

Deep learning observables in computational fluid dynamics

K. Lye and S. Mishra and Kjetil Lye

Research Report No. 2019-13

March 2019

Latest revision: December 2019

Seminar für Angewandte Mathematik
Eidgenössische Technische Hochschule
CH-8092 Zürich
Switzerland

Deep learning observables in computational fluid dynamics

Kjetil O. Lye ^{*}, Siddhartha Mishra [†] and Deep Ray [‡]

December 16, 2019

Abstract

Many large scale problems in computational fluid dynamics such as uncertainty quantification, Bayesian inversion, data assimilation and PDE constrained optimization are considered very challenging computationally as they require a large number of expensive (forward) numerical solutions of the corresponding PDEs. We propose a machine learning algorithm, based on deep artificial neural networks, that predicts the underlying *input parameters to observable* map from a few training samples (computed realizations of this map). By a judicious combination of theoretical arguments and empirical observations, we find suitable network architectures and training hyperparameters that result in robust and efficient neural network approximations of the parameters to observable map. Numerical experiments are presented to demonstrate low prediction errors for the trained network networks, even when the network has been trained with a few samples, at a computational cost which is several orders of magnitude lower than the underlying PDE solver.

Moreover, we combine the proposed deep learning algorithm with Monte Carlo (MC) and Quasi-Monte Carlo (QMC) methods to efficiently compute uncertainty propagation for nonlinear PDEs. Under the assumption that the underlying neural networks generalize well, we prove that the deep learning MC and QMC algorithms are guaranteed to be faster than the baseline (quasi-) Monte Carlo methods. Numerical experiments demonstrating one to two orders of magnitude speed up over baseline QMC and MC algorithms, for the intricate problem of computing probability distributions of the observable, are also presented.

1 Introduction

Many interesting fluid flows are modeled by so-called *convection-diffusion* equations [29] i.e, nonlinear partial differential equations (PDEs) of the generic form,

$$\mathbf{U}_t + \operatorname{div}_x(\mathbf{F}(\mathbf{U})) = \nu \operatorname{div}_x(\mathbf{D}(\mathbf{U})\nabla_x \mathbf{U}), \quad (x, t) \in D \subset \mathbb{R}^{d_s} \times \mathbb{R}_+, \quad (1.1)$$

with $\mathbf{U} \in \mathbb{R}^m$ denoting the vector of unknowns, $\mathbf{F} = (F_i)_{1 \leq i \leq d_s}$ the *flux vector*, $\mathbf{D} = (D_{ij})_{1 \leq i, j \leq d_s}$ the *diffusion matrix* and ν a small scale parameter representing kinematic viscosity.

Prototypical examples for (1.1) include the compressible Euler equations of gas dynamics, shallow water equations of oceanography and the magnetohydrodynamics (MHD) equations of plasma physics. These PDEs are *hyperbolic systems of conservation laws* i.e, special cases of (1.1) with $\nu = 0$ and with the flux Jacobian $\partial_{\mathbf{U}}(\mathbf{F} \cdot \mathbf{n})$ having real eigenvalues for all normal vectors \mathbf{n} . Another important example for (1.1) is provided by the incompressible Navier-Stokes equations, where $0 < \nu \ll 1$ and the flux function is *non-local* on account of the divergence-free constraint on the velocity field.

It is well known that solutions to convection-diffusion equations (1.1) can be very complicated. These solutions might include singularities such as shock waves and contact discontinuities in the case of hyperbolic systems of conservation laws. For small values of ν , these solutions can be generically unstable, even turbulent, and contain structures with a large range of spatio-temporal scales [29].

^{*}Seminar for Applied Mathematics (SAM), D-Math
ETH Zürich, Rämistrasse 101, Zürich-8092, Switzerland

[†]Seminar for Applied Mathematics (SAM), D-Math
ETH Zürich, Rämistrasse 101, Zürich-8092, Switzerland

[‡]Department of Computational & Applied Mathematics, Rice University, Houston, Texas, USA.

Numerical schemes play a key role in the study of fluid flows and a large variety of robust and efficient numerical methods have been designed to approximate them. These include (conservative) finite difference [31], finite volume [21, 24], discontinuous Galerkin (DG) finite element [24] and spectral (viscosity) methods [50]. These methods have been extremely successful in practice and are widely used in science and engineering today.

The exponential increase in computational power in the last decades provides us with the opportunity to solve very challenging large scale problems in computational fluid dynamics, such as uncertainty quantification (UQ) [4, 19], (Bayesian) inverse problems [48] and real-time optimal control, design and PDE constrained (shape) optimization [6, 51]. In such problems, one is not always interested in computing the whole solution field \mathbf{U} of (1.1). Rather and in analogy with experimental measurements, one is interested in computing the so-called *observables* (functionals or quantities of interest) for the solution \mathbf{U} of (1.1). These can be expressed in the generic form,

$$L(\mathbf{U}) = \int_{D \times \mathbb{R}_+} \psi(x, t) g(\mathbf{U}(x, t)) dx dt. \quad (1.2)$$

Here $\psi : D \times \mathbb{R}_+ \rightarrow \mathbb{R}$ and $g : \mathbb{R}^m \rightarrow \mathbb{R}$ are suitable test functions. Prototypical examples of such observables (functionals) are provided by body forces, such as the *lift* and *drag* in aerodynamic simulations, and by the runup height (at certain probe points) in simulations of tsunamis.

Moreover in practice, one is interested not just in a single value, but rather in the *statistics* of such observables. Typical statistical quantities of interest are the mean, variance, higher-moments and probability density functions (pdfs) of the observable (1.2). Such statistical quantities quantify uncertainty in the solution, propagated from possible uncertainty in the underlying inputs i.e. fluxes, diffusion coefficients, initial and boundary data. They might also stem from the presence of a statistical spread in design parameters such as those describing the geometry of the computational domain. It is customary to represent the resulting solution field as $\mathbf{U} = \mathbf{U}(y)$, with $y \in Y \subset \mathbb{R}^d$, a possibly very high dimensional parameter space. Thus, the goal of many CFD simulations is to compute (statistics of) the so-called *parameters to observable* map $y \mapsto L(\mathbf{U}(y))$.

Calculating a single realization of this parameters to observable map might require a very expensive forward solve of (1.1) with a CFD solver and a quadrature to calculate (1.2). However, in UQ, Bayesian inversion or optimal design and control, one needs to evaluate a large number of instances of this map, necessitating a very high computational cost, even on state of the art HPC systems. As a concrete example, we consider a rather simple yet prototypical situation. We are interested in computing statistics of body forces such as lift and drag for a model two-dimensional RAE2822 airfoil [25] (see figure 2). A single forward solve of the underlying compressible Euler equations on this airfoil geometry on a mesh of high resolution (see figure 2), with a state of the art high-resolution finite volume solver [43], takes approximately 7 wall-clock hours on a HPC cluster (see table 5). This cost scales up in three space dimensions proportionately. However, a typical UQ problem such as Bayesian inversion with a state of the art MCMC algorithm, for instance the Random walk Metropolis-Hastings algorithm [48], might need upto $10^5 - 10^6$ such realizations, rendering even a two-dimensional Bayesian inversion infeasible! This example illustrates the fact that the high-computational cost of the parameters to observable map makes problems such as forward UQ, Bayesian inversion, and optimal control and design very challenging. Hence, *we need numerical methods which allow ultrafast (several order of magnitude faster than state of the art) computations of the underlying parameters to observable map.*

Machine learning, in the form of artificial neural networks (ANNs), has become extremely popular in computer science in recent years. This term is applied to methods that aim to approximate functions with layers of units (neurons), connected by (affine) linear operations between units and nonlinear activations within units, [20] and references therein. *Deep learning*, i.e an artificial neural network with a large number of intermediate (hidden) layers has proven extremely successful at diverse tasks, for instance in image processing, computer vision, text and speech recognition, game intelligence and more recently in protein folding [15], see [30] and references therein for more applications of deep learning. A key element in deep learning is the *training* of tunable parameters in the underlying neural network by (approximately) minimizing suitable *loss functions*. The resulting (non-convex) optimization problem, on a very high dimensional underlying space, is customarily solved with variants of the stochastic gradient descent method [45].

Deep learning is being increasingly used in the context of numerical solution of partial differential equations. Given that neural networks are very powerful universal function approximators [10, 26, 3, 36, 54], it is natural to consider the space of neural networks as an ansatz space for approximating solutions of PDEs. First proposed in [28] on an underlying collocation approach, it has been successfully used recently in different contexts in [40, 41, 35, 12, 13, 23] and references therein. Given spatial and temporal locations as inputs, the (deep) neural networks, proposed by these authors, approximate the solution of the underlying PDE, by outputting function values. This approach appears to work quite well for problems with high regularity (smoothness) of the underlying solutions (see [46, 23]) and/or if the solution of the underlying PDE possesses a representation formula in terms of integrals [13, 23]. Given their low regularity, it is unclear if solutions of (1.1), realized as functions of space and time, can be efficiently learned by deep neural networks [49].

Several papers applying deep learning techniques in the context of CFD advocate embedding deep learning modules within existing CFD codes to increase their efficiency [49, 42, 33]. In this paper, we adopt a different approach and propose to use *deep neural networks to learn and predict the input parameters to observable map*. Our algorithm will be based on fully connected networks (multi-layer perceptrons) which output values of the observable (1.2) for different input parameters $y \in Y \subset \mathbb{R}^d$ with $d \gg 1$. The network will be trained on data, generated from a few, say $\mathcal{O}(100)$, *samples* i.e, realizations of (1.2) with expensive CFD solvers.

However, the task of designing deep neural networks that will approximate the parameters to observable map with reasonable accuracy is far from straightforward on account of the following issues,

- Approximation results for deep neural networks [54, 5, 39], stipulate a network size of $\mathcal{O}\left(\varepsilon^{-\frac{d}{s}}\right)$ for attaining an error of $\mathcal{O}(\varepsilon)$, with s denoting the Sobolev regularity of the underlying map. However, and as explained in section 2.3, the parameters to observable map for most CFD problems is atmost Lipschitz continuous. Thus, we might require unrealistically large networks for approximating the underlying maps in this context.
- Another relevant metric for estimating errors with deep neural network is the so-called *generalization error* of the trained network, (see section 2.3), that measures the errors of the network in predicting *unseen data*. The best available bounds for this error scale (inversely) with the square-root of the number of training samples. However, in this context, our training samples are generated from expensive CFD simulations. Hence, we might end up with large prediction errors in this *data poor* regime.

Hence, *it is quite challenging to find deep neural networks that can accurately approximate maps of low Sobolev regularity in a data poor regime*. Moreover, there are several hyperparameters that need to be specified in this framework, for instance size of the networks, choice of which variant of the stochastic gradient algorithm that one uses, choice of loss functions and regularizations thereof etc. A priori, the prediction errors can be sensitive to these choices. We propose an *ensemble training* procedure to search the hyperparameter space systematically and identify network architectures that are efficient in our context. By a combination of theoretical considerations and rigorous empirical experimentation, we provide a recipe for finding appropriate deep neural networks to compute the parameters to observable map i.e, those network architectures which ensure a low generalization error, even for relatively few training samples.

A second aim of this paper is to employ these trained neural networks, in conjunction with (Quasi-)Monte Carlo methods to compute statistical quantities of interest, particularly the underlying probability distributions of the observable and to demonstrate that these *deep learning (quasi)-Monte Carlo* methods significantly outperform the baseline algorithms.

The rest of the paper is organized as follows: the deep learning algorithm is presented in section 2. The deep learning (Quasi-)Monte Carlo algorithm for forward UQ is presented in section 3 and some details for the implementation of both sets of algorithms are provided in section 4. In section 5, we present numerical experiments illustrating the proposed algorithms and the results of the paper are summarized and discussed in section 6.

2 Deep learning algorithm

2.1 The problem

We consider the following very general form of a *parameterized* convection-diffusion equation,

$$\begin{aligned} \partial_t \mathbf{U}(t, x, y) + \operatorname{div}_x(\mathbf{F}(y, \mathbf{U})) &= \nu \operatorname{div}_x(\mathbf{D}(y, \mathbf{U}) \nabla_x \mathbf{U}), \quad \forall (t, x, y) \in [0, T] \times D(y) \times Y, \\ \mathbf{U}(0, x, y) &= \bar{\mathbf{U}}(x, y), \quad \forall (x, y) \in D(y) \times Y, \\ L_b \mathbf{U}(t, x, y) &= \mathbf{U}_b(t, x, y), \quad \forall (t, x, y) \in [0, T] \times D(y) \times Y \end{aligned} \quad (2.1)$$

Here, Y is the parameter space and without loss of generality, we assume it to be $Y = [0, 1]^d$, for some $d \in \mathbb{N}$.

The spatial domain is labeled as $y \rightarrow D(y) \subset \mathbb{R}^{d_s}$ and $\mathbf{U} : [0, T] \times D \times Y \rightarrow \mathbb{R}^m$ is the vector of unknowns. The flux vector is denoted as $\mathbf{F} = (F_i)_{1 \leq i \leq d_s} : \mathbb{R}^m \times Y \rightarrow \mathbb{R}^m$ and $\mathbf{D} = (D_{ij})_{1 \leq i, j \leq d_s} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is the *diffusion matrix*.

The operator L_b is a *boundary* operator that imposes boundary conditions for the PDE, for instance the no-slip boundary condition for incompressible Navier-Stokes equations or characteristic boundary conditions for hyperbolic systems of conservation laws. Additional conditions such as hyperbolicity for the flux function \mathbf{F} and positive-definiteness of the diffusion matrix \mathbf{D} might also be imposed, depending on the specific problem that is being considered.

We remark that the parameterized PDE (2.1) will arise in the context of (both forward and inverse) UQ, when the underlying convection-diffusion equation (1.1) contains uncertainties in the domain, the flux and diffusion coefficients and in the initial and boundary data. This input uncertainty propagates into the solution. Following [4] and references therein, it is customary to model such random inputs and the resulting solution uncertainties by *random fields*. Consequently, one can parameterize the probability space on which the random field is defined in terms of a parameter space $Y \subset \mathbb{R}^d$, for instance by expressing random fields in terms of (truncated) Karhunen-Loeve expansions. By normalizing the resulting random variables, one may assume $Y = [0, 1]^d$, with possibly a large value of the parameter dimension d . Moreover, there exists a measure, $\mu \in \operatorname{Prob}(Y)$, with respect to which the data from the underlying parameter space is drawn.

We point out that the above framework is also relevant in problems of optimal control, design and PDE constrained optimization. In these problems, the parameter space Y represents the set of design or control parameters.

For the parameterized PDE (2.1), we aim to compute observables of the following general form,

$$L_g(y, \mathbf{U}) := \int_0^T \int_{D_y} \psi(x, t) g(\mathbf{U}(t, x, y)) dx dt, \quad \text{for } \mu \text{ a.e } y \in Y. \quad (2.2)$$

Here, $\psi \in L^1_{\text{loc}}(D_y \times (0, T))$ is a *test function* and $g \in C^s(\mathbb{R}^m)$, for $s \geq 1$. Most interesting observables encountered in experiments, such as the lift and the drag, can be cast in this general form.

For fixed functions ψ, g , we also define the *parameters to observable* map:

$$\mathcal{L} : y \in Y \rightarrow \mathcal{L}(y) = L_g(y, \mathbf{U}), \quad (2.3)$$

with L_g being defined by (2.2).

We also assume that there exist suitable numerical schemes for approximating the convection-diffusion equation (2.1) for every parameter vector $y \in Y$. These schemes could be of the finite difference, finite volume, DG or spectral type, depending on the problem and on the baseline CFD code. Hence for any mesh parameter (grid size, time step) Δ , we are assuming that for any parameter vector $y \in Y$, a high-resolution approximate solution $\mathbf{U}^\Delta(y) \approx \mathbf{U}(y)$ is available. Hence, there exists an approximation to the *input to observable* map \mathcal{L} of the form,

$$\mathcal{L} : y \in Y \rightarrow \mathcal{L}(y) = L_g(y, \mathbf{U}^\Delta), \quad (2.4)$$

with the integrals in (2.2) being approximated to high accuracy by quadratures. Therefore, the original input parameters to observable map \mathcal{L} is approximated by \mathcal{L}^Δ to very high accuracy i.e, for every value of a tolerance $\varepsilon > 0$, there exists a $\Delta \ll 1$, such that

$$\|\mathcal{L}(y) - \mathcal{L}^\Delta(y)\|_{L_\mu^p(Y)} < \varepsilon, \quad (2.5)$$

for some $1 \leq p \leq \infty$ and weighted norm,

$$\|f\|_{L_\mu^p(Y)} := \left(\int_Y |f(y)|^p d\mu(y) \right)^{\frac{1}{p}},$$

for $1 \leq p < \infty$. The L_μ^∞ norm is analogously defined.

2.2 Deep learning the parameters to observable map

As stated earlier, it can be very expensive to compute the map $\mathcal{L}^\Delta(y)$ for each single realization, $y \in Y$, as a high-resolution CFD solver, possibly entailing a very large number of degrees of freedom, needs to be used. We propose instead, to *learn* this map by deep neural networks. This process entails the following steps,

2.2.1 Training set.

As is customary in supervised learning [20] and references therein, we need to generate or obtain data to train the network. To this end, we select a set of parameters $\mathcal{S} = \{y_i\}_{1 \leq i \leq N}$, with each $y_i \in Y$. The points in \mathcal{S} can be chosen as,

- (i.) randomly from Y , independently and identically distributed with the underlying probability distribution μ .
- (ii.) from a suitable set of quadrature points in Y , for instance the so-called *low discrepancy sequences* that arise in Quasi-Monte Carlo (QMC) quadrature algorithms [8]. Examples of such sequences include Sobol or Halton QMC quadrature points [8].

We emphasize that the generation of QMC points is very cheap, particularly for Sobol or Halton sequences. Moreover, these points are better spread out over the parameter space than a random selection of points and might provide more detailed information about it [8]. Hence, a priori QMC points might be a better choice for sampling the data. One can also replace QMC points with other hierarchical algorithms such as nodes of sparse grids (Smolyak quadrature points) [7] to form the set \mathcal{S} .

Once the training parameter set \mathcal{S} is chosen, we perform a set of high-resolution CFD simulations to obtain $\mathcal{L}^\Delta(y)$, for all $y \in \mathcal{S}$. As each high-resolution CFD simulation could be very expensive, we will require that $N = \#\mathcal{S}$ will not be very large. It will typically be of at most $\mathcal{O}(100)$.




Figure 1: An illustration of a (fully connected) deep neural network. The red neurons represent the inputs to the network and the blue neurons denote the output layer. They are connected by hidden layers with yellow neurons. Each hidden unit (neuron) is connected by affine linear maps between units in different layers and then with nonlinear (scalar) activation functions within units.

2.2.2 Neural network.

Given an input vector, a feedforward neural network (also termed as a multi-layer perceptron), shown in figure 1, consists of layer of units (neurons) which compose of either affine-linear maps between units (in successive layers) or scalar non-linear activation functions within units, culminating in an output [20]. In our framework, for any input vector $z \in Y$, we represent an artificial neural network as,

$$\mathcal{L}_\theta(z) = C_K \circ \sigma \circ C_{K-1} \dots \dots \dots \circ \sigma \circ C_2 \circ \sigma \circ C_1(z). \quad (2.6)$$

Here, \circ refers to the composition of functions and σ is a scalar (non-linear) activation function. A large variety of activation functions have been considered in the machine learning literature [20]. A very popular choice is the *ReLU* function,

$$\sigma(z) = \max(z, 0). \quad (2.7)$$

When, $z \in \mathbb{R}^p$ for some $p > 1$, then the output of the ReLU function in (2.7) is evaluated componentwise.

Moreover, for any $1 \leq k \leq K$, we define

$$C_k z_k = W_k z_k + b_k, \quad \text{for } W_k \in \mathbb{R}^{d_{k+1} \times d_k}, z_k \in \mathbb{R}_{d_k}, b_k \in \mathbb{R}^{d_{k+1}}. \quad (2.8)$$

For consistency of notation, we set $d_1 = d$ and $d_K = 1$.

Thus in the terminology of machine learning (see also figure 1), our neural network (2.6) consists of an input layer, an output layer and $(K - 1)$ hidden layers for some $1 < K \in \mathbb{N}$. The k -th hidden layer (with d_k neurons) is given an input vector $z_k \in \mathbb{R}^{d_k}$ and transforms it first by an affine linear map C_k (2.8) and then by a ReLU (or another) nonlinear (component wise) activation σ (2.7). Although the neural network consists of composition of very elementary functions, its complexity and ability to learn very general functions arises from the interactions between large number of hidden layers [20].

A straightforward addition shows that our network contains $\left(d + 1 + \sum_{k=2}^{K-1} d_k\right)$ neurons.

We denote,

$$\theta = \{W_k, b_k\}, \theta_W = \{W_k\} \quad \forall 1 \leq k \leq K, \quad (2.9)$$

to be the concatenated set of (tunable) weights for our network. It is straightforward to check that $\theta \in \Theta \subset \mathbb{R}^M$ with

$$M = \sum_{k=1}^{K-1} (d_k + 1)d_{k+1}. \quad (2.10)$$

Thus, depending on the dimensions of the input parameter vector and the number (depth) and size (width) of the hidden layers, our proposed neural network can contain a large number of weights. Moreover, the neural network explicitly depends on the choice of the weight vector $\theta \in \Theta$, justifying the notation in (2.6).

Although a variety of network architectures, such as convolutional neural networks or recurrent neural networks, have been proposed in the machine learning literature, [20] and references therein, we will restrict ourselves to fully connected architectures i.e., we do not a priori assume any *sparsity* structure for our set Θ .

2.2.3 Loss functions and optimization.

For any $y \in \mathcal{S}$, we have already evaluated $\mathcal{L}^\Lambda(y)$ from the high-resolution CFD simulation. One can readily compute the output of the neural network $\mathcal{L}_\theta(y)$ for any weight vector $\theta \in \Theta$. We define the so-called *loss function* or mismatch function, as

$$J(\theta) := \sum_{y \in \mathcal{S}} |\mathcal{L}^\Lambda(y) - \mathcal{L}_\theta(y)|^p, \quad (2.11)$$

for some $1 \leq p < \infty$.

The goal of the training process in machine learning is to find the weight vector $\theta \in \Theta$, for which the loss function (2.11) is minimized. The resulting optimization (minimization) problem might lead to

searching a minimum of a non-convex loss function. So, it is not uncommon in machine learning [20] to regularize the minimization problem i.e we seek to find,

$$\theta^* = \arg \min_{\theta \in \Theta} (J(\theta) + \lambda \mathcal{R}(\theta)). \quad (2.12)$$

Here, $\mathcal{R} : \Theta \mapsto \mathbb{R}$ is a *regularization* (penalization) term. A popular choice is to set $\mathcal{R}(\theta) = \|\theta_{\mathbf{W}}\|_q^q$ for either $q = 1$ (to induce sparsity) or $q = 2$. The parameter $0 < \lambda \ll 1$ balances the regularization term with actual loss J (2.11).

The above minimization problem amounts to finding a minimum of a possibly non-convex function over a subset of \mathbb{R}^M for very large M . We can approximate the solutions to this minimization problem iteratively, either by a full batch gradient descent algorithm or by a mini-batch stochastic gradient descent (SGD) algorithm. A variety of SGD algorithms have been proposed in the literature and are heavily used in machine learning, see [45] for a survey. A generic step in a (stochastic) gradient method is of the form:

$$\theta_{r+1} = \theta_r - \eta_r \nabla_{\theta} (J(\theta_k) + \lambda \mathcal{R}(\theta_k)), \quad (2.13)$$

with η_r being the *learning rate*. The stochasticity arises in approximating the gradient in (2.13) by,

$$\nabla_{\theta} J(\theta_k) \approx \nabla_{\theta} \left(\sum_{y \in \hat{\mathcal{S}}_q} |\mathcal{L}^{\Delta}(y) - \mathcal{L}_{\theta}(y)|^p \right), \quad (2.14)$$

and analogously for the gradient of the regularization term in (2.13). Here $\hat{\mathcal{S}}_q \subset \mathcal{S}$ refers to the q -th batch, with the batches being shuffled randomly. Moreover, the SGD methods are initialized with a starting value $\theta_0 = \bar{\theta} \in \Theta$. A widely used variant of the SGD method is the so-called *ADAM* algorithm [27].

For notational simplicity, we denote the (approximate, local) minimum weight vector in (2.12) as θ^* and the underlying deep neural network \mathcal{L}_{θ^*} will be our neural network surrogate for the parameters to observable map \mathcal{L} (2.4). The algorithm for computing this neural network is summarized below,

Algorithm 2.1. Deep learning of parameters to observable map.

Inputs: *Parameterized PDE* (2.1), *Observable* (2.2), *high-resolution numerical method for solving* (2.1) and *calculating* (2.2).

Goal: *Find neural network \mathcal{L}_{θ^*} for approximating the parameters to observable map \mathcal{L} (2.4).*

Step 1: *Choose the training set \mathcal{S} and evaluate $\mathcal{L}^{\Delta}(y)$ for all $y \in \mathcal{S}$ by high-resolution CFD simulations.*

Step 2: *For an initial value of the weight vector $\bar{\theta} \in \Theta$, evaluate the neural network $\mathcal{L}_{\bar{\theta}}$ (2.6), the loss function (2.12) and its gradients to initialize the (stochastic) gradient descent algorithm.*

Step 3: *Run a stochastic gradient descent algorithm of form (2.13) till an approximate local minimum θ^* of (2.12) is reached. The map $\mathcal{L}^* = \mathcal{L}_{\theta^*}$ is the desired neural network approximating the parameters to observable map \mathcal{L} .*

2.3 Theory

2.3.1 Approximation with deep neural networks.

Universal approximation theorems [3, 26, 10] prove that neural networks can approximate any continuous (in fact any measurable) function. However, these universality results are not constructive. More recent papers such as [54, 39] provide the following constructive approximation result:

Theorem 2.2. *Let $f \in W^{s,p}([0, 1]^d, \mathbb{R})$ for some $1 \leq p \leq \infty$, such that $\|f\|_{W^{s,p}} \leq 1$, then for every $\varepsilon > 0$, there exists a neural network $NN(f)$ of the form (2.6) with the ReLU activation function, with $\mathcal{O}(1 + \log(\frac{1}{\varepsilon}))$ layers and of size (number of weights) $\mathcal{O}\left(\varepsilon^{-\frac{d}{s}} (1 + \log(\frac{1}{\varepsilon}))\right)$ such that*

$$\|f - NN(f)\|_{L^p} \leq \varepsilon, \quad 1 \leq p \leq \infty. \quad (2.15)$$

Given the above approximation result, it is essential to investigate regularity of the underlying parameters to observable map \mathcal{L} in order to estimate the size of neural networks, needed to approximate them.

2.3.2 Regularity of the parameters to observable map (2.3)

For simplicity, we assume that the underlying domain $D_y = D \subset B_R$ for a. e. $y \in Y$ and B_R is a ball of radius R . Moreover, we assume that $\psi, g^k \in L^\infty$, with g^k denoting the k -th derivative of the map g in (2.2), for some $k \geq 1$. Then, it is straightforward to obtain the following upper bound,

$$\|\mathcal{L}\|_{W^{k,p}(Y)} \leq C \|\mathbf{U}\|_{W^{k,p}(D \times (0,T) \times Y)}, \quad (2.16)$$

with the constant C depending only on the initial data, Ψ and g . Thus, regularity in the parameters to observable map is predicated on the space-time and parametric regularity of the underlying solution field \mathbf{U} .

Unfortunately, it is well-known that one cannot expect such space-time or parametric regularity for solutions of generic convection-diffusion equations (1.1). In particular, even for the simple case of scalar one-dimensional conservation laws, at best we can show that the solution $\mathbf{U} \in L^\infty((0,T); BV(D \times Y))$. Therefore, one can show readily that

$$\|\mathcal{L}\|_{BV(Y)} \leq C. \quad (2.17)$$

It can be rightly argued that deriving regularity estimates on \mathcal{L} in terms of the underlying solution field \mathbf{U} can lead to rather pessimistic results and might overestimate possible cancellations. This is indeed the case, at least for scalar conservation laws with random initial data i.e, (2.1) with $m = 1, \nu = 0$ and $\mathbf{F}(y, \mathbf{U}) = \mathbf{F}(\mathbf{U})$. In this case, we have the following theorem,

Theorem 2.3. *Consider the scalar conservation law, i.e parameterized convection-diffusion PDE (2.1) with $m = 1$ and $\nu = 0$, in domain $D = \mathbb{R}^{d_s}$ and time period $[0, T]$ and assume that the initial data $u_0 \in W^{1,\infty}(Y; L^1(D))$ and support of u_0 is compact. Assume, furthermore that $\psi \in L^\infty(D \times (0, T))$ and $g \in W^{1,\infty}(\mathbb{R})$. Then, the parameters to observable map \mathcal{L} , defined in (2.3) satisfies,*

$$\|\mathcal{L}\|_{W^{1,\infty}(Y)} \leq C. \quad (2.18)$$

for some constant C , depending on the initial data u_0 , ψ and g . Moreover, if the numerical solution u^Δ is generated by a monotone numerical method, then the approximate parameters to observable map \mathcal{L}^Δ (2.4) satisfies

$$\|\mathcal{L}^\Delta\|_{W^{1,\infty}(Y)} \leq C. \quad (2.19)$$

The proof of this theorem is a consequence of the L^1 stability (contractivity) of the solution operator for scalar conservation laws [11]. In fact, for any $y, y^* \in Y$, a straightforward calculation using the definition (2.3) and the Lipschitz regularity of g yields,

$$\begin{aligned} |\mathcal{L}(y) - \mathcal{L}(y^*)| &\leq \|\psi\|_\infty \|g\|_{\text{Lip}} \int_0^T \|u(t, \cdot, y) - u(t, \cdot, y^*)\|_1 dt, \\ &\leq \|\psi\|_\infty \|g\|_{\text{Lip}} \int_0^T \|u_0(\cdot, y) - u_0(\cdot, y^*)\|_1 dt, \quad \text{by } L^1 \text{-contractivity} \\ &\leq \|\psi\|_\infty \|g\|_{\text{Lip}} T \|u_0\|_{W^{1,\infty}(Y, L^1(D))} \leq C \end{aligned}$$

The above proof also makes it clear that bounds such as (2.18) and (2.19) will hold for systems of conservation laws and the incompressible Navier-Stokes equations as long as there is some stability of the field with respect to the input parameter vector. On the other hand, we cannot expect any such bounds on the higher parametric derivatives of \mathcal{L} , due to the lack of differentiability of the underlying solution field.

Given a bound such as (2.18) or (2.19), we can illustrate the difficulty of approximating the map \mathcal{L} by considering a prototypical problem, namely that of learning the lift and the drag of the RAE2822 airfoil (see section 5.1). For this problem, the underlying parameter space is six-dimensional. Assuming that $s = 1$ in theorem 2.2 and requiring that the approximation error is at most one percent relative error i.e $\varepsilon = 10^{-2}$ in (2.15), yields a neural network of size $\mathcal{O}(10^{12})$ tunable parameters and at least 6 layers. Such a large network is clearly unreasonable as it will be very difficult to train and expensive to evaluate.

2.3.3 Trained networks and generalization error

It can be argued that approximation theoretic results can severely overestimate errors with *trained networks*. Rather, the relevant measure is the so-called *generalization error* i.e, once a (local) minimum θ^* of (2.12) has been computed, we are interested in the following error,

$$\mathcal{E}_G(\theta^*) := \left(\int_Y |\mathcal{L}^\Delta(y) - \mathcal{L}_{\theta^*}(y)|^p d\mu(y) \right)^{\frac{1}{p}}, \quad 1 \leq p \leq \infty. \quad (2.20)$$

In practice, one estimates the generalization error on a so-called *test set* i.e, $\mathcal{T} \subset Y$, with $\#(\mathcal{T}) \gg N = \#(\mathcal{S})$. For instance, \mathcal{T} could consist of i.i.d random points in Y , drawn from the underlying distribution μ . In this case, the generalization error (2.20) can be estimated by the *prediction error*,

$$\mathcal{E}_P(\theta^*) := \left(\frac{1}{\#(\mathcal{T})} \sum_{y \in \mathcal{T}} |\mathcal{L}^\Delta(y) - \mathcal{L}_{\theta^*}(y)|^p \right)^{\frac{1}{p}} \quad (2.21)$$

It turns out that one can estimate the generalization error (2.20) by using tools from machine learning theory [47] as,

$$\mathcal{E}_G(\theta^*) \leq \sqrt{\frac{U}{N}}, \quad (2.22)$$

with $N = \#(\mathcal{S})$ being the number of training samples. The numerator U is typically estimated in terms of the Rademacher complexity or the Vapnik-Chervonenkis (VC) dimension of the network, see [47] and references therein for definitions and estimates.

Unfortunately, such bounds on U are very pessimistic in practice and over estimate the generalization error by several (tens of) orders of magnitude [2, 55, 38]. Recent papers such as [2] provide far sharper bounds by estimating the *compression* i.e number of effective parameters to total parameters in a network.

Nevertheless, even if we make a stringent requirement of $U \sim \mathcal{O}(1)$, we might end up with unacceptably high prediction errors of 30 – 100% for the $\mathcal{O}(10^2 - 10^3)$ training samples that we can afford to generate in practice.

Summarizing, theoretical considerations outlined above indicate the challenges of *finding deep neural networks to accurately learn maps of low regularity in a data poor regime*. We illustrate these difficulties with a simple numerical example.

2.3.4 An illustrative numerical experiment.

In this experiment we consider the parameter space $Y = [0, 1]^6$ the following two maps,

$$\mathcal{L}^1(y) = \sum_{i=1}^6 \sin(4\pi y_i), \quad \mathcal{L}^2(y) = \sum_{i=1}^6 \frac{1}{i^3} \sin(4\pi y_i). \quad (2.23)$$

Both maps are infinitely differentiable but with high amplitudes of the derivatives.

We define a neural network of the form (2.6) and with the architecture specified in table 3. In order to generate the training set, we select the first 128 Sobol points in $[0, 1]^6$ and sample the functions (2.23) at these points. The training is performed with ADAM algorithm and hyperparameters defined in table 4 (first line). The performance of the neural network is ascertained by choosing the first 8192 Sobol points in $[0, 1]^6$ as the test set \mathcal{T} . The results of evaluating the trained networks on the test set is shown in table 1 where we present the relative prediction error (2.21) as a percentage and also present the standard deviation (on the test set) of the prediction error. Clearly, the errors are unacceptably high, particularly for the unscaled sum of sines \mathcal{L}^1 in (2.23). Although still high, the errors reduce by an almost order of magnitude for the scaled sum of sines \mathcal{L}^2 . This is to be expected as the derivatives for this map are smaller. This example illustrates the difficulty to approximating (even very regular) functions, in moderate to high dimensions, by neural networks when the training set is relatively small.

Map	Error (Mean) in %	Error (Std.) in %
$\sum_{i=1}^6 \sin(4\pi x_i)$	133.95	417.83
$\sum_{i=1}^6 \frac{\sin(4\pi x_i)}{i^3}$	43.66	41.26

Table 1: Relative percentage Prediction errors (2.21) with $p = 2$, for the trained neural networks in the sum of sines experiment (2.23).

2.3.5 Pragmatic choice of network size

We estimate network size based on the following very deep but easy to prove fact about training neural networks [55],

Lemma 2.4. *For the map \mathcal{L}^Δ and for the training set \mathcal{S} , with $\#\mathcal{S} = N$, there exists a weight vector $\hat{\theta} \in \Theta$, and a resulting neural network of the form \mathcal{L}_θ , with $\Theta \subset \mathbb{R}^M$ and $M = \mathcal{O}(d + N)$, such that the following holds,*

$$\mathcal{L}^\Delta(z) = \mathcal{L}_{\hat{\theta}}(z), \quad \forall z \in \mathcal{S}. \quad (2.24)$$

The above lemma implies that there exists a neural network with weight vector $\hat{\theta}$ of size $\mathcal{O}(d + N)$ such that the training error defined by,

$$\mathcal{E}_T(\hat{\theta}) := \frac{1}{\#\mathcal{S}} \sum_{y \in \mathcal{S}} |\mathcal{L}^\Delta(y) - \mathcal{L}_{\hat{\theta}}(y)|^p = 0. \quad (2.25)$$

Thus, it is reasonable to expect that a network of size $\mathcal{O}(d + N)$ can be trained by a gradient descent method to achieve a very low training error. It is customary to monitor the generalization capacity of a trained neural network by computing a so-called *validation set* $\mathcal{V} \subset Y$ with $\mathcal{V} \cap \mathcal{S} = \emptyset$, and evaluating the so-called validation loss,

$$J_{\mathcal{V}}(\theta) := \frac{1}{\#\mathcal{V}} \sum_{y \in \mathcal{V}} \|\mathcal{L}^\Delta(y) - \mathcal{L}_\theta(y)\|_p^p. \quad (2.26)$$

A low validation loss is observed to correlate with low generalization errors. In order to generate the validation set, one can set aside a small proportion, say 10 – 20% of the training set as the validation set.

2.4 Hyperparameters and Ensemble training.

Given the theoretical challenges of determining neural networks that can provide low prediction errors in our data poor regime, a suitable choice of the hyperparameters in the training process becomes imperative. To this end, we devise a simple yet effective *ensemble training algorithm*. To this end, we consider the set of hyperparameters, listed in table 2. A few comments regarding the listed hyperparameters are in order. We need to choose the number of layers and width of each layer such that the overall network size, given by (2.10) is $\mathcal{O}(N + d)$. The exponents p, q in the loss function and regularization terms (2.12) usually take the values of 1 or 2 and the regularization parameter λ is required to be small. The starting value $\bar{\theta}$ is chosen randomly.

For each choice of the hyperparameters, we realize a single sample in the hyperparameter space. Then, the machine learning algorithm 2.1 is run with this sample and the resulting loss function minimized. Further details of this procedure are provided in section 4. We remark that this procedure has many similarities to the *active learning* procedure proposed recently in [56].

3 Uncertainty quantification in CFD with deep learning.

We will employ the trained deep neural networks that approximate the parameters to observable map \mathcal{L}^Δ , to efficiently quantify uncertainty in the underlying map. We are interested in computing the entire measure (probability distribution) of the observable. To this end, we assume that the underlying probability distribution on the input parameters to the problem (2.1) is given by $\mu \in \text{Prob}(Y)$. In the context of forward UQ, we are interested in how this initial measure is changed by the parameters to observable

1. Number of hidden layers ($K - 1$)
2. Number of units in k -th layer (d_k)
3. Exponent p in the loss function (2.11).
4. Exponent q in the regularization term in (2.12)
5. Value of regularization parameter λ in (2.12)
6. Choice of optimization algorithm (optimizer) – either standard SGD or ADAM.
7. Initial guess $\bar{\theta}$ in the SGD method (2.13).

Table 2: Hyperparameters in the training algorithm

map \mathcal{L} (or its high-resolution numerical surrogate \mathcal{L}^Δ). Hence, we consider the *push forward* measure $\hat{\mu}^\Delta \in \text{Prob}(\mathbb{R})$ given by

$$\hat{\mu}^\Delta := \mathcal{L}^\Delta \# \mu, \quad \Rightarrow \quad \int_{\mathbb{R}} h(z) d\hat{\mu}^\Delta(z) = \int_Y h(\mathcal{L}^\Delta(y)) d\mu(y), \quad (3.1)$$

for any μ -measurable function $h : \mathbb{R} \mapsto \mathbb{R}$. It should be emphasized that any statistical moment of interest can be computed by integrating an appropriate test function with respect to this measure $\hat{\mu}^\Delta$. For instance, letting $h(w) = w$ in (3.1) yields the mean of the observable. Similarly, the variance can be computed from the mean and letting $h(w) = w^2$ in (3.1). *The task of efficiently computing this probability distribution is significantly harder than just estimating the mean and variance of the underlying map.*

The simplest algorithm for computing this measure $\hat{\mu}^\Delta$ is to use the Monte Carlo algorithm ([8] and references therein). It consists of selecting J samples i.e points $\{y_j\}$ with each $y_j \in Y$ and $1 \leq j \leq J$, that are independent and identically distributed (*iid*). Then, the Monte Carlo approximation of the push-forward measure is given by,

$$\hat{\mu}_{mc} = \frac{1}{J} \sum_{j=1}^J \delta_{\mathcal{L}^\Delta(y_j)} \quad \Rightarrow \quad \int_{\mathbb{R}} h(z) d\hat{\mu}_{mc}(z) = \frac{1}{J} \sum_{j=1}^J h(\mathcal{L}^\Delta(y_j)). \quad (3.2)$$

It can be shown that $\hat{\mu}_{mc} \approx \hat{\mu}^\Delta$ with an error estimate that scales (inversely) as the square root of the number of samples J . Hence, the Monte Carlo algorithm can be very expensive in practice.

3.1 Quasi-Monte Carlo (QMC) methods

3.1.1 Baseline QMC algorithm.

Quasi-Monte Carlo (QMC) methods are deterministic quadrature rules for computing integrals, [8] and references therein. The key idea underlying QMC is to choose a (deterministic) set of points on the domain of integration, designed to achieve some measure of *equidistribution*. Thus, the QMC algorithm is expected to approximate integrals with higher accuracy than MC methods, whose nodes are randomly chosen [8].

For simplicity, we set the underlying measure μ to be the scaled Lebesgue measure on Y . Our aim is to approximate the push forward measure $\hat{\mu}^\Delta$ with the QMC algorithm. To this end, we choose a set of points $\mathcal{J}_q = \{y_j\} \subset Y$ with $1 \leq j \leq J_q = \#(\mathcal{J}_q)$. Then the measure $\hat{\mu}^\Delta$ can be approximated by,

$$\hat{\mu}_{qmc} = \frac{1}{J_q} \sum_{j=1}^{J_q} \delta_{\mathcal{L}^\Delta(y_j)} \quad \Rightarrow \quad \int_{\mathbb{R}} h(z) d\hat{\mu}_{qmc}(z) = \frac{1}{J_q} \sum_{j=1}^{J_q} h(\mathcal{L}^\Delta(y_j)), \quad (3.3)$$

for any measurable function $h : Y \rightarrow \mathbb{R}$.

We want to estimate the difference between the measures $\hat{\mu}^\Delta$ and $\hat{\mu}_{qmc}$ and need some metric on the space of probability measures, to do so. One such metric is the so-called Wasserstein metric [53]. To define this metric, we need the following,

Definition 3.1. Given two measures, $\nu, \sigma \in \text{Prob}(\mathbb{R})$, a transport plan $\pi \in \text{Prob}(\mathbb{R}^2)$ is a probability measure on the product space such that the following holds for all measurable functions F, G ,

$$\int_{\mathbb{R} \times \mathbb{R}} (F(u) + G(v)) d\pi(u, v) = \int_{\mathbb{R}} F(u) d\nu(u) + \int_{\mathbb{R}} G(v) d\sigma(v). \quad (3.4)$$

The set of transport plans is denoted by $\Pi(\nu, \sigma)$ ■

Then the 1-Wasserstein distance [53] between $\hat{\mu}^\Delta$ and $\hat{\mu}_{qmc}$ is defined as,

$$W_1\left(\hat{\mu}^\Delta, \hat{\mu}_{qmc}\right) := \inf_{\pi \in \Pi(\hat{\mu}^\Delta, \hat{\mu}_{qmc})} \int_{\mathbb{R} \times \mathbb{R}} |u - v| d\pi(u, v) \quad (3.5)$$

It can be proved using the Koksma-Hlawka inequality and some elementary optimal transport techniques that the error in the Wasserstein metric behaves as,

$$W_1\left(\hat{\mu}^\Delta, \hat{\mu}_{qmc}\right) \sim V^\Delta \mathcal{D}(\mathcal{J}_q). \quad (3.6)$$

Here, V^Δ is some measure of *variation* of \mathcal{L}^Δ . It can be bounded above by

$$V^\Delta \leq C \int_Y \left| \frac{\partial^d}{\partial y_1 y_2 \cdots y_d} \mathcal{L}^\Delta(y) dy \right|, \quad (3.7)$$

the Hardy-Krause variation of the function \mathcal{L}^Δ . This upper bound requires some regularity for the underlying map in terms of mixed partial derivatives. However, it is well-known that the bound (3.7) is a large overestimate of the integration error [8].

The $\mathcal{D}(\mathcal{J}_q)$ in (3.6) is the so-called *discrepancy* of the point sequence \mathcal{J}_q , defined formally in [8] and references therein. It measures how equally is the point sequence \mathcal{J}_q spread out in Y . The whole objective in the development of quasi-Monte Carlo (QMC) methods is to design *low discrepancy sequences*. In particular, there exists many popular QMC sequences such as Sobol, Halton or Niederreiter [8] that satisfy the following bound on discrepancy,

$$\mathcal{D}(\mathcal{J}_q) \sim \frac{(\log(J_q))^d}{J_q}. \quad (3.8)$$

Thus, the integration error with QMC quadrature based on these rules behaves log-linearly with respect to the number of quadrature points. For simplicity, we assume that the dimension is moderate enough to replace the logarithms in (3.6) by a suitable power leading to

$$W_1\left(\hat{\mu}^\Delta, \hat{\mu}_{qmc}\right) \sim \left(\frac{V^\Delta}{J_q}\right)^\alpha, \quad (3.9)$$

for some $1/2 < \alpha \leq 1$. Thus, for moderate dimensions, we see that the QMC algorithm outperforms the baseline MC algorithm.

Requiring that the Wasserstein distance is of size $\mathcal{O}(\varepsilon)$ for some tolerance $\varepsilon > 0$, entails choosing $J_q \sim \frac{V^\Delta}{\varepsilon^{\frac{1}{\alpha}}}$ and leads to a cost of

$$\hat{C}_{qmc} \sim \mathcal{C} \frac{V^\Delta}{\varepsilon^{\frac{1}{\alpha}}}, \quad (3.10)$$

with \mathcal{C} being the computational cost of a single CFD forward solve. This can be very expensive, particularly for smaller values of α , as the cost of each forward solve is very high.

3.1.2 Deep learning quasi-Monte Carlo (DLQMC)

In order to accelerate the baseline QMC method, we describe an algorithm that combines the QMC method with a deep neural network approximation of the underlying parameters to observable map.

Algorithm 3.2. Deep learning Quasi-Monte Carlo (DLQMC).

Inputs: *Parameterized PDE* (2.1), *Observables* (2.2), *high-resolution numerical method for solving* (2.1) and *calculating* (2.2).

Goal: *Approximate the push-forward measure $\hat{\mu}^\Delta$.*

Step 1: For $J_L \in \mathbb{N}$, select $\mathcal{J}_L = \{y_j\} \subset Y$, with $1 \leq j \leq J_L$, as the first J_L points from a Quasi-Monte Carlo low-discrepancy sequence such as Sobol, Halton etc.

Step 2: For some $N \ll J_L$, denote $\mathcal{S} = \{y_j\} \subset \mathcal{J}_L$, with $1 \leq j \leq N$ i.e, first N points in \mathcal{J}_L , and evaluate the map $h(\mathcal{L}^\Delta(y_j))$, for all $y_j \in \mathcal{J}_L$ with the high-resolution CFD simulation. As a matter of fact, any consecutive set of N QMC points can serve as the training set.

Step 3: With \mathcal{S} as the training set, compute an optimal neural network \mathcal{L}^* for the parameters to observable map \mathcal{L}^Δ , by using algorithm 2.1.

Step 4: Approximate the measure $\hat{\mu}^\Delta$ by

$$\hat{\mu}_{qmc}^* = \frac{1}{J_L} \sum_{j=1}^{J_L} \delta_{\mathcal{L}^*(y_j)}, \quad (3.11)$$

The complexity of the DLQMC algorithm is analyzed in the following theorem.

Theorem 3.3. For any given tolerance $\varepsilon > 0$ and under the assumption that the generalization error (2.20) is of size $\mathcal{O}(\varepsilon)$ and the baseline QMC estimator (3.3) follows the error estimate (3.9), the speedup Σ_{dlqmc} , defined as the ratio of the cost of baseline QMC algorithm and the cost of the DLQMC algorithm 3.2 for approximating the push-forward measure $\hat{\mu}^\Delta$ to an error of size $\mathcal{O}(\varepsilon)$ in the 1-Wasserstein metric, with the DLQMC algorithm 3.2, satisfies

$$\frac{1}{\Sigma_{dlqmc}} \sim \frac{N}{J_q} + \frac{\mathcal{C}_* V^*}{\mathcal{C} V^\Delta}. \quad (3.12)$$

Here, $\mathcal{C}, \mathcal{C}_*$ are the computational cost of computing $\mathcal{L}^\Delta(y)$ (high resolution CFD simulation) and $\mathcal{L}^*(y)$ (deep neural network) for any $y \in Y$ and $V^* = V(\mathcal{L}^*)$ is an estimate on the variation that arises in a QMC estimate such as (3.9).

Proof. We define the measure $\hat{\mu}^* = \mathcal{L}^* \# \mu \in \text{Prob}(\mathbb{R})$ and estimate,

$$W_1(\hat{\mu}^\Delta, \hat{\mu}_{qmc}^*) \leq W_1(\hat{\mu}^\Delta, \hat{\mu}^*) + W_1(\hat{\mu}^*, \hat{\mu}_{qmc}^*).$$

We claim that assuming that the generalization error $\mathcal{E}_G \sim \varepsilon$ implies that $W_1(\hat{\mu}^\Delta, \hat{\mu}^*) \sim \varepsilon$. To see this, we define a transport plan $\pi^* \in \text{Prob}(\mathbb{R}^2)$ by $\pi^* = \hat{\mu}^\Delta \otimes \hat{\mu}^*$. Then,

$$\begin{aligned} \int_{\mathbb{R}^2} |u - v| d\pi^*(u, v) &= \int_{\mathbb{R}^2} |u - v| d(\mathcal{L}^\Delta \# \mu \times \mathcal{L}^* \# \mu)(u, v), \\ &= \int_Y |\mathcal{L}^\Delta(y) - \mathcal{L}^*(y)| d\mu(y) := \mathcal{E}_G \\ &\sim \varepsilon. \end{aligned}$$

This provides an upper bound on the Wasserstein distance (3.5) and proves the claim. Similarly, we use an estimate of the type (3.9) to obtain $W_1(\hat{\mu}^*, \hat{\mu}_{qmc}^*) \sim \left(\frac{V^*}{J_L}\right)^\alpha$. Requiring this error to be $\mathcal{O}(\varepsilon)$ yields $J_L \sim \frac{V^*}{\varepsilon^{\frac{1}{\alpha}}}$ and leads to the following estimate on the total cost of the DLQMC algorithm,

$$\hat{C}_{dlqmc} \sim \mathcal{C}N + \mathcal{C}_* \frac{V^*}{\varepsilon^{\frac{1}{\alpha}}}, \quad (3.13)$$

with $\mathcal{C}, \mathcal{C}_*$ being the computational cost of a single CFD forward solve and evaluation of the deep neural network \mathcal{L}^* , respectively.

Dividing (3.13) with (3.10) and using $J_q \sim \frac{V^\Delta}{\varepsilon^{\frac{1}{\alpha}}}$ leads to the speedup estimate (3.12). \square

Remark 3.4. In particular, we know that the cost $\mathcal{C} = \mathcal{O}\left(\Delta^{-\frac{1}{d+1}}\right)$ is very high for a small mesh size Δ . On the other hand the cost $\mathcal{C}_* = \mathcal{O}(M)$ with M being the number of neurons (2.10) in the network. This cost is expected to be much lower than \mathcal{C} . Moreover, it is reasonable to expect that $V^* \approx V^\Delta$. As long as we can guarantee that $N \ll J_q$, it follows from (3.12) that the speedup Σ_{dlqmc} , which measures the gain of using the deep learning based DLQMC algorithm 3.2 over the baseline Quasi-Monte Carlo method, can be quite substantial. \blacksquare

4 Implementation of ensemble training

In the following numerical experiments, we use fully connected neural networks with ReLU activation functions. Essential hyperparameters are listed in table 2. For each experiment, we use a reference network architecture i.e number of hidden layers and width per layer, such that the total network size is $\mathcal{O}(d + N)$. This ensures consistency with the prescriptions of lemma 2.4.

We use two choices of the exponent p in the loss function (2.11) i.e either $p = 1$ or $p = 2$, denoted as mean absolute error (MAE) or mean square error (MSE), respectively. Similarly the exponent q in the regularization term (2.12) is either $q = 1$ or $q = 2$. In order to keep the ensemble size tractable, we chose four different values of the regularization parameter λ in (2.12) i.e, $\lambda = 7.8 \times 10^{-5}, 7.8 \times 10^{-6}, 7.8 \times 10^{-7}, 0$, corresponding to different orders of magnitude for the regularization parameter. For the optimizer of the loss function, we use either a standard SGD algorithm or ADAM. Both algorithms are used in the *full batch* mode. This is reasonable as the number of training samples are rather low in our case.

A key hyperparameter is the starting value θ in the SGD or ADAM algorithm (2.13). We remark that the loss function (2.12) is non-convex and we can only ensure that the gradient descent algorithms converge to a local minimum. Therefore, the choice of the initial guess is quite essential as different initial starting values might lead to convergence to local minima with very different loss profiles and generalization properties. Hence, it is customary in machine learning to *retrain* i.e, start with not one, but several starting values and run the SGD or ADAM algorithm in parallel. Then, one has to select one of the resulting trained networks. The ideal choice would be to select the network with the best generalization error. However, this quantity is not accessible with the data at hand. Hence, we rely on the following surrogates for predicting the best generalization error,

1. **Best trained network.** (*train*): We choose the network that leads to the lowest training error (2.25), once training has been terminated.
2. **Best validated network.** (*val*): We choose the network that leads to the lowest validation loss (2.26), once training has been terminated. One disadvantage of this approach is to sacrifice some training data for the validation set.
3. **Best trained network wrt mean.** (*mean-train*): We choose the network that has minimized the error in the mean of the training set i.e,

$$\mathcal{E}_{mean} := \left| \frac{1}{N} \sum_{j=1}^N \mathcal{L}^\Delta(y_j) - \frac{1}{N} \sum_{j=1}^N \mathcal{L}^*(y_j) \right|, \quad \forall y_j \in \mathcal{S}. \quad (4.1)$$

4. **Best trained network wrt Wasserstein.** (*wass-train*): We choose the network that has minimized the error in the Wasserstein metric with respect to the training set \mathcal{S} :

$$\mathcal{E}_{wass} := W_1 \left(\frac{1}{N} \sum_{j=1}^N \delta_{\mathcal{L}^\Delta(y_j)}, \frac{1}{N} \sum_{j=1}^N \delta_{\mathcal{L}^*(y_j)} \right), \quad \forall y_j \in \mathcal{S}. \quad (4.2)$$

We note that Wasserstein metric for measures on \mathbb{R} can be very efficiently computed with the Hungarian algorithm [37].

Another hyperparameter (property) is whether we choose randomly distributed points (Monte Carlo) or more uniformly distributed points, such as quasi-Monte Carlo (QMC) quadrature points, to select the training set \mathcal{S} . We use both sets of points in the numerical experiments below in order to ascertain which choice is to be preferred in practice. Other important hyperparameters in machine learning are the *learning rate* in (2.13) and the number of training epochs. Unless otherwise stated, we keep them fixed to the values $\eta_r = 0.01, \forall r$, for both the SGD and ADAM algorithms and to 500000 epochs, respectively.

4.1 Code

All the following numerical experiments use finite volume schemes as the underlying CFD solver. The machine learning and ensemble training runs are performed by a collection of Python scripts and Jupyter

notebooks, utilizing Keras [57] and Tensorflow [58] for machine learning and deep neural networks. The ensemble runs over hyperparameters are parallelized as standalone processes, where each process trains the network for one choice of hyperparameters. Jupyter notebooks for all the numerical experiments can be downloaded from https://github.com/kjetil-lye/learning_airfoils, under the MIT license. A lot of care is taken to ensure reproducibility of the numerical results. All plots are labelled with the git commit SHA code, that produced the plot, in the upper left corner of the plot.




Figure 2: The shape of the RAE airfoil and the underlying primary and voronoi dual meshes.

5 Numerical results

In the following numerical experiments, we consider the compressible Euler equations of gas dynamics as the underlying PDE (1.1). In two dimensions, these equations are

$$\mathbf{U}_t + \mathbf{F}^x(\mathbf{U})_x + \mathbf{F}^y(\mathbf{U})_y = 0, \quad \mathbf{U} = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ E \end{pmatrix}, \quad \mathbf{F}^x(\mathbf{U}) = \begin{pmatrix} \rho u \\ \rho v^2 + p \\ \rho uv \\ (E + p)u \end{pmatrix}, \quad \mathbf{F}^y(\mathbf{U}) = \begin{pmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ (E + p)v \end{pmatrix} \quad (5.1)$$

where $\rho, \mathbf{u} = (u, v)$ and p denote the fluid density, velocity and pressure, respectively. The quantity E represents the total energy per unit volume

$$E = \frac{1}{2}\rho|\mathbf{u}|^2 + \frac{p}{\gamma - 1},$$

where $\gamma = c_p/c_v$ is the ratio of specific heats, chosen as $\gamma = 1.4$ for our simulations. Additional important variables associated with the flow include the speed of sound $a = \sqrt{\gamma p/\rho}$, the Mach number $M = |\mathbf{u}|/a$ and the fluid temperature T evaluated using the ideal gas law $p = \rho RT$, where R is the ideal gas constant.

5.1 Flow past an airfoil

5.1.1 Problem description.

We consider flow past an RAE 2822 airfoil, whose shape is shown in Figure 2 (Left). The flow past this airfoil is a benchmark problem for UQ in aerodynamics [25]. The upper and lower surface of the airfoil are denoted at $(x, S^U(x))$ and $(x, S^L(x))$, with $x \in [0, 1]$. We consider the following parametrized (free-stream) initial conditions and perturbed airfoil shapes

$$\begin{aligned} T^\infty(y) &= 1 + \varepsilon_1 G_1(y), & M^\infty(y) &= 0.729(1 + \varepsilon_1 G_2(y)), & \alpha(y) &= [2.31(1 + G_3(y))]^\circ, \\ p^\infty(y) &= 1 + \varepsilon_1 G_4(y), & \hat{S}^U(x; y) &= S^U(x)(1 + \varepsilon_2 G_5(y)), & \hat{S}^L(x; y) &= S^L(x)(1 + \varepsilon_2 G_6(y)), \end{aligned} \quad (5.2)$$

where α is the angle of attack, the gas constant $R = 1$, $\varepsilon_1 = 0.1$, $\varepsilon_2 = 0.2$, $y \in Y = [0, 1]^6$ is a parameter and $G_k(y) = 2y_k - 1$ for $k = 1, \dots, 6$.

The total drag and lift experienced on the airfoil surface $\hat{S}(y) = \hat{S}^L(y) \cup \hat{S}^U(y)$ are chosen as the observables for this experiment. These are expressed as

$$\text{Lift} = \mathcal{L}_1(y) = \frac{1}{\mathcal{K}(y)} \int_{\hat{S}(y)} p(y) (\mathbf{n}(y) \cdot \boldsymbol{\psi}_l(y)) ds, \quad \text{Drag} = \mathcal{L}_2(y) = \frac{1}{\mathcal{K}(y)} \int_{\hat{S}(y)} p(y) (\mathbf{n}(y) \cdot \boldsymbol{\psi}_d(y)) ds, \quad (5.3)$$

where $\mathcal{K}(y) = \rho^\infty(y) |\mathbf{u}^\infty(y)|^2 / 2$ is the free-stream kinetic energy, while $\boldsymbol{\psi}_l(y) = [-\sin(\alpha(y)), \cos(\alpha(y))]$ and $\boldsymbol{\psi}_d(y) = [\cos(\alpha(y)), \sin(\alpha(y))]$.

Layer	Width (Number of Neurons)	Number of parameters
Hidden Layer 1	12	84
Hidden Layer 2	12	156
Hidden Layer 3	10	130
Hidden Layer 4	12	132
Hidden Layer 5	10	130
Hidden Layer 6	12	156
Hidden Layer 7	10	130
Hidden Layer 8	10	110
Hidden Layer 9	12	132
Output Layer	1	13
		1149

Table 3: Reference neural network architecture for the flow past airfoil problem.

The domain is discretized using triangles to form the primary mesh, while the control volumes are the voronoi dual cells formed by joining the circumcenters of the triangles to the face mid-points. A zoomed view of the primary and dual meshes are shown in Figure 2. The solutions are approximated using a second-order finite volume scheme, with the mid-point rule used as the quadrature to approximate the cell-boundary integrals. The numerical flux is chosen as the kinetic energy preserving entropy stable scheme [43], with the solutions reconstructed at the cell-interfaces using a linear reconstruction with the minmod limiter. For details on flux expression and the reconstruction procedure, we refer interested readers to [43]. The scheme is integrated using the SSP-RK3 time integrator. The contour plots for the Mach number with two different realizations of the parameter $y \in [0, 1]^6$ are shown in Figure 3. It should be emphasized that the stochastic perturbations in (5.2) are rather strong, representing on an average a 10 – 20 percent standard deviation (over mean). Hence, there is a significant statistical spread in the resulting solutions. As shown in figure 3, one of the configurations lead to a transsonic shock around the airfoil while another leads to a shock-free flow around the airfoil.

5.1.2 Generation of training data.

In order to generate the training, validation and test sets, we denote \mathcal{Q}^{sob} as the first 1001 Sobol points on the domain $[0, 1]^6$. For each point $y_j \in \mathcal{Q}^{sob}$, the maps or (rather their numerical surrogate) $\mathcal{L}_{1,2}^\Delta(y_j)$ is then generated from high-resolution numerical approximations $\mathbf{U}^\Delta(y)$ and the corresponding lift and drag are calculated by a numerical quadrature. This forms the test set \mathcal{T} . For this numerical experiment, we take the first 128 Sobol points to generate the training set \mathcal{S} and the next 128 points to generate the validation set \mathcal{V} .

5.1.3 Results of the Ensemble training procedure.

To set up the ensemble training procedure of the previous section, we select a fully connected network, with number of layers and size of each layer listed in table 3. Our network has 1149 tunable parameters and this choice is clearly consistent with the prescriptions of Lemma 2.4. Varying the hyperparameters as described in section 4 led to a total of 114 configurations and each was *retrained* 5 times with different random initial choices of the weights and the biases. Thus, we trained a total of 570 networks in parallel.




Figure 3: Numerical solutions (Mach number) for the flow past an RAE airfoil, with two different realizations of the parameter y .

For each configuration, we select the retraining that minimizes the set selection criteria for the configuration. Once the network with best retraining is selected, it is evaluated on the whole test set $\mathcal{T} = \mathcal{Q}^{sob}$ and the percentage relative L^2 prediction error (2.21) computed. Histograms, describing the probability distribution over all hyperparameter samples for the lift and the drag are shown in figure 4. From this ensemble, we choose the hyperparameter configuration that led to the smallest prediction error and list them in table 4. As seen from this table, there is a difference in the best performing network corresponding to lift and to drag. However, both networks use ADAM as the optimizer and regularize with the L^2 -norm of the weights in the loss function (2.12). We plot the training (and validation) loss, for these

Obs	Opt	Loss	L^1 -reg	L^2 -reg	Selection	BP . Err mean (std)
Lift	ADAM	MSE	0.0	7.8×10^{-6}	wass-train	0.786 (0.010)
Drag	ADAM	MAE	0.0	7.8×10^{-6}	mean-train	1.847 (0.022)

Table 4: The hyperparameter configurations that correspond to the best performing network (one with least mean prediction error) for the two observables of the flow past airfoils.

best performing networks with respect to the number of epochs of the ADAM algorithm in figure 5 (Top) and observe a three orders (for lift) and two orders (for drag) reduction in the underlying loss function during training.

From figure 5 (Bottom) and table 4, we note that the prediction errors for the best performing networks are very small i.e, less than 1% relative error for the lift and less than 2% relative error for the drag, even for the relatively small number (128) of training samples. This is orders of magnitude less than the sum of sines regression problem (table 1). These low prediction errors are particularly impressive when viewed in term of the theory presented in section 2.3. A crude upper bound on the generalization error is given by $\sqrt{\frac{M}{N}}$, with M, N being the number of parameters in the network and number of training samples respectively. Setting these numbers for this problem yields a relative error of approximately 300%. However, our prediction errors are 100–200 times (two orders of magnitude) less, demonstrating significant *compression* in the trained networks [2]. These results are further contextualized, when considered with reference to the computational costs, shown in table 5. In this table, we present the costs of generating a single sample, the average training time (over a subset of hyperparameter configurations) and the time to evaluate a trained network. Each sample is generated on the EULER high-performance cluster of ETH-




Figure 4: Histograms depicting distribution of the percentage relative mean L^2 prediction error (2.21) (X-axis) over number of hyperparameter configurations (samples, Y-axis) for the lift and the drag in the flow past airfoil problem. Note that the range of X-axis is different in the plots.

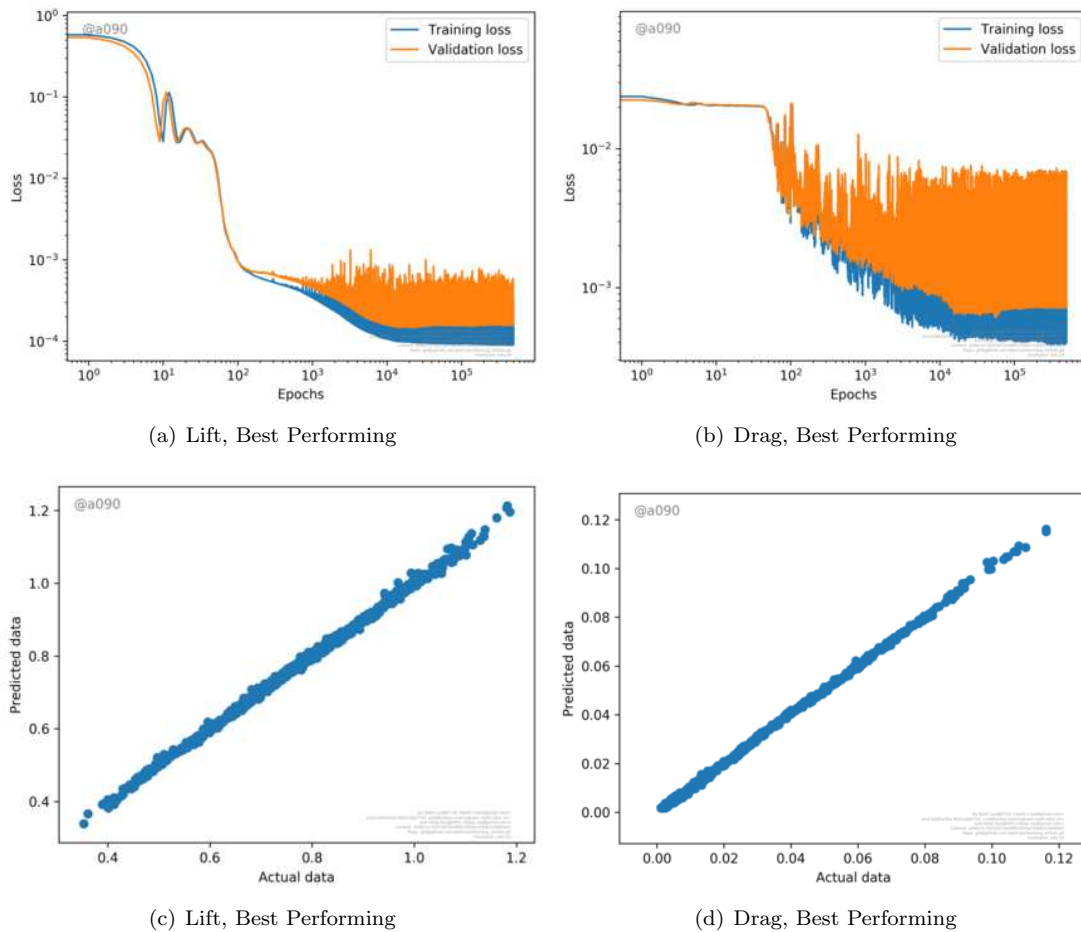


Figure 5: Training loss (2.25) and Validation loss (2.26) (Top) and prediction errors (Ground truth vs. prediction) (Bottom) for the best performing (table 4) neural networks for the flow past airfoils problem as a function of the number of epochs.

	Time (in secs)
Sample generation	24000
Training (Lift)	700
Evaluation (\mathcal{L}_1)	9×10^{-6}
Training (Drag)	840
Evaluation (\mathcal{L}_2)	10^{-5}

Table 5: Computational times (cost) of (single) sample generation, network training and (single) network evaluation, for the flow past airfoils problem. Each sample was generated using a parallelized finite-volume solver on 16 Intel(R) Xeon(R) Gold 5118 @ 2.30GHz processor cores. The entire run took 1500 secs on the cluster and translates into a total wall clock time of $1500 \times 16 = 24000$ secs. The training and evaluation of the neural networks were all performed on a Intel(R) Core(TM) i7-8700K CPU @ 3.70GHz machine. The training and evaluation times are approximations of the average runtimes over all hyperparameter configurations.

Zurich. In this case, the training time is a very small fraction of the time to generate a single sample, let alone the whole training set whereas the computational cost of evaluating the neural networks is almost negligible as it is nine orders of magnitude less than evaluating a single training sample. A priori, it is unclear if the underlying problem has some special features that makes the compression, and hence high accuracy possible. In particular, learning an affine map is relatively straightforward for a neural network. To test this possibility, we approximate the underlying parameters to observable map, for both lift and drag, with classical *linear least squares algorithm* i.e., $\mathcal{L}_j^{\text{lsq}} : \mathbb{R}^6 \rightarrow \mathbb{R}$ be affine map defined by

$$\mathcal{L}_j^{*,\text{lsq}} = \arg \min \left\{ \sum_{y \in \mathcal{S}} (f(y) - \mathcal{L}_j(y))^2 \mid f : \mathbb{R}^6 \rightarrow \mathbb{R} \text{ is affine} \right\}, \quad j = 1, 2. \quad (5.4)$$

The resulting convex optimization problem is solved by the standard gradient descent. The errors with the linear least squares algorithm are shown in figure 4. From this figure we can see that the prediction error with the linear least squares is approximately 4% for lift and 36% for drag. Hence, the best performing neural networks are approximately 5 times more efficient for lift and 20 times more efficient for drag, than the linear least squares algorithm. The gain over linear least squares implicitly measures how *nonlinear* the underlying map is. Clearly in this example, the drag is more nonlinear than the lift and thus more difficult to approximate.

5.1.4 Network sensitivity to hyperparameters.

We summarize the results of the sensitivity of network performance to hyperparameters below,

- *Overall sensitivity.* The overall sensitivity to hyperparameters is depicted as histograms, representing the distribution, over samples (hyperparameter configurations), shown in figure 4. As seen from the figure, the prediction errors with different hyperparameters are spread over a significant range for each observable. There are a few outliers which perform very poorly, particularly for the drag. On the other hand, a large number of hyperparameter configurations concentrate around the best performing network. Moreover all (most) of the hyperparameter choices led to better prediction errors for the lift (drag), than the linear least squares algorithm.
- *Choice of optimizer.* The difference between ADAM and the standard SGD algorithm is shown in figure 6. As in the Sod shock tube problem, ADAM is clearly superior to the standard SGD algorithm. Moreover, there are a few outliers with the ADAM algorithm that perform very poorly for the drag. We call these outliers as the *bad set of ADAM* and they correspond to the configurations with MSE loss function and L^1 regularization, or L^2 regularization, with a regularization parameter $\lambda > 10^{-5}$, or MAE loss function without any regularization. The difference between the bad set of ADAM and its complement i.e., the good set of ADAM is depicted in figure 7.


- *Choice of loss function.* The difference between L^1 mean absolute error (MAE) and L^2 root mean square error (MSE) i.e, exponent $p = 1$ or $p = 2$ in the loss function (2.11) is plotted in the form of a histogram over the error distributions in figure 8. The difference in performance, with respect to choice of loss function is minor. However, there is an outlier, based on MSE, for the drag.
- *Choice of type of regularization.* The difference between choosing an L^1 or L^2 regularization term in (2.12) is shown in figure 9. Again, the differences are minor but the L^2 regularization seems to be slightly better, particularly for the drag.
- *Value of regularization parameter.* The variation of network performance with respect to the value of the regularization parameter λ is shown in figure 11. We observe from this figure that a little amount of regularization works better than no regularization. However, the variation in prediction errors wrt respect to the non-zero values of λ is minor.
- *Choice of selection criteria.* The distribution of prediction errors with respect to the choice of selection criteria for retraining, namely *train, val, mean-train, wass-train* is shown in figure 10. There are very minor differences between *mean-train* and *wass-train*. On the other hand, both these criteria are clearly superior to the other two selection criteria.
- *Sensitivity to retrainings.* The variation of performance with respect to retrainings i.e, different random initial starting values for ADAM, is provided in table 6. In this table, we list the minimum, maximum, mean and standard deviation of the relative prediction error, over 5 retrainings, for the best performing networks for each observable. As seen from this table, the sensitivity to retraining is rather low for the lift. On the other hand, it is more significant for the drag.
- *Variation of Network size.* We consider the dependence of performance with respect to variation of network size by varying the width and depth for the best performing networks. The resulting prediction errors are shown in table 7 where a 3×3 matrix for the errors (rows represent width and columns depth) is tabulated. All the network sizes are consistent with the requirements of lemma 2.4. As observed from the table, the sensitivity to network size is greater for the drag than for the lift. It appears that increasing the width for constant depth results in lower errors for the drag. On the other hand, training clearly failed for some of the larger networks as the number of training samples appears to be not large enough to train these larger networks.

5.1.5 Network performance with respect to number of training samples.

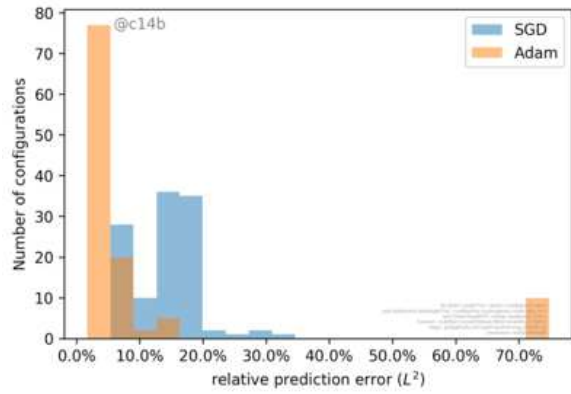
All the above results were obtained with $N = 128$ training samples. We have repeated the entire ensemble training with different samples sizes i.e $N = 32, 64, 256$ also and present a summary of the results in figure 12, where we plot the mean (percentage relative) L^2 prediction error with respect to the number of samples. For each sample size, three different errors are plotted, namely the minimum (maximum) prediction errors and the mean prediction error over a subset of the hyperparameter space, that omits the *bad set of ADAM* as the optimizer. Moreover, we only consider errors with selected (optimal) retrainings in each case. As seen from the figure, the prediction error *decreases with sample size N* as predicted by the theory. The decay seems of the of form $\left(\frac{U_i}{N}\right)^{\alpha_i}$ with $\alpha_1 = 0.81$ and $U_1 = 0.77$ for the lift and $\alpha_2 = 0.92$ and $U_2 = 3.42$ for the drag. In both cases, the decay of error with respect to samples, is at a higher rate than that predicted by the theory in (2.22). Moreover, the constants are much lower than the number of parameters in the network. Both facts indicate a very high degree of compressibility (at least in this possible pre-asymptotic regime) for the trained networks and explains the unexpectedly low prediction errors.

5.1.6 UQ with deep learning

Next, we approximate the underlying probability distributions (measures) (3.1), with respect to the lift and the drag in the flow past airfoils problem with the DLQMC algorithm 3.2. A reference measure is computed from the whole test set of 1001 samples and the corresponding histograms are shown in figure 13 to visualize the probability distributions. In the same figure, we also plot corresponding histograms,




(a) Lift

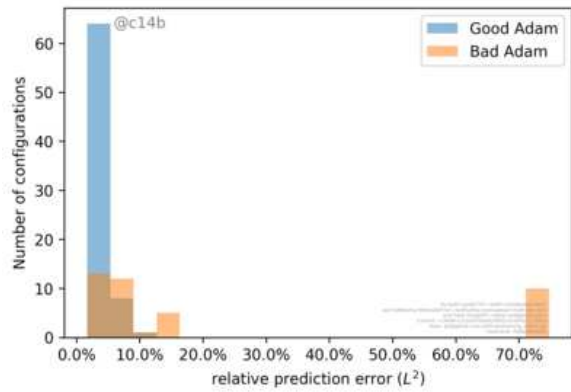


(b) Drag

Figure 6: Histograms for the relative prediction errors comparing the ADAM and SGD algorithms on the flow past airfoils.

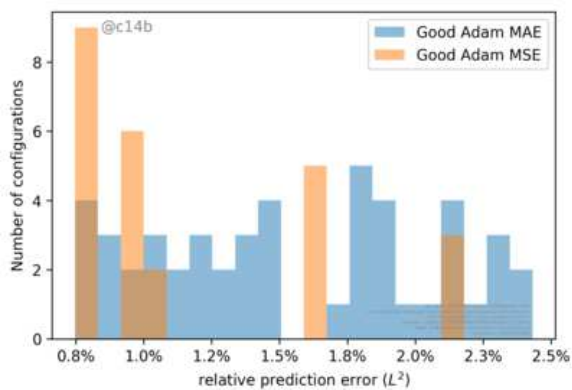


(a) Lift

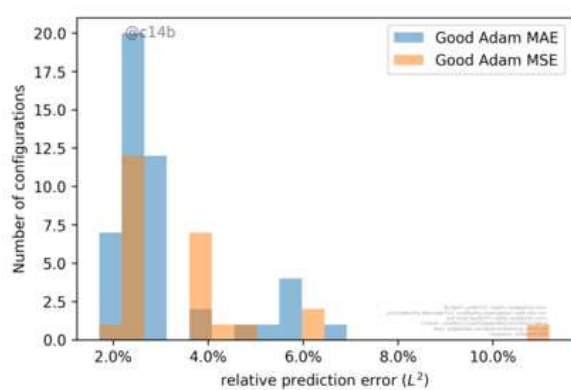


(b) Drag

Figure 7: Histograms of the relative prediction errors, comparing *Good ADAM* and *bad ADAM*, for the flow past airfoils.



(a) Lift



(b) Drag

Figure 8: Histograms for the relative prediction errors, comparing the mean absolute error and mean square error loss functions, for the flow past airfoils

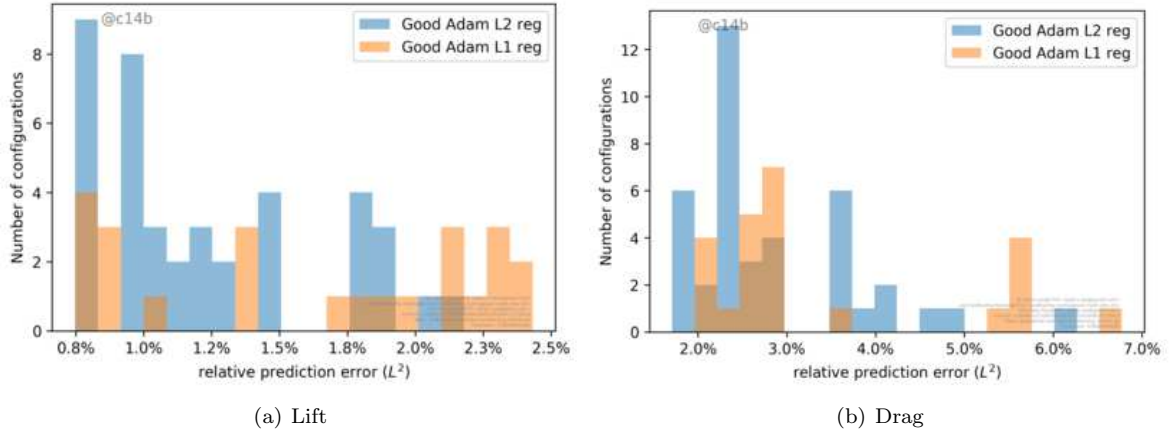


Figure 9: Histograms for the relative prediction error, comparing L^1 and L^2 regularizations of the loss function, for the flow past airfoils

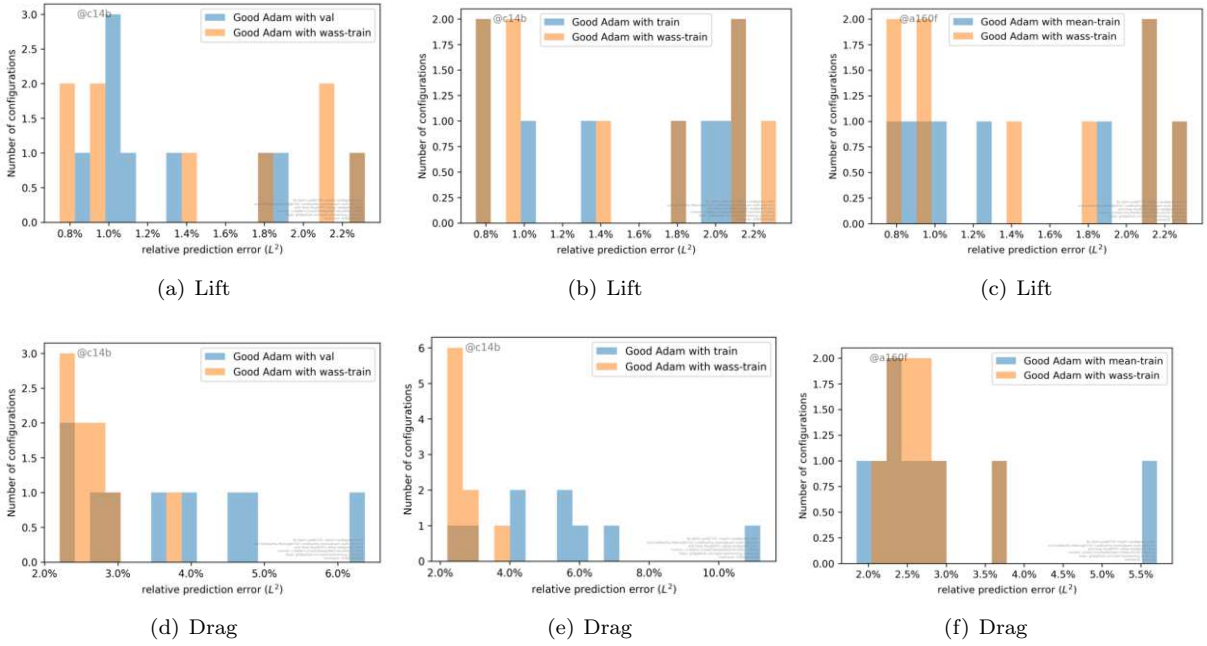


Figure 10: Network performance with respect to selection criteria for retrainings (see section 4 i.e., *train, val, mean-train, wass-train* for the flow past airfoils problem. We compare *wass-train* with each of the other three criteria. Histograms for prediction error (X-axis) with number of configurations (Y-axis) are shown.

Obs	Err (min)	Err (max)	Err (mean)	Err (std)
Lift	0.786	1.336	1.004	0.187
Drag	1.847	8.016	3.841	2.332

Table 6: Sensitivity of the best performing networks (listed in table 4) to retrainings i.e starting values for ADAM for flow past airfoil. All errors are relative mean L^2 prediction error (2.21) in percentage and we list the minimum, maximum, mean and standard deviation of the error over 5 retrainings.

Width/Depth	6	12	24
4	[1.21, 4.37]	[0.83, 1.74]	[1.09, 2.48]
8	[0.84, 3.66]	[0.86, 2.02]	[0.81, 1.69]
16	[0.80, -]	[-, 74.72]	[-, 5.11]

Table 7: Performance of trained neural networks with respect to the variation of network size for the flow past airfoils problem. The rows represent variation with respect to Width (number of neurons per hidden layer) and the columns represent variation with respect to Depth (number of hidden layers). For each entry of the Width-Depth matrix, we tabulate the vector of relative percentage mean prediction error for the best performing networks. The components of each vector represent error in [Lift, Drag] and - is used to indicate that the training procedure failed for this particular configuration.




Figure 11: Variation of prediction error (Y-axis) with respect to the size of the regularization parameter λ in (2.12) (X-axis) for the flow past airfoils problems. The minimum and mean of prediction error (over all hyperparameter configurations) is shown




Figure 12: Relative percentage mean square prediction error (2.21) (Y-axis) for the neural networks approximating the flow past airfoils problem, with respect to number of training samples (X-axis). For each number of training samples, we plot the mean, minimum and maximum (over hyperparameter configurations) of the prediction error. Only the selected (optimal) retraining is shown.

computed with the best performing networks for each observable (listed in table 4) and observe very good qualitative agreement between the outputs of the DLQMC algorithm and the reference measure.

In order to quantify the gain in computational efficiency with the DLQMC algorithm over the baseline QMC algorithm in approximating probability distributions, we will compute the *speedup*. To this end, we observe from figure 14 that the baseline QMC algorithm converges to the reference probability measure with respect to the Wasserstein distance (3.9), at a rate of $\alpha = 0.81$, for both the lift and drag. Next, from table 5, we see that once the training samples have been generated, the cost of training the network and performing evaluations with the trained network is essentially free. Consequently, if the number of training samples is N , we compute a *raw speed-up* (for each observable) given by,

$$\sigma_{i,dlqmc}^{raw} := \frac{W_1(\hat{\mu}_{i,qmc,N}, \hat{\mu}_i^\Delta)}{W_1(\hat{\mu}_{i,qmc}^*, \hat{\mu}_i^\Delta)}. \quad (5.5)$$

Here for $i = \text{Lift, Drag}$, $\hat{\mu}_i^\Delta$ is the reference measure (probability distribution) computed from the test set \mathcal{T} , $\hat{\mu}_{i,qmc,N}$ is the measure (3.3) with the baseline QMC algorithm and N QMC points, and $\hat{\mu}_{i,qmc}^*$ is the measure computed with the DLQMC algorithm (3.11). The real speed up is then given by,

$$\sigma_{i,dlqmc}^{real} = \left(\sigma_{i,dlqmc}^{raw}\right)^{\frac{1}{\alpha_i}}, \quad (5.6)$$

in order to compensate for the rate of convergence of the baseline QMC algorithm. In other words, our *real speed up* compares the costs of computing errors in the Wasserstein metric, that are of the same magnitude, with the DLQMC algorithm vis a vis the baseline QMC algorithm. Henceforth, we denote σ_{dlqmc}^{real} as the *speedup* and list the speedups (for each observable) with the DLQMC algorithm for $N = 128$ training samples in table 8. We observe from this table that we obtain speedups ranging between half an order to an order of magnitude for the lift and the drag. The speedups for drag are slightly better than that for lift.

In figure 15, we plot the speedup of the DLQMC algorithm (over the baseline QMC algorithm) as a function of the number of training samples in the deep learning algorithm 2.1. As seen from the figure, the best speed ups are obtained in case of 64 training samples for the lift and 128 training samples for the drag. This is on account of a complex interaction between the prediction error which decreases (rather fast) with the number of samples and the fact that the errors with the baseline QMC algorithm also decay with an increase in the number of samples.

5.1.7 Comparison with Monte Carlo algorithms

We have computed a Monte Carlo approximation of the probability distributions of the lift and the drag by randomly selecting $N = 128$ points from the parameter domain $Y = [0, 1]^6$ and computing the probability measure $\hat{\mu}_{mc} = \frac{1}{N} \sum_{j=1}^N \delta_{\mathcal{L}^\Delta(y_j)}$. We compute the Wasserstein error $W_1(\hat{\mu}_{mc}, \hat{\mu}^\Delta)$, with respect to the reference measure $\hat{\mu}^\Delta$ and divide it with the error obtained with the DLQMC algorithm to obtain a *raw speedup* of the DLQMC algorithm over MC. As MC errors converge as a square root of the number of samples, the *real speedup* of DLQMC over MC is calculated by squaring the raw speedup. We present this real speedup over MC in table 9 and observe from that the speedups of the DLQMC algorithm over the baseline MC algorithm are very high and amount to at least two orders of magnitude. This is not surprising as the baseline QMC algorithm significantly outperforms the MC algorithm for this problem. Given that the DLQMC algorithm was an order of magnitude faster than the baseline QMC algorithm, the cumulative effect leads to a two orders of magnitude gain over the baseline MC algorithm.

5.1.8 Choice of training set.

For the sake of comparison with our choice of Sobol points to constitute the training set \mathcal{S} in this problem, we chose $N = 128$ random points in Y as the training set. With this training set, we repeated the ensemble training procedure and found the best performing networks. The relative percentage mean L^2 -prediction error (2.21) (with respect to a test set of 320 random points) was computed and is presented

in table 10. As seen from this table, the prediction errors with respect to the best performing networks are considerably (an order of magnitude) higher than the corresponding errors obtained for the QMC training points (compare with table 4). Hence, at least for this problem, Sobol points are a much better choice for the training set than random points. An intuitive reason for this could lie in the fact that the Sobol points are better distributed in the parameter domain than random points. So on an average, a network trained on them will generalize better to unseen data, than it would if trained on random points.

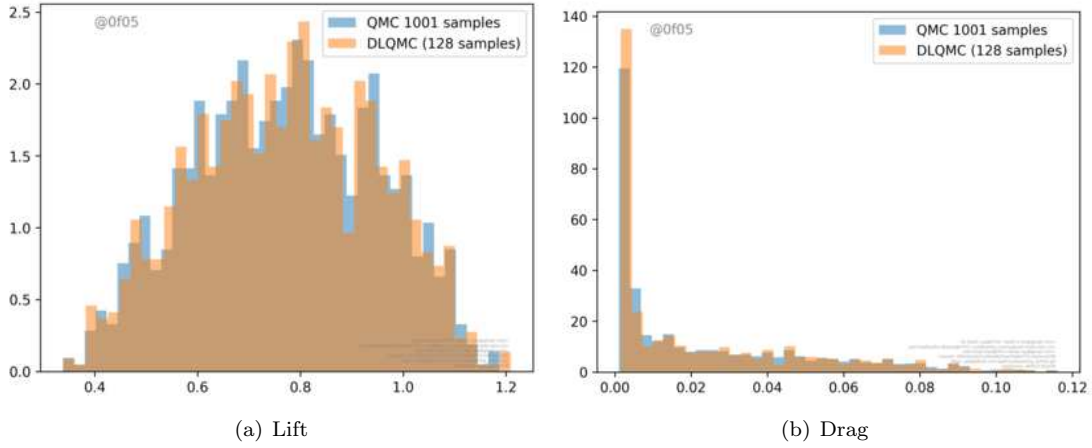


Figure 13: Empirical histograms representing the probability distribution (measure) for the lift and the drag in the flow past airfoils problem. We compare the reference histograms (computed with the test set \mathcal{T}) and the histograms computed with the DLQMC algorithm.

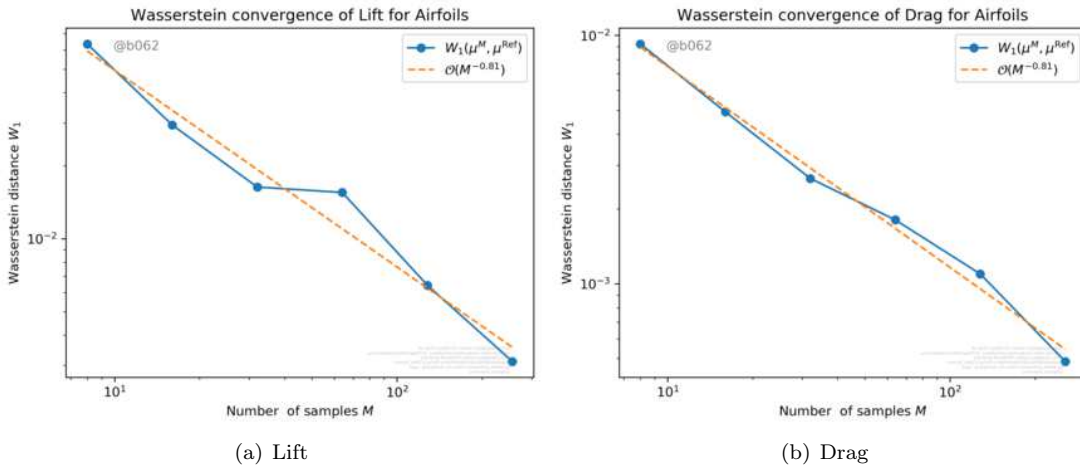


Figure 14: Convergence of the Wasserstein distance $W_1(\hat{\mu}_{i,qmc}, \hat{\mu}_i^\Delta)$ (Y-axis) for the of the lift and the drag for the flow past airfoils, with respect to number of QMC points (X-axis)

5.2 A stochastic shock tube problem

As a second numerical example, we consider a shock tube problem with the one-dimensional version of the compressible Euler equations (5.1).

Observable	Speedup
Lift	6.64
Drag	8.56

Table 8: Real speedups (5.6), for the lift and the drag in the flow past airfoils problem, comparing DLQMC with baseline QMC algorithm.

Observable	Speedup over MC
Lift	246.02
Drag	179.54

Table 9: Real speedups, for the lift and the drag in the flow past airfoils problem, comparing DLQMC with the baseline MC algorithm.




Figure 15: Real speedups (5.6) for the DLQMC algorithm over the baseline QMC algorithm (Y-axis) with respect to number of training samples (X-axis) for the flows past airfoils problem. The maximum and mean speed up (over all hyperparameter configurations) are shown.

Obs	Opt	Loss	L^1 -reg	L^2 -reg	Selection	BP . Err mean
Lift	SGD	MAE	0.0	7.8×10^{-7}	mean-train	8.487
Drag	ADAM	MAE	7.8×10^{-7}	0.0	train	20.25

Table 10: The hyperparameter configurations that correspond the best performing network (one with least mean prediction error) for the two observables of the flow past airfoils, trained on random (Monte Carlo) training set.

5.2.1 Problem description

The underlying computational domain is $[-5, 5]$ and initial conditions are prescribed in terms of left state (ρ_L, u_L, p_L) and a right state (ρ_R, u_R, p_R) at the initial discontinuity x_0 . These states are defined as,

$$\begin{aligned} \rho_L &= 0.75 + 0.45G_1(y) & \rho_R &= 0.4 + 0.3G_2(y) \\ w_L &= 0.5 + 0.5G_3(y) & w_R &= 0 \\ p_L &= 2.5 + 1.6G_4(y) & p_R &= 0.375 + 0.325G_5(y) \\ x_0 &= 0.5G_6(y). \end{aligned} \tag{5.7}$$

$y \in Y = [0, 1]^6$ is a parameter and $G_k(y) = 2y_k - 1$ for $k = 1, \dots, 6$. We use the Lebesgue measure on Y as the underlying measure.

The observables for this problem are chosen as the average integral of the density over fixed intervals

$$\mathcal{L}_j(y) = \frac{1}{|I_j|} \int_{I_j} \rho(x, T_f; y) dx, \quad j = 1, 2, 3, \tag{5.8}$$

where $I_1 = [-1.5, -0.5]$, $I_2 = [0.8, 1.8]$, $I_3 = [2, 3]$ are the underlying regions of interest.

A priori, this problem appears to be simpler as it is only one-dimensional when compared to the two-dimensional airfoil. However, the difficulty stems from the *large variance* of the initial data. As seen from (5.7), the ratio of the standard deviation to the mean, for the underlying variables, is very high and can be as high as 100%. This is almost an order of magnitude larger than the corresponding initial stochasticity of the flow past airfoil problem and we expect this large initial variance to be propagated into the solution. This is indeed borne out by the results shown in figure 16, where we plot the mean (and standard deviation) of the density, both initially and at time $T_f = 1.5$. Moreover, the initial stochastic inputs span a wide range of possible scenarios. Particular choices of the parameters i.e., $(y_1, y_2, y_3, y_5, y_6) = (7/9, 1/24, 0, 1/32, 1/13, 0)$ and $(y_1, y_2, y_3, y_5, y_6) = (0.1611, 2/3, 0.698, 0.8213, 0.8015, 0)$, correspond to the classical Sod shock tube and Lax shock tube problems [16]. Thus, we can think of this stochastic shock tube problem as more or less generic (scaled) Riemann problem for the Euler equations.

In a recent paper [32], the authors have analyzed the role of the underlying variance in the generalization errors for deep neural networks trained to regress functions. In particular, high variance can lead to large generalization errors. In this sense, this shock tube problem might be harder to predict with neural networks than the flow past airfoil.

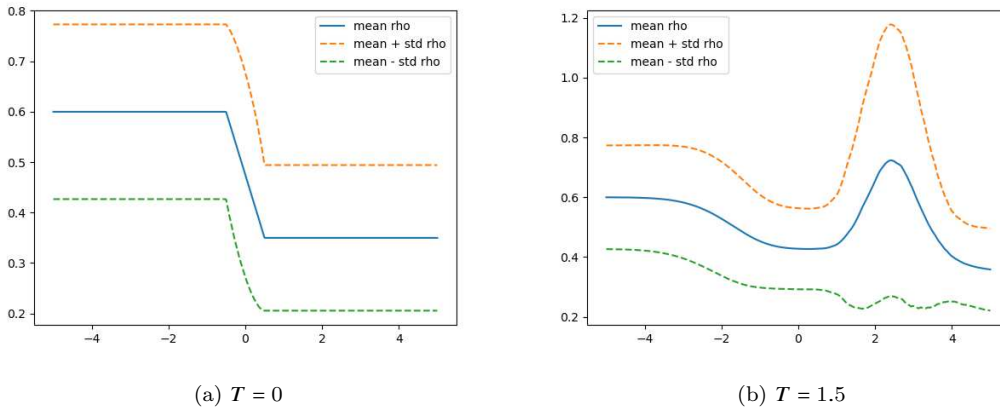


Figure 16: Mean (\pm standard deviation) for the density of the stochastic shock tube problem.

5.2.2 Generation of training data.

In order to generate the training and test sets, we denote Q^{sob} as the first 16256 Sobol points on the domain $[0, 1]^6$. The training set \mathcal{S} consists of the first 128 of these Sobol points whereas as the test

set is $\mathcal{T} = \mathcal{Q}^{sob}$. For each point $y_j \in \mathcal{Q}^{sob}$, the maps or (rather their numerical surrogate) $\mathcal{L}_{1,2,3}^\Delta(y_j)$ are generated from high-resolution numerical approximations $\mathbf{U}^\Delta(y)$ obtained using a second-order finite volume scheme. The domain is discretized into a uniform mesh of $K = 2048$ disjoint intervals, with the HLL numerical flux and with the left and right cell-interface values obtained using a second-order WENO reconstruction. Time marching is performed using the second-order Runge-Kutta method with a CFL number of 0.475. Once the finite volume solution is computed, the corresponding observables (5.8) at each y are approximated as

$$\mathcal{L}_j^\Delta(y) = \frac{\sum_{i \in \Lambda_j} \rho_i(y)}{\#(\Lambda_j)}, \quad \Lambda_j = \{i \mid x_i \in I_j\}, \quad j = 1, 2, 3, \quad (5.9)$$

where x_i is the barycenter of the cell Ω_i .

5.2.3 Results.

We follow the ensemble training procedure outlined in section 4. The network architecture, specified in table 3 is used together with the 114 hyperparameter configurations of the previous section, with 5 retrainings for each hyperparameter configuration. The best performing networks for each observable are identified as the ones with lowest prediction error (2.21) and these networks are presented in table 11. From the table, we observe very slight differences for the best performing networks for the three observables. However, these networks are indeed different from the best performing networks for the flow past airfoil (compare with table 4). Moreover, the predictions errors are larger than the flow past airfoil case, ranging between 2% for the observable \mathcal{L}_1 to approximately 6% for the other two observables. This larger error is expected as the variance of the underlying maps is significantly higher. Nevertheless, a satisfactory prediction accuracy is obtained for only 128 training samples. Results of the ensemble training

Obs	Opt	Loss	L^1 -reg	L^2 -reg	Selection	BP. Err mean
\mathcal{L}_1	ADAM	MSE	7.8×10^{-6}	0	wass-train	2.212
\mathcal{L}_2	ADAM	MSE	7.8×10^{-6}	0	wass-train	6.575
\mathcal{L}_3	ADAM	MSE	7.8×10^{-5}	0	mean-train	5.467

Table 11: The hyperparameter configurations that correspond to the best performing network (one with least mean prediction error) for the three observables of the stochastic shock tube problem.

procedure are depicted in the form of histograms for the prediction errors in figure 17. From this figure, we observe that although there are some outliers, most of the hyperparameter configurations resulted in errors comparable to the best performing networks (see table 11). Furthermore, a large majority of the hyperparameters resulted in substantially (a factor of 3 – 4) smaller prediction errors, when compared to the linear least squares regression (5.4).

A sensitivity study (not presented here) indicates very similar robustness results as for the flow past the airfoil. In figure 18, we plot the prediction error with respect to the number of samples and observe that the mean error (over hyperparameter configurations) decays with an exponent of 0.5 for the observables $\mathcal{L}_{2,3}$. On the other hand, the decay is a bit slower for \mathcal{L}_1 . However, the low value of the constants still indicates significant compression allowing us to approximate this rather intricate problem with accurate neural networks.

Finally, we compute the underlying probability distributions for each observable with the DLQMC algorithm 3.2. In figure 19, we plot the maximum and mean real speedup (5.6), over hyperparameter configurations, for each observable, with respect to the number of training samples. From this figure, we observe maximum speedups of about 16 for the observables $\mathcal{L}_{1,2}$ and 5 for the observable \mathcal{L}_3 , indicating that the DLQMC algorithm is significantly superior to the baseline QMC algorithm, even for this problem.




Figure 17: Histograms depicting distribution of the percentage relative mean L^2 prediction error (2.21) (X-axis) over number of hyperparameter configurations (samples, Y-axis) for the observables of the stochastic shock-tube problem.




Figure 18: Relative percentage mean square prediction error (2.21) (Y-axis) for the neural networks approximating the stochastic shock tube problem, with respect to number of training samples (X-axis). For each number of training samples, we plot the mean, minimum and maximum (over hyperparameter configurations) of the prediction error. Only the selected (optimal) retraining is shown.




Figure 19: Real speedups (5.6) for the DLQMC algorithm over the baseline QMC algorithm (Y-axis) with respect to number of training samples (X-axis) for the stochastic shock tube problem. The maximum and mean speed up (over all hyperparameter configurations) are shown.

6 Discussion

Problems in CFD such as UQ, Bayesian Inversion, optimal design etc are of the *many query* type i.e, solving them requires evaluating a large number of realizations of the underlying input parameters to observable map (2.3). Each realization involves a call to an expensive CFD solver. The cumulative total cost of the many-query simulation can be prohibitively expensive.

In this paper, we have tried to harness the power of machine learning by proposing using *deep fully connected neural networks*, of the form (2.6), to learn and predict the underlying parameters to observable map. However, a priori, this constitutes a formidable challenge. In section 2.3, we argue by a combination of theoretical considerations and numerical experiments that the crux of this challenge is to find neural networks to learn maps of *low regularity in a data poor regime*. We are in this regime as the many parameters to observable maps (2.3) in fluid dynamics, can be at best, Lipschitz continuous. Moreover, given the cost, we can only afford to compute a few training samples.

We overcome these difficulties with the following *novel* ideas, first, we chose to focus on learning observables rather than the full solution of the parametrized PDE (2.1). Observables are more regular than fields (see section 2.3) and might have less underlying variance and are thus, easier to learn, see [32] for a more recent theoretical justification. Second, we propose using *low-discrepancy sequences* to generate the training set \mathcal{S} . The equi-distribution property of these points might ensure lower generalization errors than those resulting from training data generated at random points, see the forthcoming paper [34] for further details. Finally, we propose an *ensemble training procedure* (see section 4) to systematically scan the hyperparameter space in order to winnow down the best performing networks as well as to test sensitivity of the trained networks to hyperparameter choices.

We presented two prototypical numerical experiments for the compressible Euler equations in order to demonstrate the efficacy of our approach. We observed from the numerical experiments that the trained networks were indeed able to provide low prediction errors, despite being trained on very few samples. Moreover, a large number of hyperparameter configurations were close in performance to the best performing networks, indicating a significant amount of robustness to most hyperparameters. More crucially, the obtained generalization errors followed the theoretical predictions of section 2.3 and indicating a high degree of compression, explaining why the networks generalized well. It is essential to state that these low prediction errors are obtained even when the cost of evaluating the neural network is several orders of magnitude lower than the full CFD solve.

As a concrete application of our deep learning algorithm 2.1, we combined it with a Quasi-Monte Carlo (QMC) method to propose a deep learning QMC algorithm 3.2, for performing forward UQ, by approximating the underlying probability distribution (3.1) of the observable. In theorem 3.3, we prove that the DLQMC algorithm is guaranteed to out-perform the baseline QMC algorithm as long as the underlying neural network generalizes well. This is indeed borne out in the numerical experiments, where we observed about an order of magnitude speedup of the DLQMC algorithm over the corresponding QMC algorithm and two orders of magnitude speedup over a Monte Carlo algorithm.

Thus, we have demonstrated the viability of using neural networks to learn observables and to solve challenging problems of the many-query type. Our approach can be compared to a standard model order reduction algorithm [52]. Here, generating the training data and training the networks is the *offline* step whereas evaluating the network constitutes the *online* step. Our results compare favorably with attempts to use Model order reduction for hyperbolic problems [1, 9], particularly when it comes to the low costs of training and evaluation. On the other hand, typical MOR algorithms provide a surrogate for the solution field, where we only predict the observable of interest in the approach presented here.

Our results can be extended in many directions; we can adapt this setting to learn the full solution field, instead of the observables. However, this might be harder in practice as the field is even less regular than the observable. Our approach is entirely data driven, in the sense that we do not use any specific information about the underlying parameterized PDE (2.1). Thus, the approach can be readily extended to other forms of (2.1), for instance the Navier-Stokes equations or other elliptic and parabolic PDEs. Finally, we plan to apply the trained neural networks in the context of Bayesian inverse problems and shape optimization under uncertainty, in forthcoming papers.

Acknowledgements

The research of SM is partially support by ERC Consolidator grant (CoG) NN 770880 COMANFLO. A large proportion of computations for this paper were performed on the ETH compute cluster EULER.

References

- [1] R. Abgrall and R. Crisovan. Model reduction using L1-norm minimization as an application to nonlinear hyperbolic problems. *Internat. J. Numer. Methods Fluids* 87 (12), 2018, 628651.
- [2] Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 254-263. PMLR, Jul 2018.
- [3] A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Trans. Inform. Theory*, 39(3), 930-945, 1993.
- [4] H. Bijl, D. Lucor, S. Mishra and Ch. Schwab. (editors). *Uncertainty quantification in computational fluid dynamics.*, Lecture notes in computational science and engineering 92, Springer, 2014.
- [5] H. Bölcskei, P. Grohs, G. Kutyniok, and P. Petersen. Optimal approximation with sparsely connected deep neural networks. *arXiv preprint*, available from arXiv:1705.01714, 2017.
- [6] A. Borzi and V. Schulz. *Computational optimization of systems governed by partial differential equations.* SIAM (2012).
- [7] H. J. Bungartz and M. Griebel. Sparse grids. *Acta Numer.*, 13, 2004, 147-269.
- [8] R. E. Caflisch. Monte Carlo and quasi-Monte Carlo methods. *Acta. Numer.*, 1, 1988, 1-49.
- [9] Crisovan, R.; Torlo, D.; Abgrall, R.; Tokareva, S. Model order reduction for parametrized nonlinear hyperbolic problems as an application to uncertainty quantification. *J. Comput. Appl. Math.* 348 (2019), 466489
- [10] G. Cybenko. Approximations by superpositions of sigmoidal functions. *Approximation theory and its applications.*, 9 (3), 1989, 17-28
- [11] Constantine M. Dafermos. *Hyperbolic Conservation Laws in Continuum Physics (2nd Ed.)*. Springer Verlag (2005).
- [12] W. E and B. Yu. The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems. *Commun. Math. Stat.* 6 (1), 2018, 1-12.
- [13] W. E, J. Han and A. Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Commun. Math. Stat.* 5 (4), 2017, 349-380.
- [14] W. E, C. Ma and L. Wu. A priori estimates for the generalization error for two-layer neural networks. *ArXIV preprint*, available from arXiv:1810.06397, 2018.
- [15] R. Evans et. al. De novo structure prediction with deep learning based scoring. *Google DeepMind working paper*, 2019.
- [16] U. S. Fjordholm, S. Mishra and E. Tadmor, Arbitrarily high-order accurate entropy stable essentially nonoscillatory schemes for systems of conservation laws. *SIAM J. Numer. Anal.*, 50 (2012), no. 2, 544573.
- [17] U. S. Fjordholm, R. Käppeli, S. Mishra and E. Tadmor, Construction of approximate entropy measure valued solutions for hyperbolic systems of conservation laws. *Found. Comput. Math.*, 17 (3), 2017, 763-827.

- [18] U. S. Fjordholm, K. O. Lye, S. Mishra and F. R. Weber, Numerical approximation of statistical solutions of hyperbolic systems of conservation laws. *In preparation*, 2019.
- [19] R. Ghanem, D. Higdon and H. Owhadi (eds). *Handbook of uncertainty quantification*, Springer, 2016.
- [20] I. Goodfellow, Y. Bengio and A. Courville. *Deep learning*, MIT Press, (2016), available from <http://www.deeplearningbook.org>
- [21] Edwige Godlewski and Pierre A. Raviart. *Hyperbolic Systems of Conservation Laws*. Mathematiques et Applications, Ellipses Publ., Paris (1991).
- [22] I. G. Graham, F. Y. Kuo, J. A. Nichols, R. Scheichl, Ch. Schwab and I. H. Sloan. Quasi-Monte Carlo finite element methods for elliptic PDEs with lognormal random coefficients. *Numer. Math.* 131 (2), 2015, 329368.
- [23] J. Han, A. Jentzen and W. E. Solving high-dimensional partial differential equations using deep learning. *PNAS*, 115 (34), 2018, 8505-8510.
- [24] J S. Hesthaven. *Numerical methods for conservation laws: From analysis to algorithms*. SIAM, 2018.
- [25] C. Hirsch, D. Wunsch, J. Szumbarksi, L. Laniewski-Wollk and J. Pons-Prats (editors). *Uncertainty management for robust industrial design in aeronautics* Notes on numerical fluid mechanics and multidisciplinary design (140), Springer, 2018.
- [26] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5), 359-366, 1989.
- [27] Diederik P. Kingma and Jimmy Lei Ba. Adam: a Method for Stochastic Optimization. *International Conference on Learning Representations*, 1-13, 2015.
- [28] I. E. Lagaris, A. Likas and D. I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*, 9 (5), 1998, 987-1000.
- [29] L. D. Landau and E. M. Lipschitz. Fluid Mechanics, 2nd edition, *Butterworth Heinemann*, 1987, 532 pp.
- [30] Y. LeCun, Y. Bengio and G. Hinton. Deep learning. *Nature*, 521, 2015, 436-444.
- [31] R. J. LeVeque. *Finite difference methods for ordinary and partial differential equations, steady state and time dependent problems*. SIAM (2007).
- [32] K. O. Lye, S. Mishra and R. Molinaro. A Multi-level procedure for enhancing accuracy of machine learning algorithms. *Preprint*, 2019, available from ArXiv:1909:09448.
- [33] S. Mishra. A machine learning framework for data driven acceleration of computations of differential equations, *Math. in Engg.*, 1 (1), 2018, 118-146.
- [34] S. Mishra and T. K. Rusch. Enhancing accuracy of machine learning algorithms by training on low-discrepancy sequences. *In preparation*, 2020.
- [35] T. P. Miyanawala and R. K. Jaiman. An efficient deep learning technique for the Navier-Stokes equations: application to unsteady wake flow dynamics. *Preprint*, 2017, available from arXiv:1710.09099v2.
- [36] H. N. Mhaskar. Neural networks for optimal approximation of smooth and analytic functions, *Neural Comput.*, 8 (1), 1996, 164-177.
- [37] J. Munkres. Algorithms for the Assignment and Transportation Problems *Journal of the Society for Industrial and Applied Mathematics*, 5 (1), 1957, 32-38.

- [38] Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. Towards understanding the role of over-parametrization in generalization of neural networks. *arXiv preprint arXiv:1805.12076*, 2018.
- [39] P. Petersen and F. Voightlaender. Optimal approximation of piecewise smooth functions using deep ReLU neural networks. *ArXIV preprint*, available from arXiv:1709.05289, 2018.
- [40] M. Raissi and G. E. Karniadakis. Hidden physics models: machine learning of nonlinear partial differential equations. *J. Comput. Phys.*, 357, 2018, 125-141.
- [41] M. Raissi, A. Yazdani and G. E. Karniadakis. Hidden fluid mechanics: A Navier-Stokes informed deep learning framework for assimilating flow visualization data. *ArXIV preprint*, available from arXiv:1808.04327, 2018.
- [42] D. Ray and J. S. Hesthaven. An artificial neural network as a troubled cell indicator. *J. Comput. Phys.*, 367, 2018, 166-191
- [43] D. Ray, P. Chandrasekhar, U. S. Fjordholm and S. Mishra. Entropy stable scheme on two-dimensional unstructured grids for Euler equations, *Commun. Comput. Phys.*, 19 (5), 2016, 1111-1140.
- [44] D. Ray. Entropy-stable finite difference and finite volume schemes for compressible flows, *Doctoral thesis*, 2017, available from <https://deepray.github.io/thesis.pdf>
- [45] S. Ruder. An overview of gradient descent optimization algorithms. Preprint, 2017, available from arXiv:1609.04747v2.
- [46] C. Schwab and J. Zech. Deep learning in high dimension. *Technical Report 2017-57*, Seminar for Applied Mathematics, ETH Zürich, 2017.
- [47] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [48] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19, 2010, 451-559.
- [49] J. Tompson, K. Schlachter, P. Sprechmann and K. Perlin. Accelerating Eulerian fluid simulation with convolutional networks. *Preprint*, 2017. Available from arXiv:1607.03597v6.
- [50] L. N. Trefethen. *Spectral methods in MATLAB*, SIAM, (2000).
- [51] F. Tröltzsch. *Optimal control of partial differential equations*. AMS, (2010).
- [52] A. Quateroni, A. Manzoni and F. Negri. *Reduced basis methods for partial differential equations: an introduction*, Springer Verlag (2015).
- [53] C. Villani. *Topics in Optimal Transportation*. American Mathematical Society, Graduate Studies in Mathematics, Vol. 58 (2013)
- [54] D. Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94, 2017, 103-114
- [55] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *In International Conference on Learning Representations*, 2017.
- [56] Linfeng Zhang, De-Ye Lin, Han Wang, Roberto Car, and Weinan E. Active Learning of Uniformly Accurate Inter-atomic Potentials for Materials Simulation. *ArXIV preprint*, available from arXiv:1810.11890, 2018.
- [57] Chollet, François and others Keras <https://keras.io>, 2015.
- [58] MartinAbadi et al, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems <https://www.tensorflow.org/>, 2015.