent is used to enable efficient Bayesian learning. The proposed approach is able to reconstruct the flow field with estimated uncertainties particularly when data are corrupted with measurement noise. Numerical experiments were conducted on a number of flow reconstruction cases with idealized vascular geometries, where synthetic data are used to evaluate the performance of the proposed method. We have demonstrated that the constraints of a physical model can significantly improve the reconstruction results from limited clean data. When the data are noisy, our proposed PC-BNN can accurately predict the mean flow field, meanwhile reasonably estimate the prediction uncertainties corresponding to different data noise levels.

Acknowledgement

References

[1] C.M. Lawley, K.M. Broadhouse, F.M. Callaghan, et al., 4D flow magnetic resonance imaging: role in pediatric congenital heart disease, Asian Cardiovascular and Thoracic Annals 26 (2018) 28–37.

[2] F.M. Callaghan, S.M. Grieve, Spatial resolution and velocity field improvement of 4D-flow MRI, Magnetic Resonance in Medicine 78 (2017) 1959–1968.

[3] R. Krishnamurthy, A. Choukulkar, R. Calhoun, et al., Coherent Doppler lidar for wind farm characterization, Wind Energy 16 (2013) 189–206.

[4] L. Sirovich, Turbulence and the dynamics of coherent structures. I. coherent structures, Quarterly of Applied Mathematics 45 (1987) 561–571.

[5] P.J. Schmid, Dynamic mode decomposition of numerical and experimental data, Journal of Fluid Mechanics 656 (2010) 5–28.

[6] J.H. Tu, C.W. Rowley, D.M. Luchtenburg, et al., On dynamic mode decomposition: Theory and applications, Journal of Computational Dynamics 1 (2014) 391–421.

[7] D. Venturi, G.E. Karniadakis, Gappy data and reconstruction procedures for flow past a cylinder, Journal of Fluid Mechanics 519 (2004) 315–336.

[8] T. Bui-Thanh, M. Damodaran, K. Willcox, Proper orthogonal decomposition extensions for parametric applications in compressible aerodynamics, in: 21st AIAA Applied Aerodynamics Conference (2003) 4213.

[9] B. Podvin, Y. Fraigneau, F. Lusseyran, et al., A reconstruction method for the flow past an open cavity, Journal of Fluids En-

gineering 128 (2006) 531–540.

[10] A. Yakhot, T. Anor, G.E. Karniadakis, A reconstruction method for gappy and noisy arterial flow data, IEEE Transactions on Medical Imaging 26 (2007) 1681–1697.

[11] A.I. Moreno, A.A. Jarzabek, J.M. Perales, et al., Aerodynamic database reconstruction via gappy high order singular value decomposition, Aerospace Science and Technology 52 (2016) 115–128.

[12] M. Mifsud, A. Vendl, L.-U. Hansen, et al., Fusing wind-tunnel measurements and CFD data using constrained gappy proper orthogonal decomposition, Aerospace Science and Technology 86 (2019) 312–326.

[13] K. Lee, K.T. Carlberg, Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders, Journal of Computational Physics (2019) 108973.

[14] J.L. Callaham, K. Maeda, S.L. Brunton, Robust flow reconstruction from limited measurements via sparse representation, Physical Review Fluids 4 (2019) 103907.

[15] K. Manohar, B.W. Brunton, J.N. Kutz, et al., Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns, IEEE Control Systems Magazine 38 (2018) 63–86.

[16] D.P. Foures, N. Dovetta, D. Sipp, et al., A data-assimilation method for Reynolds-averaged Navier-Stokes-driven mean flow reconstruction, Journal of Fluid Mechanics 759 (2014) 404–431.

[17] B. Combés, D. Heitz, A. Guibert, et al., A particle filter to reconstruct a free-surface flow from a depth camera, Fluid Dynamics Research 47 (2015) 051404.

[18] R. Kikuchi, T. Misaka, S. Obayashi, Assessment of probability density function based on pod reduced-order model for ensemble-based data assimilation, Fluid Dynamics Research 47 (2015) 051403.

[19] V. Mons, J.-C. Chassaing, T. Gomez, et al., Reconstruction of unsteady viscous flows using data assimilation schemes, Journal of Computational Physics 316 (2016) 255–280.

[20] S. Symon, N. Dovetta, B.J. McKeon, et al., Data assimilation of mean velocity from 2D PIV measurements of flow over an idealized airfoil, Experiments in Fluids 58 (2017) 61.

[21] J.-X. Wang, H. Xiao, Data-driven CFD modeling of turbulent flows through complex structures, International Journal of Heat and Fluid Flow 62 (2016) 138–149.

[22] K. Zhang, W. Zuo, Y. Chen, et al., Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising, IEEE Transactions on Image Processing 26 (2017) 3142–3155.

[23] C. Dong, C.C. Loy, K. He, et al., Learning a deep convolutional network for image super-resolution, in: European Conference on Computer Vision, Springer (2014) 184-199.

[24] J. Yu, J.S. Hesthaven, Flowfield reconstruction method using artificial neural network, AIAA Journal 57 (2018) 482–498.

[25] N.B. Erichson, L. Mathelin, Z. Yao, et al., Shallow learning for fluid flow reconstruction with limited sensors and limited data, arXiv preprint arXiv: 1902.07358.

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

[27] Y. Zhu, N. Zabaras, P.-S. Koutsourelakis, et al., Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data, Journal of Computational Physics 394 (2019) 56–81.

[28] L. Sun, H. Gao, S. Pan, et al., Surrogate modeling for fluid flows based on physicsconstrained deep learning without simulation data, Computer Methods in Applied Mechanics and Engineering. doi: https://doi.org/10.1016/j.cma.2019.112732.

[29] H. Gao, J.-X. Wang, M.J. Zahr, Non-intrusive model reduction of large-scale, nonlinear dynamical systems using deep learning, arXiv preprint arXiv: 1911.03808.

[30] M. Raissi, A. Yazdani, G.E. Karniadakis, Hidden fluid mechanics: A navier-stokes informed deep learning framework for assimilating flow visualization data, arXiv preprint arXiv: 1808.04327.

[31] D. Zhang, L. Lu, L. Guo, et al., Quantifying total uncertainty in physicsinformed neural networks for solving forward and inverse stochastic problems, Journal of Computational Physics 397 (2019) 108850.

[32] Y. Yang, P. Perdikaris, Adversarial uncertainty quantification in physics-informed neural networks, Journal of Computational Physics 394 (2019) 136–152.

[33] N. Geneva, N. Zabaras, Modeling the dynamics of PDE systems with physics-constrained deep auto-regressive networks, Journal of Computational Physics 403 (2020) 109056.

[34] J. Wu, J.-X. Wang, S.C. Shadden, Adding constraints to bayesian inverse problems, in: Proceedings of the AAAI Conference on Artificial Intelligence 33 (2019) 1666-1673.

[35] Q. Liu, D. Wang, Stein variational gradient descent: A general purpose bayesian inference algorithm, in: Advances in Neural Information Processing Systems (2016) 2378-2386.

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

[37] C. Blundell, J. Cornebise, K. Kavukcuoglu, et al., Weight uncertainty in neural networks, in: Proceedings of the 32nd International Conference on International Conference on Machine Learning 37 (2015) 1613-1622.

[38] Z. Chen, R. Zhang, J. Zheng, et al., Sparse bayesian learning for structural damage identification, Mechanical Systems and Signal Processing 140 (2020) 106689.

[39] Y. Zhu, N. Zabaras, Bayesian deep convolutional encoder-decoder networks for surrogate modeling and uncertainty quantification, Journal of Computational Physics 366 (2018) 415–447.

[40] P. Ramachandran, B. Zoph, Q.V. Le, Searching for activation functions, arXiv preprint arXiv: 1710.05941.

[41] A. Paszke, S. Gross, F. Massa, et al., Pytorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems (2019) 8024-8035.