Shock Wave Prediction in Transonic Flow Fields using Domain-Informed Probabilistic Deep Learning

Bilal Mufti,¹ Anindya Bhaduri^{*},^{2, a)} Sayan Ghosh,² Liping Wang,² and Dimitri N. Mavris¹

¹⁾Daniel Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, Georgia, 30332, USA

²⁾Probabilistic Design & Materials Informatics, GE Aerospace Research, Niskayuna, New York, 12309, USA

(*Electronic mail: anindya.bhaduri@ge.com)

(Dated: 12 December 2023)

Transonic flow fields are marked by shock waves of varying strength and location and are crucial for the aerodynamic design and optimization of high-speed transport aircraft. While deep learning methods offer the potential for predicting these fields, their deterministic outputs often lack predictive uncertainty. Moreover, their accuracy, especially near critical shock regions, needs better quantification. In this paper, we introduce a domain-informed probabilistic (DIP) deep learning framework tailored for predicting transport flow fields with shock waves called DIP-ShockNet. This methodology utilizes Monte Carlo Dropout (MCD) to estimate predictive uncertainty and enhances flow field predictions near the wall region by employing the inverse wall distance function (IWDF) based input representation of the aerodynamic flow field. The obtained results are benchmarked against the signed distance function (SDF) and the geometric mask input representations. The proposed framework further improves prediction accuracy in shock wave areas using a domain-informed loss function. To quantify the accuracy of our shock wave predictions, we developed metrics to assess errors in shock wave strength and location, achieving errors of 6.4% and 1%, respectively. Assessing the generalizability of our method, we tested it on different training sample sizes and compared it against the proper orthogonal decomposition (POD)-based reduced order model (ROM). Our results indicate that DIP-ShockNet outperforms POD-ROM by 60% in predicting the complete transonic flow field.

Keywords: transonic flow field, shock wave, probabilistic deep learning, uncertainty quantification, UNet, domaininformed loss, Monte Carlo dropout

I. INTRODUCTION

High-speed transport aircraft design remains a principal strategic thrust within the aviation community^{1,2}. Such aircraft are designed to operate efficiently at transonic speeds. This necessitates improvements in cruise efficiency and drag divergence. Consequently, the aerodynamic design of these aircraft must undergo rigorous analysis and evaluation for the transonic flow regime during the early conceptual design phase. Transonic flow fields are inherently complex, marked by variations in shock wave strength and position, and intricate interactions between shock waves and boundary layers. In the aerospace industry, Reynolds-averaged Navier-Stokes (RANS) based computational fluid dynamics (CFD) analysis is the conventional approach for evaluating high-speed compressible flow fields. However, for aerodynamic design, CFD analysis often entails a many-query context where the analysis is reiterated numerous times. Such many-query analyses include numerical optimization, uncertainty quantification, design space exploration, and more^{3,4}. Resorting to computationally intensive RANS CFD analysis in a many-query context quickly becomes computationally intractable, despite the current advancements in computational capabilities. This underscores the need for methods that can rapidly and efficiently evaluate flow fields.

a) corresponding author

Non-intrusive, data-driven ROMs offer a promising alternative to high-fidelity CFD simulations⁵. They provide a computationally efficient approximation of complex numerical simulations and, once trained, can quickly estimate highdimensional flow fields. ROMs operate by determining a low-dimensional representation, often referred to as the latent space or subspace, of the original high-dimensional field solution. POD stands out as a widely utilized ROM technique, seeing extensive application in aerodynamics and aircraft design^{b-12}. The appeal of POD stems from its simple formulation, reduced computational requirements, minimal training data needs, and a well-defined back mapping. However, its limitations surface when accurately predicting nonlinearities such as shock waves within the flow field^{5,13,14}.

Recently, many studies within fluid mechanics and aerodynamics have turned to deep learning (DL) methods to develop ROMs for highly non-linear systems and capture the complex relationships between inputs and outputs. These techniques have been used to study unsteady flows around bluff bodies^{15–20}, address turbulence in flows^{21–23}, heat transfer modeling^{24,25} and reconstruction of pressure fluctuations in supersonic shock–boundary layer interaction²⁶. DL has also found applications in examining flow fields around high-speed aerodynamic bodies. Thuerey *et al.*²² used the UNet architecture to predict pressure and velocity fields on various airfoils for different Reynolds numbers and angles of attack. Bhatnagar *et al.*²⁷ improved their convolutional neural network (CNN) predictions on airfoils by using the SDF, introduced by Guo *et al.*²⁸, as an input channel. Their approach utilized

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0185370

single and shared decoder streams for pressure and velocity field predictions. In their work, Duru et al.^{29,30} presented a detailed solution for CNN-based predictions of flow fields around transonic airfoils, using a training dataset of 4614 samples from 204 different airfoils at M = 0.7 across various angles of attack. They achieved notable accuracy in predicting flow fields and shock regions. However, a common theme across these studies was the mapping of the original CFD grid to a Cartesian grid for neural network (NN) model training. This method can sometimes lead to interpolation errors and loss of information in near-wall regions since the Cartesian grid might not represent the geometry boundary accurately, leading to oscillations in pressure field close to aerodynamic body boundaries³¹.

The mesh transformation technique can be employed to mitigate the challenges posed by mapping onto a Cartesian grid³². Hu and Zhang³³ introduced a Mesh-Conv convolutional operator. They transformed the O-type grid into a Cartesian grid, which was then used as an input to the CNN. Their model was tested to predict the flow field around the NACA 0012 airfoil at a low subsonic Mach number of 0.5. In a similar vein, Chen and Thuerey³⁴ applied an encoding technique to convert a body-fitted, structured grid from Cartesian coordinates to curvilinear coordinates. Following this transformation, a U-Net architecture was used to analyze flow fields around various airfoils under different flow conditions. Deng et al.35 devised a vision transformer (ViT)-based architecture, which utilized a transformed-encoded geometric input. Their approach achieved good prediction accuracy for flow around transonic airfoils using a tailored loss function. However, a common constraint with these methods is the need for the underlying CFD grid to possess a specific topological order. This stipulation restricts their applicability to grids with a structured representation. Generating these structured grids can be particularly challenging, especially for aerodynamic bodies with complex shapes.

Graph neural networks (GNNs) have recently emerged as a tool for predicting flow fields across diverse grids^{36,37}. Pfaff et al.38 introduced MeshGraphNets, a GNN-based approach, for predicting inviscid flow fields over airfoils. However, the utilization of GNNs faces constraints related to the theory of graph representations, hardware limitations, and challenges in distributed solving^{34,35}. In a notable contribution, Kashefi et al.39 developed a deep learning framework that leverages the PointNet architecture to capture point cloud locations of grid vertices and use them as network inputs. This approach enabled them to predict flow fields in irregular domains where the field solution is influenced by the object's shape. After testing their network on a range of simple shapes, they achieved impressive results. Yet, their methodology has not been evaluated for predicting high-speed flow fields on dense, unstructured grids. It's important to note that the studies discussed so far adopt a deterministic approach, predicting only the mean values of flow fields and their associated errors. This poses limitations for applications where aircraft designers must also account for uncertainties in their designs. The existing studies also do not explicitly quantify errors in shock waves regions and the performance of the methods in regions of shock waves is analyzed qualitatively.

In the present study, we introduce a probabilistic deep learning framework DIP-ShockNet, that utilizes a domaininformed Monte Carlo Dropout (MCD)-UNet architecture tailored for predicting transonic flow fields with an emphasis on capturing shock wave regions accurately. This approach enables us to not only provide mean predictions of flow fields but also to quantify the uncertainty associated with these predictions by estimating the variance of the predicted flow fields. For input structuring, we employ the Cartesian grid mapping strategy. While acknowledging its inherent limitations, this strategy proves advantageous in our context, because it enables us to swiftly estimate gradients, thereby refining shock wave predictions; as well as utilize information from highfidelity but highly unstructured CFD grids. Another potential advantage of the grid mapping strategy is that it can decrease training costs through a multi-fidelity approximation of the model by mapping grids of varying densities onto a common grid. To improve predictions near the wall, we use an inverse wall distance function (IWDF). This aids our model in better learning field distributions close to the aerodynamic surface, leading to better accuracy in near-wall predictions. Furthermore, we quantify shock wave errors by assessing shock location and strength. We evaluate the robustness of our model to the number of training samples by evaluating it for different training data sizes. We compare these results against those derived from POD-based models. The importance of comparing the outcomes from computationally demanding DL methods with simpler POD-based techniques has been welldocumented40.

The remainder of this paper is organized as follows: In Section II, the proposed DIP-ShockNet framework is discussed where we present the inverse wall distance function, detail the design of the MCD-UNet network architecture, and introduce the domain-informed loss function, all tailored to enhance the prediction accuracy of our model. Section III introduces the application problem used in this research and describes the shape parametrization technique employed to develop various designs. This section also sheds light on data generation, and preprocessing methodologies for creating training data, and introduces metrics devised to quantify errors in field predictions as well as errors in shock wave strength and location. Section IV examines the influence of input representation, the domain-informed loss function, and changes in training data size on our model's predictions. We wrap up in Section V by summarizing the research and offering concluding insights.

II. METHODOLOGY

In this section, we discuss in detail the proposed Domain-Informed Probabilistic ShockNet (DIP-ShockNet) deep learning framework used to predict the transonic flow fields with shock waves. The framework has three key components: the IWDF as the input representation of the aerodynamic flow field, the MCD-UNet architecture, and the domain-informed loss function. We start by introducing the IWDF and give its mathematical details. Then, we explain the structure and

design of the MCD-UNet architecture. We finish by introducing a customized loss function that uses domain knowledge to better predict the areas with shock waves.

A. Inverse Wall Distance Function

To improve prediction accuracy near aerodynamic surfaces in our proposed model, we introduce the *inverse wall distance function* (IWDF). The essence of IWDF lies in assigning higher weightage to regions where significant flow variations, such as in boundary layers, are anticipated.

Inspired by the *SDF* methodology, the first step is to represent the boundary of the aerodynamic surface using a zero-level set *Z* represented as:

$$Z = \{(i, j) \in \mathbb{R}^2 : \Omega(i, j) = 0\}$$
(1)

where $\Omega(i, j)$ is the level-set function defined over the domain of interest. The condition $\Omega(i, j) = 0$ defines the surface of the aerodynamic body. The nearest wall distance function, D(i, j), from surface of aerodynamic body is given by:

$$D(i,j) = \min_{(i',j') \in \mathbb{Z}} \left\| (i,j) - (i',j') \right\|_2$$
(2)

where $||.||_2$ denotes the L_2 or Euclidean norm. Subsequently, the IWDF is computed as

$$IWDF(i,j) = \left| 1 - \frac{D(i,j)}{\max(D(i,j)) + \varepsilon_{\rho}} \right| \operatorname{sign}(\Omega(i,j))$$
(3)

where

$$\operatorname{sign}(\Omega(i,j)) = \begin{cases} -1 & \Omega(i,j) < 0 \\ 0 & \Omega(i,j) = 0 \\ 1 & \Omega(i,j) > 0 \end{cases}$$
(4)

In Eq. 3, the offset $\varepsilon_o > 0$, ensures that the IWDF value does not reduce to zero at the domain's extremities. This distinction allows our model to differentiate sharply between the aerodynamic surface and the far-field regions. In this study we use a value of 0.05 for ε_o .

The efficacy and versatility of IWDF, especially in comparison with other prevalent geometric and distance representation methodologies, will be demonstrated in the section IV of this paper.

B. Design of MCD-UNet architecture

Deterministic surrogate models^{41–45} provide output predictions without any information about the uncertainty of the predictions. Probabilistic surrogate models^{46–48}, on the other hand, are useful in assessing model confidence on output predictions as they provide a measure of uncertainty. In deep learning models, uncertainty can be incorporated with Bayesian neural networks (BNNs)^{49,50} where the tunable network parameters are modeled as distributions rather than scalars. For most complex problems, BNNs are often infeasible to implement as the evaluation of the marginal distribution of the weights during training becomes computationally intractable.

Bayesian Variational Inference (VI) methods have been proposed as a practical alternative to BNNs. In VI-based deep neural networks, Gaussian distribution is commonly assumed as the marginal distribution for the weights. However, this requires double the number of parameters of a network to represent its uncertainty which leads to increased computational costs^{51,52}. It has been shown⁵² that the use of the regular dropout in neural networks can be interpreted as an efficient approximation to BNNs. Traditionally, dropout has been used only during model training as a regularization technique to avoid over-fitting. Monte Carlo Dropout (MCD) is a technique that proposes using dropout during model inference and this enables DL methods to represent model form uncertainty efficiently. The MCD layer randomly deactivates some of the neurons (weights) of a neural network layer, thus nullifying their contributions to the final output. Thus, during inference, when the same input is fed into a trained DL model multiple times, the outputs obtained are slightly different from each other due to the deactivation of a different set of neurons (weights) each time. The mean of the output predictions is then used as the estimate, and the variance as a measure of its uncertainty.

At the core of the DIP-ShockNet framework is the MCD-UNet architecture. UNet is a popular CNN architecture consisting of an encoder-decoder structure that learns complex image-to-image mapping efficiently. The role of the encoder is to apply a series of linear and non-linear operations to the given input image and come up with an abstract reduced-order feature representation. This intermediate reduced order output is then passed onto the decoder which through further sequential linear and non-linear operations reshapes it back to the output image. Further details about the UNet architecture can be found in the existing literature^{53–55}. The additional MCD feature in the MCD-UNet architecture enables the estimation of uncertainty in the prediction of the flow fields, along with the mean estimates.

Design of the MCD-UNet architecture with optimal selection of the hyperparameters is an important step in predicting the flow field accurately and efficiently. The training hyperparameters considered are the initial learning rate, the exponential decay rate constant of the learning rate, the dropout rate, and the training batch size while the architecture-based hyperparameters considered are the depth of the architecture, the number of channels at each depth level, and the type of normalization applied in each convolutional layer. A random sampling-based hyperparameter tuning is performed by generating 100 hyperparameter configurations. The MCD-UNet architecture is then trained for each of these configurations within the DIP-ShockNet framework and the optimal configuration is selected based on the validation data accuracy. Table I lists the hyperparameter configuration selected for our study. It is noted that the selected normalization type is group normalization with a group size of 256 which is equivalent to layer normalization as the maximum number of channels

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0185370

Accepted to Phys. Fluids 10.1063/5.0185370



FIG. 1. Architecture of the MCD-UNet used in the DIP-ShockNet framework.

used in the architecture is also 256. The optimal MCD-UNet architecture based on the selected configuration is shown in Fig. 1. The input to the deep learning architecture is a 2channel image representation that contains information about the airfoil geometry and the angle of attack. It is noted that a single-channel representation may also act as an input to this architecture as discussed later in section IV.

The encoder part of the architecture consists of 5 repeating blocks. Each block consists of a 2×2 max pooling operation with stride 2 for downsampling (MaxPool 2×2), followed by the application of 2 successive 3×3 2D convolutions (Conv 3×3), each followed by a layer normalization (LayerNorm) and a leaky rectified linear unit (LeakyReLU) operation. The only exception to this block structure is the first encoder block which does not have the max pooling layer upfront. Also, the last 3 encoder blocks have an additional MCD laver before the max pooling layer. The decoder part consists of 5 repeating blocks where each block (except the first and the last block) consists of two successive $3 \times 32D$ convolutions (Conv 3×3), each followed by a layer normalization (LayerNorm) and a ReLU operation, and it is then followed by a transpose convolution operation (TransposeConv3 \times 3). The first decoder block has only the transpose convolution layer and no standard convolution layers, while the last decoder block has the two successive Conv 3×3 -LayerNorm-ReLU layers but no transpose convolution. The first 3 decoder blocks have an additional MCD layer before the transpose convolution layer. The encoder and decoder blocks are followed by a final 1×1 convolutional layer (Conv 1×1) which maps the 64-channel decoder output to a single-channel output which represents the predicted flow field. It is noted that the training of the MCD-UNet architecture using the DIP-ShockNet framework is conTABLE I. Hyperparameters used in the DIP-ShockNet framework for the MCD-UNet architecture training for flow field predictions

4

Hyperparameters	Values
Initial learning rate	7.94e-4
Decay rate constant	0.994
Training batch size	13
Dropout rate	0.20
Architecture depth	5
No. of channels	[64, 128, 128, 256, 256]
Normalization type	Group normalization (group size = 256)

ducted on an NVIDIA-RTX-A6000 GPU in a local computer.

C. Domain-informed loss function

The training of the weights within our deep learning architecture requires the minimization of a loss function. This function is responsible for comparing various features of both the predicted and the true flow field map, as derived from our training data set. Within the scope of our research, our primary goal is to achieve a high degree of accuracy in predicting global flow field features. Concurrently, we also aim to capture the more intricate local features that showcase pronounced variations in flow.

For this purpose, we opted for the weighted mean squared error (wMSE) as our choice of the loss function. The wMSE calculates the error by using the absolute true field values at each node point within a given flow field. These values act as weights that enable the estimation of a weighted mean squared error for the entire flow field. The wMSE is formally defined

(5)

(6)

(7)

as:

$$L_{wMSE} = \frac{\sum_{i=1}^{h} \sum_{j=1}^{w} |Y^{i,j}| (\tilde{Y}^{i,j} - Y^{i,j})^2}{\sum_{i=1}^{h} \sum_{j=1}^{w} |Y^{i,j}|}$$

Here, *h* and *w* represent the height and width of the flow field, respectively. $Y^{i,j}$ denotes the actual field value at the node location (i, j), whereas $\tilde{Y}^{i,j}$ signifies its predicted counterpart. The intrinsic advantage of such a loss function is its heightened emphasis on higher absolute field values, leading to a robust characterization of both vital global and local features within the flow field.

However, while the wMSE performs adequately on a global scale, giving equal emphasis across the entire flow domain, it occasionally struggles with highly localized features in the flow, such as shock waves. Recognizing this limitation, we refined our approach by developing a custom loss function, enriched with domain-specific information.

Shock waves are characterized by sudden changes in flow properties, predominantly in the flow direction. To effectively capture these nonlinear features, we compute the gradient of the flow field using finite differencing, albeit with an understanding of the inherent approximations involved due to grid transformation and the finite discretization of the Cartesian grid. Figure 2 demonstrates this through contour plots of gradient fields along the *w* (width) and *h* (height) directions of the flow field. The gradient in the flow-wise direction (width) is particularly instrumental in locating the shock wave, while the gradient along the height direction reveals the curvature associated with the shock wave.

In our improved loss function, we introduce an additional term: the weighted mean squared error of the flow field gradient $(wMSE_g)$, derived using central differencing techniques across both the height and width of the flow field, denoted as ΔY_h and ΔY_w .

$$\Delta Y_h^{i,j} = rac{(Y^{i+1,j}-Y^{i-1,j})}{2\Delta h} \ \Delta Y_w^{i,j} = rac{(Y^{i,j+1}-Y^{i,j-1})}{2\Delta w}$$

where Δh and Δw represent the grid spacing in the *h* and *w* directions respectively.

It's important to clarify that our primary aim here is not to calculate exact gradients, but rather to utilize these gradient approximations to locate regions of shock waves, where there is a pronounced jump in field values. By incorporating these approximate gradients into our domain-informed loss function, the DIP-ShockNet model is trained to focus more on areas with shock waves, thereby enhancing predictive accuracy in these crucial regions. Using these gradients, we can define the gradient-focused loss function as:

$$L_{wMSE_g} = rac{\sum_{i=1}^{h}\sum_{j=1}^{w}|\Delta Y_h^{i,j}|(\Delta \tilde{Y}_h^{i,j} - \Delta Y_h^{i,j})^2}{\sum_{i=1}^{h}\sum_{j=1}^{w}|\Delta Y_h^{i,j}|} + rac{\sum_{i=1}^{h}\sum_{j=1}^{w}|\Delta Y_w^{i,j}|(\Delta \tilde{Y}_w^{i,j} - \Delta Y_w^{i,j})^2}{\sum_{i=1}^{h}\sum_{w=1}^{w}|\Delta Y_w^{i,j}|}$$

where $\Delta \tilde{Y}_h$ and $\Delta \tilde{Y}_w$ represents the predicted gradients.

Building upon these individual loss functions, our proposed domain-informed loss function is a combination, taking into account both the wMSE and $wMSE_g$. This domain-informed loss function is described as:

$$L_{DI} = L_{wMSE} + w_0 * L_{wMSE_a} \tag{8}$$

Within this equation, w_0 is a weight factor, determining the relative contribution of the gradient loss to the domaininformed loss function.

III. APPLICATION PROBLEM

We assess the efficacy of our proposed model, particularly in predicting aerodynamic flow fields and shock waves, using the transonic flow over the RAE2822 airfoil as a test case. This airfoil represents the typical profiles found in transonic wings of commercial transport aircraft. Moreover, it serves as the foundational geometry for the second benchmark problem set forth by the AIAA Aerodynamic Design Optimization Discussion Group (ADODG)^{56,57}.

A. Shape parametrization

To generate various airfoil designs for evaluating our proposed method, we employ a geometric parameterization technique, using the RAE2822 airfoil as the baseline. Several shape parameterization methods have been proposed in the literature for the effective parameterization of airfoils and wings^{58–60}. A detailed review of these techniques has been provided by Masters *et al.*⁶¹. Among these methods, we adopt the free-form deformation (FFD) technique for this study. The fundamental principle of FFD involves enclosing the aerodynamic body within an FFD box. Within this box, we define specific control points, termed as FFD nodes. Altering the position of these FFD nodes consequently modifies the shape of the enclosed geometry. A linear elasticity approach is then employed to disseminate the geometric displacements throughout the grid⁶².

For the scope of this study, the RAE2822 airfoil is encapsulated within a 4×2 FFD box, as illustrated in Fig. 3. The four corner nodes of the FFD box remain stationary, while the remaining nodes are constrained to vertical motion, subjected to a maximum displacement limit of ± 0.03 times the chord length.

B. Data generation and preprocessing

For the aerodynamic analysis around the airfoil, we employ an unstructured grid due to its inherent advantages: enhanced flexibility, capability for local grid refinement, and efficiency in terms of grid generation time and effort. A boundary layer is constructed adjacent to the airfoil, with the initial cell of this

Physics of Fluids

Publishing

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

Accepted to Phys. Fluids 10.1063/5.0185370





FIG. 3. Shape parametrization of RAE2822 using FFD technique: (a) original geometry and, (b) deformed geometry.

layer positioned to ensure $y^+ < 1$. A grid independence study was performed to ascertain the optimal grid configuration for our test case. Detailed insights from the grid independence study can be found in the work by Mufti *et al.*^{63,64}. The final grid selection contains 954 points on the airfoil surface and a total grid size of 96,913 nodes as shown in Fig. 4.

The RANS-based CFD simulations are executed using the open-source SU2 code⁶⁵. Simulations are conducted under free stream flight conditions of M = 0.729 and $Re = 6.5 \times 10^6$. To discern the behavior of transonic flow over the airfoil and to capture shock waves of varying magnitudes and positions, we utilized an angle of attack (α) sweep ranging from 0° to 4°. The Spalart-Allmaras model was chosen for turbulence modeling, while a backward Euler scheme ensured the attainment of steady flow conditions.

To populate our design space, Latin hypercube sampling (LHS) was employed, generating a total of 2754 samples from a design space defined by 5 design variables (comprising 4 FFD nodes and α). It is pertinent to note that each sample in our dataset represents a unique airfoil configuration, stemming from a distinct combination of the 4 FFD design variables. This diversity ensures a broad spectrum of airfoil shapes, adding complexity to the problem at hand.

To prepare the data for our model, the CFD flow fields were transferred to a structured Cartesian grid. As highlighted in section I, such a mapping can introduce interpolation errors and potentially result in the loss of information, especially within the boundary layer close to the airfoil surface. To assess the fidelity of the flow field on the Cartesian grid postmapping, we compared the pressure coefficient (C_p) field near the airfoil surface for Cartesian grids of resolutions 64×64 . 128×128 and 256×256 with those obtained from the unstructured CFD grid. From Fig. 5, it is observed that for the 64×64 grid, the C_p distribution and shock wave representation do not align well with the CFD grid results. However, as the grid size is increased to 128×128 and then to 256×256 , the alignment of the flow field features with the CFD results improves significantly. Comparing the C_p distribution, shock wave strength, and location for the 128×128 and 256×256 grids demonstrates that the flow features are not sensitive to the grid size. Based on these findings, the 128×128 structured grid was selected for subsequent analyses. Utilizing this structured grid not only aids in enhancing the model's capability to predict shock waves but also facilitates the development of metrics to quantify discrepancies in shock wave predictions.

C. Performance metrics

In this study, the C_p field is selected as our primary output due to its pivotal role in aircraft design. It critically influences aerodynamic coefficients and is vital for aero-structural design and optimization. While our focus is on the C_p field in this work, the proposed framework is versatile, allowing for adaptation to other aerodynamic fields such as temperature and velocity. These fields, also essential in capturing shock wave dynamics, can be predicted with minor modifications to our domain-informed loss function and performance metrics. Additionally, the MCD-UNet deep learning architecture can also be used to predict multiple output fields of interest simultaneously by only modifying the final output channel dimensions. For example, if C_p , temperature and velocity are the three output fields of interest, the number of output channels should be set to 3. In the present study, since C_p is the only field considered, the number of output channels is set to 1.

1. Field prediction error

To assess the global predictive capability of our model within the chosen flow domain, the root-mean-squared error (RMSE) is employed. Given a testing set of n_t designs, which were excluded during both training and validation phases, the RMSE is formulated as:

$$E(\mathbf{Y}) = \sqrt{\frac{\sum_{i=1}^{n_t} \|\mathbf{Y}_i - \tilde{\mathbf{Y}}_i\|_2^2}{n_t}}$$
(9)

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0185370

Accepted to Phys. Fluids 10.1063/5.0185370



FIG. 4. Unstructured grid used for analysis: (a) full domain, (b) close-up of grid around airfoil



FIG. 5. Comparison of C_p distribution at airfoil surface for different grids for three randomly selected samples from the complete dataset.

Here, Y_i represents the ground truth i-th field solution in the testing set, while \tilde{Y}_i denotes the corresponding DL field prediction. Equation (9) can be used to compute the RMSE for either the mean of the multiple output instances of the field or a certain predictive instance When applied to the mean prediction, it is termed the mean field prediction error. In order to facilitate error comparisons between different methodologies, the prediction error is normalized by the standard deviation of the testing dataset, expressed as:

$$\widehat{E}(\boldsymbol{Y}) = \frac{E(\boldsymbol{Y})}{\sqrt{\sum_{i=1}^{n_t} \|\boldsymbol{Y}_i - \overline{\boldsymbol{Y}}\|_2^2 / n_t}}$$
(10)

It is imperative to note that we avoid utilizing relative error for field accuracy quantification. This is particularly relevant as the C_p field values hover around the free-stream C_p value of 0. Dividing by a value that closely approaches zero might lead to a significantly inflated relative error, even if the deviation between the actual and predicted pressures is minimal. Introducing a bias term could potentially circumvent the near-zero values, but such a term is dependent on the scale²⁹.

2. Shock wave location and strength error

Traditionally, errors in shock wave regions have predominantly been assessed qualitatively. Some recent studies attempted a more quantitative assessment by utilizing heuristics, identifying nodes in proximity to the shockwave region where the deviations in predicted values exceeded a defined threshold^{13,66}. This heuristic method, however, carries an implicit assumption that models inherently perform a lot worse in the shock wave region compared to elsewhere in the flow field. This assumption might not hold, especially for models demonstrating adeptness in shock wave prediction. Therefore, in this study, we introduce metrics aimed at rigorously quantifying prediction accuracy concerning shock wave location and strength.

Consider the C_p distribution along a horizontal line above the airfoil surface where shockwaves are prevalent as shown in Fig. 6. On this line, shockwaves can be identified through a sudden increase in the C_p value. When computing the gradient of C_p along this line, a pronounced spike identifies the shock wave position. To rigorously ascertain the shock wave



FIG. 6. Representation of shock wave on a horizontal line plotted over airfoil surface

location, x^s , along the airfoil chord, we first identify the index *i* such that:

(20)

$$\left(\frac{\frac{\partial C_p}{\partial x}}{\frac{\partial C_p}{\partial x}}\right)_i > \gamma \tag{11}$$

where γ is a threshold. For this study, $\gamma = 0.1$ was selected based on engineering judgment and after analyzing various flow fields of the actual data set. We can now compute the error in shock wave location prediction by comparing actual x^s with the predicted \bar{x}^s . By estimating this error across multiple horizontal lines, drawn at varied Y-coordinates (along the height), a mean error estimate can be derived. This is representative of the shock wave propagation through the outer flow field. Therefore, the shock location error for the test set is:

$$\widehat{E}_{sl}(x^{s}(i)) = \frac{1}{n_{l} \cdot n_{a}} \sum_{j=1}^{n_{l}} \sum_{k=1}^{n_{a}} \left[\frac{|x^{s}(i) - \bar{x}^{s}(i)|}{|x^{s}(i)|} \right]_{j,k}$$
(12)

where n_a is the number of horizontal lines used to derive an average shock wave location estimate for the entire flow field.

To further characterize the shockwave, the strength of the shock wave can be estimated by finding its end location, x^e , which can be identified where the gradient reverts below our threshold after the start index

$$\frac{\left(\frac{\partial C_p}{\partial x}\right)_{i^*}}{\left(\frac{\partial C_p}{\partial x}\right)_{max}} < \gamma \tag{13}$$

This allows us to compute the shock strength, δC_p , as the difference in C_p values at x^s and x^e . The error in shock strength

can then be quantified as:

$$\widehat{E}_{ss}(\delta C_p) = \frac{1}{n_l \cdot n_a} \sum_{j=1}^{n_l} \sum_{k=1}^{n_a} \left[\frac{|\delta C_p - \delta \check{C}_p|}{|\delta C_p|} \right]_{j,k}$$
(14)

where $\delta C_p = |C_p(x^s) - C_p(x^e)|$. Equations (12) and (14) can be applied either to individual predictive instances or to the mean field prediction. For the mean field, the resulting errors are referred to as the mean shock wave strength and mean location errors, respectively. In calculating the shock wave strength and location error across the entire testing dataset, we exclude cases devoid of a shock wave to ensure the error metric is not unduly influenced. The error metrics introduced in this section will be further leveraged in section IV to assess the precision of our proposed methodology in forecasting shock waves with diverse strengths and positions.

IV. RESULTS AND DISCUSSION

This section presents the performance of the DIP-ShockNet model on the RAE2822 airfoil test case. To assess the impact of different input representations and the domain-informed loss function, we partitioned the generated dataset into training, validation, and test sets with 2100, 400, and 274 samples, respectively. Later in this section, we investigate the influence of varying training sample sizes and compare our findings with the POD-based ROM. This analysis involves modifying the number of samples in the training and validation sets while maintaining a consistent test set.

A. Effect of input representations

To understand the effect of input representation on the prediction capability of the proposed method, we explored three different input representations, as illustrated in Fig. 7 and detailed in Table II. Initially, we used an α -mask to capture information about the airfoil geometry and the free stream conditions. In this mask, the value inside the airfoil is set to 0, while the free stream α is distributed across the rest of the domain. This α -mask was then combined with both SDF and IWDF, leading to two input configurations with two channels each.

From the data in Table II, it is evident that the model using only the α -mask input incurs the highest mean field prediction error of 5.95%. While the α -mask input informs the model about the airfoil shape and angle of attack, it lacks insights into the flow domain and distance. Combining the α mask with distance maps enhances the model's spatial comprehension of the flow domain, thus elevating prediction accuracy. Of all representations, the one employing IWDF records the lowest prediction error, standing at 3.45%. The IWDF offers enriched physical and mathematical details, assisting the model in differentiating between the airfoil surface, external flow domain, and the airfoil's interior. Consequently, this leads to an approximate 8% improvement in mean field prediction error when compared with the SDF representation.

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0185370

8

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0185370

2 2

Accepted to Phys. Fluids 10.1063/5.0185370



FIG. 7. Different input representations: (a) airfoil mask, (b) SDF, and, (c) IWDF.

3.75% (+8%)

3.45%

TABLE II. Effect of input representation on mean field error.					
	Input Channel	\widehat{E}			
No.	Туре				
1	α-mask	5.95% (+42%)			
2	α -mask + SDF	3.75% (+8%)			

 α -mask + IWDF

TABLE III. Variation of mean field, mean shock wave strength and mean location prediction errors with the weight of domain-informed loss function

9

Loss function	\widehat{E}	\widehat{E}_{ss}	\widehat{E}_{sl}
L _{wMSE}	3.45%	7.0%	2.01%
$L_{DI} (w_0 = 0.01)$	3.40%	6.65%	1.93%
$L_{DI} (w_0 = 0.1)$	3.23%	6.40%	0.99%
$L_{DI} (w_0 = 1)$	3.97%	6.59%	1.03%

Examining Fig. 8, we observed that the C_p distribution on the airfoil surface for the IWDF representation showed fewer oscillations compared to the single-channel α -mask and the two-channel SDF. IWDF assigns a higher positive weight to the flow field's boundary layer and a stronger negative weight near the inner wall of the airfoil. This distinction helps the model differentiate between the airfoil's surface and the surrounding flow, leading to a more accurate representation of the flow inside the boundary layer. As a result, the C_p predictions using IWDF align more closely with the physics of the problem under study and is used as the input representation for the rest of this paper.

B. Effect of the domain-informed loss function

Our introduction of the domain-informed loss function, as shown in Eq. (8), aimed to enhance the model's capability in recognizing local shifts in C_p , especially those arising from shock waves. The importance of gradient loss within the overall loss is controlled by the weight factor w_0 . To truly gauge the impact of the domain-informed loss function, and to ascertain the best value for w_0 , we carried out multiple training sessions of our model using different wo values. Subsequently, we matched the outcomes with a model that was trained using only the weighted mean squared error loss function.

The results are compiled in Table III, detailing the mean field, mean shock wave strength, and mean location errors for the different loss functions and various w_0 values. With a small weight factor set at $w_0 = 0.01$, the model exhibited marked improvement in predicting both the location and strength of shock waves. This is primarily because the gradient-loss portion of our domain-informed loss function gives special attention to regions where there's a noticeable gap between the actual and predicted gradients. Such regions typically align with the shock wave areas. Consequently, this enhanced accuracy in the shock wave zones contributed to a decrease in the global field error.

By adjusting w_0 to 0.1, we observed a further dip in prediction errors. Specifically, the mean field, mean shock wave strength, and mean location errors were logged at 3.23%, 6.40%, and 0.99% respectively. However, increasing w₀ to 1 swung the balance, causing the model to lean heavily towards gradient loss. While this heightened focus produced accurate shock wave predictions, it inadvertently elevated errors in other field regions that were not as close to the shock wave areas. Taking these observations into account, and after a series of tests, we finalized using a weight factor of 0.1 for our DIP-ShockNet model throughout the remainder of the paper.

C. Prediction of shock waves in the flow field

To evaluate the general performance of our model, we examined the flow field predictions for various airfoil designs and combinations of α . Fig. 9 presents the actual fields, mean predicted fields, error $(\tilde{C}_p - \tilde{C}_p)$ fields, and uncertainty $\sigma(\tilde{C}_p)$ fields for four distinct cases, each characterized by shock waves of different intensities and positions. In scenarios where no shock wave is detected (case I) or where a weak normal shock appears near the trailing edge (case II), the DIP-ShockNet model demonstrates a high level of accuracy in its flow field predictions. The low uncertainty in these predictions indicates a strong confidence in the model's results.

For cases exhibiting a strong normal shock at the center of the airfoil's top surface (case III), a region of shock-induced boundary layer separation is apparent near the airfoil surface. This represents a challenging phenomenon. Nonetheless, the



FIG. 8. Effect of input representation on C_p distribution at airfoil surface: (a) airfoil mask, (b) airfoil mask + SDF, and, (c) airfoil mask + IWDF.



FIG. 9. Contour plots for four different cases: (a)-(d) CFD C_p field, (e)-(h) mean predicted \tilde{C}_p field, (i)-(l) mean error $C_p - \tilde{C}_p$ field, and, (m)-(p) uncertainty $\sigma(\tilde{C}_p)$ field.

model succeeds in capturing both the normal shock wave and the shock-induced separation. The error field reveals a slender region of prediction error close to the shock, though the error magnitude remains modest. Apart from this narrow area, the model's predictions exhibit certainty throughout the entire flow domain. Only when the shock wave assumes a complicated, curved shape (case IV) do we notice a broader region of error and increased prediction uncertainty near the shock wave. On the whole, the model performs admirably in predicting the entire flow field and shock waves of diverse strengths.

To better understand the model's predictive accuracy in the shock wave area, we plotted the C_p distribution on a horizontal line situated above the airfoil's top surface. Figure 10 illustrates the C_p distribution at y = 0.1 for the four cases depicted

in Fig. 9. For all these cases, the actual C_p distribution lies within the model's 95% confidence interval (CI). A sudden spike in the C_p value along this line indicates the presence of a shock wave. The model predicts the shock wave's location with notable precision, with slight discrepancies observed only in the predicted strength of the shock wave.

10

D. Effect of variation in training sample size

Predictive models for evaluating aerodynamic flow fields necessitate access to a significant amount of training data. However, in many design applications, the quantity of training data that can be generated is constrained by the computational

 This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset



FIG. 10. Variation of C_p along a horizontal line drawn at y = 0.1 for four cases given in Fig. 9: (a) case I, (b) case II, (c) case III, and, (d) case IV. The shock wave can be visualized by a sudden jump in C_p

TABLE IV. Comparison of mean field, mean shock wave strength and mean location prediction errors between DIP-ShockNet and POD-ROM for different training sample sizes.

<i>n</i> ₁	\widehat{E}		\widehat{E}_{ss}		\widehat{E}_{sl}	
	POD-ROM	DIP-ShockNet	POD-ROM	DIP-ShockNet	POD-ROM	DIP-ShockNet
100	13.99%	14.11% (+0.85%)	49.06%	24.22% (-50.18%)	16.3%	8.45% (-48.15%)
500	10.49%	8.19% (-21.89%)	42.82%	12.12% (-71.79%)	13.77%	4.09% (-70.29%)
1500	8.80%	4.93% (-43.88%)	33.35%	6.96% (-77.3%)	12.11%	1.89%(-84.39%)
2500	8.23%	3.23% (-60.75%)	29.63%	6.40% (-78.40%)	11.85%	0.99%(-91.64%)

budget. To replace the computationally intensive CFD simulation with a DL model, it's imperative that the model delivers reliable predictions with a limited dataset. To probe this aspect, we trained the DIP-ShockNet framework using varying training sample sizes.



FIG. 11. Comparison of mean field prediction error with the number of training samples for DIP-ShockNet and POD-ROM.

Our approach is compared with the POD-based ROM technique, which, as discussed in section I, has demonstrated proficiency in predicting flow fields with high accuracy, even with a constrained number of training samples. Figure 11 illustrates the variation of field prediction error as a function of training samples for both DIP-ShockNet and POD-ROM. For DIP-ShockNet, this includes both training and validation sets with a train/validation split of 84%/16%. The test set remains consistent across all cases, with both models evaluated on a dataset comprising 274 samples. Within Fig. 11, the solid line for DIP-ShockNet denotes the mean field prediction error, with the error bars signifying the uncertainty in the prediction. With a limited number of training samples, i.e., $n_1 = 100$, the field prediction performance of the two models is almost similar, with POD-ROM slightly outperforming DIP-ShockNet. Given its linear nature, POD finds a linear low-dimensional subspace of the high-dimensional field solution. This enables good prediction accuracy across the field even with limited data. Conversely, DIP-ShockNet, a non-linear DL technique, seeks a non-linear subspace of the high-dimensional solution, encountering challenges in establishing non-linear relationships with sparse data. The pronounced height of the error bar for DIP-ShockNet symbolizes significant model uncertainty when training samples are few.

11

Increasing the number of training samples leads to a decline in the field error for both models. However, after reaching a specific threshold, the prediction error curve for POD-ROM starts to plateau, indicating that further samples don't significantly refine the model's learning. In contrast, DIP-ShockNet exhibits a continuous decline in mean field error and a reduction in uncertainty with additional samples, signifying its capability to learn intricate non-linear relationships with more extensive training.

Table IV presents the mean field prediction error for both techniques. When trained on just 100 samples, POD-ROM surpasses DIP-ShockNet by 0.85%. However, as the training sample size increases, the predictive accuracy gap between the two methods widens. The DIP-ShockNet accuracy when

Physics of Fluids

Publishing

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset



FIG. 12. Comparison of shock wave prediction errors with the number of training samples for DIP-ShockNet and POD-ROM: (a) mean strength error, and, (b) mean location error.



FIG. 13. Contour plots for various training sample sizes for weak shock case: (a) CFD C_p field, (b)-(e) DIP-ShockNet predicted \tilde{C}_p field, (f)-(i) DIP-ShockNet Error $C_p - \tilde{C}_p$ field, (j)-(m) POD- ROM predicted \tilde{C}_p field, and, (n)-(q) POD-ROM error $C_p - \tilde{C}_p$ field.

compared to POD-ROM improves by 21.89% at $n_1 = 500$, and this difference further magnifies to 60.75% for the complete training set. The relationship between shock wave location and strength

error and the number of training samples is illustrated in

Fig. 12. With a limited dataset of 100 samples, our proposed

model encounters challenges in pinpointing the shock wave, manifesting in a strength and location error of 24.22% and 8.45% respectively. Nevertheless, it surpasses the POD-ROM by nearly 50% for both error metrics as depicted in Table IV. The linear nature of POD-ROM limits its capacity to accurately represent the nonlinear shock wave phenomenon. Ex-

12

ing Physics of Fluids

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0185370

AIP Publishing Phys



FIG. 14. Comparison of C_p along a horizontal line drawn at y = 0.1 for weak shock case and different training sizes: (a) $n_1 = 100$, (b) $n_1 = 500$, (c) $n_1 = 1500$, and, (d) $n_1 = 2500$.

panding the training dataset leads to reductions in both shock strength and location errors, enabling our model to capture flow intricacies more proficiently. We also observe a reduction in predictive uncertainty for both error metrics with an increase in training data size. Utilizing the complete training set, the DIP-ShockNet model reduces the shock wave strength error to 6.40% and attains a location error below 1%.

To evaluate the influence of training sample size on predicting the entire flow field and shock waves, flow field contours were plotted for varying training dataset sizes for both DIP-ShockNet and POD-ROM. Figure 13 presents contours for a scenario depicting a weak shock wave on the airfoil's upper surface near its leading edge. At $n_1 = 100$, both techniques face difficulties in representing the flow, leading to pronounced field errors, particularly in shock wave regions. For this scenario, POD-ROM seemingly outperforms the DL approach. Increasing n_1 improves field prediction, shrinking the error-prone areas. Training DIP-ShockNet on the entire dataset yields predictions closely aligned with the actual field, whereas POD-ROM continues to exhibit errors near shock wave regions.

Figure 14 plots the C_p distribution on a horizontal line, highlighting that with fewer samples, both models struggle to capture the pressure jump in the shock wave region. Even with the full dataset, POD-ROM struggles to predict the accurate strength and position of the shock wave. Figure 15 showcases contour plots for a scenario characterized by a pronounced shock wave at the airfoil's midsection. With fewer training samples, substantial error contours emerge around the shock wave for both DIP-ShockNet and POD-ROM. Increasing the dataset size diminishes these errors for both models. With the full dataset, our DIP-ShockNet model delivers a good representation of the shock wave. However, POD-ROM still produces some non-physical, artificial fluctuations - known as spurious oscillations, particularly in regions close to shock waves. These oscillations do not reflect actual physical phenomena but rather result from modeling limitations inherent in the POD-ROM approach, especially in capturing complex nonlinear behaviors such as shock waves. Figure 16 offers a depiction of the shock wave via the C_p jump on a horizontal line. This visualization emphasizes DIP-ShockNet's proficiency in capturing the shock wave strength and location, whereas the oscillations in POD-ROM's predictions become evident near the shock wave.

13

V. CONCLUSION

In this study, we presented DIP-ShockNet, a novel domaininformed probabilistic deep learning approach tailored to predict transonic flow fields containing shock waves. The significant contributions and outcomes of this research are summarized as:

- Uncertainty Prediction: At the heart of the DIP-ShockNet framework is the MCD-UNet architecture. This structure serves a dual purpose: accurately learning the complex mapping between the input and output fields and introducing the ability to predict uncertainty, essential for real-world aerodynamic design applications.
- 2. Inverse Wall Distance Function (IWDF): The introduction of the IWDF was pivotal in enhancing our model's performance near-wall regions. By providing distance field information, the IWDF emphasized field locations close to the airfoil surface, resulting in a marked reduction in prediction oscillations, especially in the boundary layer area. When pitted against common input methods like the airfoil geometric mask and the SDF, IWDF consistently exhibited superior performance resulting in an improvement of predictive performance by 42% and 8% respectively.
- 3. Domain-Informed Loss Function: A central feature of our approach was the domain-informed loss function. This function incorporated gradient information, enabling the model to better identify the presence of shock waves and the variations across them. Through this method, we achieved commendable results, with shock strength errors as low as 6.40% and shock location errors under 1%.

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset

Physics of Fluids

AIP Publishing This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0185370

Accepted to Phys. Fluids 10.1063/5.0185370



FIG. 15. Contour plots for various training sample sizes for strong shock case: (a) CFD C_p field, (b)-(e) DIP-ShockNet predicted \tilde{C}_p field, (f)-(i) DIP-ShockNet Error $C_p - \tilde{C}_p$ field, (j)-(m) POD- ROM predicted \tilde{C}_p field, and, (n)-(q) POD-ROM error $C_p - \tilde{C}_p$ field.



FIG. 16. Comparison of C_p along a horizontal line drawn at y = 0.1 for strong shock case and different training sizes: (a) $n_1 = 100$, (b) $n_1 = 500$, (c) $n_1 = 1500$, and, (d) $n_1 = 2500$.

4. Robustness to Training Data: We also evaluated our model's resilience across different training data sizes. The results were conclusive: DIP-ShockNet consistently outperformed the traditional POD-based ROM in predicting both shock wave location and strength. Particularly, with a comprehensive training set of 2500 samples, our model exhibited a 60% improvement in accuracy over the POD-ROM. In our future efforts, we aim to decrease the training cost associated with the DIP-ShockNet. Our plan is to integrate the technique with an adaptive sampling process. Starting with a limited set of CFD data, the adaptive sampling will identify and target regions where CFD sampling significantly improves model accuracy. Additionally, we're looking to further reduce the sample generation cost by introducing a multifidelity version of our method. This approach involves gen-

erating grids of varying fidelity for the same problem. While dense grids will be used for a select number of high-quality solutions, the less resource-intensive sparse grids will produce a larger set of training data. By aligning these grids, we expect an increase in the volume of training samples without a proportional rise in cost.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflict of interest.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

¹NASA Aeronautics Research Mission Directorate (ARMD), "NASA Aeronautics Strategic Implementation Plan 2023," Tech. Rep. (2023).
 ²IATA, "Aircraft Technology Roadmap to 2050," Tech. Rep. (2020).

- ³J. R. R. A. Martins and S. A. S. A. Ning, Engineering design optimization,
- ⁴A. I. Forrester and A. J. Keane, "Recent advances in surrogate-based optimization," Progress in Aerospace Sciences 45, 50–79 (2009), arXiv:1106.2697.
- ⁶C. Perron, D. Sarojini, D. Rajaram, J. Corman, and D. Mavris, "Manifold alignment-based multi-fidelity reduced-order modeling applied to structural analysis," Structural and Multidisciplinary Optimization 65 (2022), 10.1007/s0158-022-03274-1.
- ⁶M. Fossati, "Evaluation of aerodynamic loads via reduced-order methodology," AIAA Journal **53**, 2389–2405 (2015).
- ⁷D. Rajaram, T. G. Puranik, C. Perron, and D. N. Mavris, "Non-intrusive parametric reduced order modeling using randomized algorithms," in *AIAA Scitech 2020 Forum*, Vol. 1 PartF (American Institute of Aeronautics and Astronautics Inc, AIAA, 2020).
- ⁸A. Behere, D. Rajaram, T. G. Puranik, M. Kirby, and D. N. Mavris, "Reduced order modeling methods for aviation noise estimation," Sustainability (Switzerland) **13**, 1–19 (2021).
- ⁹M. Ghoreyshi, A. Jirásek, and R. M. Cummings, "Reduced order unsteady aerodynamic modeling for stability and control analysis using computational fluid dynamics," (2014).
- ¹⁰S. L. Brunton, B. R. Noack, and P. Koumoutsakos, "Annual Review of Fluid Mechanics Machine Learning for Fluid Mechanics," Annu. Rev. Fluid Mech. 2020 52, 477–508 (2019).
- ¹¹D. J. J. Toal, "On the potential of a multi-fidelity G-pod based approach for optimization & uncertainty quantification," Turbine technical conference and exposition (2014).
- ¹²W. Chen, B. Junqiang, and J. S. Hesthaven, "An iterative approach to improve non-intrusive reduced-order models efficiency for parameterized problems," 21st AIAA International Space Planes and Hypersonics Technologies Conference, Hypersonics 2017 (2017), 10.2514/6.2017-2348.
- ¹³K. Decker, N. Iyengar, D. Rajaram, C. Perron, and D. Mavris, "Manifold Alignment-Based Nonintrusive and Nonlinear Multifidelity Reduced-Order Modeling," AIAA Journal 61, 454–474 (2023).
- ¹⁴ D. Rajaram, R. H. Gautier, C. Perron, O. J. Pinon-Fischer, and D. Mavris, "Non-intrusive parametric reduced order models with high-dimensional inputs via gradient-free active subspace," in *AIAA AVIATION 2020 FORUM* (2020) p. 3184.

- ¹⁵K. Hasegawa, K. Fukami, T. Murata, and K. Fukagata, "Machine-learningbased reduced-order modeling for unsteady flows around bluff bodies of various shapes," Theoretical and Computational Fluid Dynamics **34**, 367– 383 (2020).
- ¹⁶H. Eivazi, H. Veisi, M. H. Naderi, and V. Esfahanian, "Deep neural networks for nonlinear model order reduction of unsteady flows," Physics of Fluids **32**, 105104 (2020).
- Fluids 32, 105104 (2020).
 ¹⁷S. Pawar, S. M. Rahman, H. Vaddireddy, O. San, A. Rasheed, and P. Vedula, "A deep learning enabler for nonintrusive reduced order modeling of fluid flows," Physics of Fluids 31 (2019), 10.1063/1.5113494/14800114/085101_I_ONLINE.PDF.
- ¹⁸J. Z. Peng, S. Chen, N. Aubry, Z. Chen, and W. T. Wu, "Unsteady reduced-order model of flow over cylinders based on convolutional and deconvolutional neural network structure," Physics of Fluids **32** (2020), 10.1063/5.0030867.
- ¹⁹S. Lee, D. Y. J. o. F. Mechanics, and u. 2019, "Data-driven prediction of unsteady flow over a circular cylinder using deep learning," cambridge.orgS Lee, D YouJournal of Fluid Mechanics, 2019-cambridge.org 879, 217–254 (2019).
- ²⁰X. Jin, P. Cheng, W. L. Chen, and H. Li, "Prediction model of velocity field around circular cylinder over various Reynolds numbers by fusion convolutional neural networks based on pressure on the cylinder," Physics of Fluids 30 (2018), 10.1063/1.5024595.
- ²¹T. Nakamura, K. Fukami, K. Hasegawa, Y. Nabae, and K. Fukagata, "Convolutional neural network and long short-term memory based reduced order surrogate for minimal turbulent channel flow," Physics of Fluids 33, 25116 (2021).
- ²⁰N. Thuerey, K. Weißenow, L. Prantl, and X. Hu, "Deep learning methods for reynolds-averaged navier-stokes simulations of airfoil flows," AIAA Journal 58, 25–36 (2020).
- ²³L. Guastoni, A. Güemes, A. Ianiro, S. Discetti, P. Schlatter, H. Azizpour, and R. Vinuesa, "Convolutional-network models to predict wall-bounded turbulence from wall quantities," cambridge.orgL Guastoni, A Güemes, A Ianiro, S Discetti, P Schlatter, H Azizpour, R VinuesaJournal of Fluid Mechanics, 2021•cambridge.org **928**, 27 (2021).
- ²⁴S. Cai, Z. Wang, S. W. J. o, H. . . . , and u. 2021, "Physics-informed neural networks for heat transfer problems," asmedigitalcollection.asme.orgS Cai, Z Wang, S Wang, P Perdikaris, GE KarniadakisJournal of Heat Transfer, 2021-asmedigitalcollection.asme.org.
 ²⁵J. Kim, C. L. J. o, F. Mechanics, and u. 2020, "Prediction of turbu-
- ²⁵J. Kim, C. L. J. o. F. Mechanics, and u. 2020, "Prediction of turbulent heat transfer using convolutional neural networks," cambridge.orgJ Kim, C LeeJournal of Fluid Mechanics, 2020-cambridge.org 882 (2019), 10.1017/jfm.2019.814.
- ²⁶K. Poulinakis, D. Drikakis, I. W. Kokkinakis, and S. M. Spottswood, "Deep learning reconstruction of pressure fluctuations in supersonic shock-boundary layer interaction," Physics of Fluids 35 (2023), 10.1063/5.0156444.
- ²⁷ S. Bhatnagar, Y. Afshar, S. Pan, K. Duraisamy, and S. Kaushik, "Prediction of aerodynamic flow fields using convolutional neural networks," Computational Mechanics 64, 525–545 (2019).
- ²⁸X. Guo, W. Li, and F. Iorio, "Convolutional neural networks for steady flow approximation," in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Vol. 13-17-August-2016 (Association for Computing Machinery, 2016) pp. 481–490.
- ²⁹C. Duru, H. Alemdar, and O. U. Baran, "A deep learning approach for the transonic flow field predictions around airfoils," Computers and Fluids 236 (2022), 10.1016/j.compfluid.2022.105312.
- ³⁰C. Duru, H. Alemdar, and U. Baran, "CNNFOIL: convolutional encoder decoder modeling for pressure fields around airfoils," Neural Computing and Applications 33, 6835–6849 (2021).
- ³¹ K. Zuo, Z. Ye, W. Zhang, X. Yuan, and L. Zhu, "Fast aerodynamics prediction of laminar airfoils based on deep attention network," Physics of Fluids 35 (2023), 10.1063/5.0140545.
- ³²D. Sun, Z. Wang, F. Qu, and J. Bai, "A deep learning based prediction approach for the supercritical airfoil at transonic speeds," Physics of Fluids **33** (2021), 10.1063/5.0060604.
- ³³J. W. Hu and W. W. Zhang, "Mesh-Conv: Convolution operator with mesh resolution independence for flow field modeling," Journal of Computational Physics 452, 110896 (2022).

Publishing

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0185370

ACCEPTED MANUSCRIP1

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset

Physics of Fluids

Publishing

- ³⁴L. W. Chen and N. Thuerey, "Towards high-accuracy deep learning inference of compressible flows over aerofoils," Computers & Fluids 250, 105707 (2023).
 ³⁵Z. Deng, J. Wang, H. Liu, H. Xie, B. K. Li, M. Zhang, T. Jia, Y. Zhang,
- Z. Deng, J. Wang, H. Liu, H. Ale, D. K. Li, M. Zhang, T. Jia, T. Zhang, Z. Wang, and B. Dong, "Prediction of transonic flow over supercritical airfoils using geometric-encoding and deep-learning strategies," Physics of Fluids 35 (2023), 10.1063/5.0155383.
- ³⁶E. J. Hall, S. Taverniers, M. A. Katsoulakis, and D. M. Tartakovsky, "GINNs: Graph-Informed Neural Networks for multiscale physics," Journal of Computational Physics 433, 110192 (2021).
 ³⁷A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and
- ⁷A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. W. Battaglia, "Learning to simulate complex physics with graph networks," proceedings.mlr.pressA Sanchez-Gonzalez, J Godwin, T Pfaff, R Ying, J Leskovec, P BattagliaInternational conference on machine learning, 2020/proceedings.mlr.press (2020).
- ³⁸T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, and P. W. Battaglia, "Learning Mesh-Based Simulation with Graph Networks," ICLR 2021 - 9th International Conference on Learning Representations (2020).
- ³⁹A. Kashefi, D. Rempe, and L. J. Guibas, "A point-cloud deep learning framework for prediction of fluid flow fields on irregular geometries," Physics of Fluids **33** (2021), 10.1063/5.0033376.
- ⁴⁰J. Li, X. Du, and J. R. R. A. Martins, "Machine Learning in Aerodynamic Shape Optimization," (2022), 10.1016/j.paerosci.2022.100849.
 ⁴¹N. Cristianini, J. Shawe-Taylor, et al., An introduction to support vector
- ⁴¹N. Cristianini, J. Shawe-Taylor, et al., An introduction to support vector machines and other kernel-based learning methods (Cambridge University Press, 2000).
- ⁴²C. K. Williams, "Prediction with gaussian processes: From linear regression to linear prediction and beyond," in *Learning in Graphical Models* (Springer, 1998) pp. 599–621.
- ⁴³Å. Bhaduri and L. Graham-Brady, "An efficient adaptive sparse grid collocation method through derivative estimation," Probabilistic Engineering Mechanics 51, 11–22 (2018).
- ⁴⁴A. Bhaduri, Y. He, M. D. Shields, L. Graham-Brady, and R. M. Kirby, "Stochastic collocation approach with adaptive mesh refinement for parametric uncertainty analysis," Journal of Computational Physics **371**, 732– 750 (2018).
- ⁴⁵A. Bhaduri, J. Gardner, C. F. Abrams, and L. Graham-Brady, "Free energy calculation using space filled design and weighted reconstruction: a modified sincle sween approach." Molecular Simulation **46**, 193–206 (2020).
- fied single sweep approach," Molecular Simulation 46, 193–206 (2020).
 ⁴⁶R. Maulik, K. Fukami, N. Ramachandra, K. Fukagata, and K. Taira, "Probabilistic neural networks for fluid flow surrogate modeling and data recovery," Physical Review Fluids 5, 104401 (2020).
- ⁴⁷S. K. Ravi, P. Pandita, S. Ghosh, A. Bhaduri, V. Andreoli, and L. Wang, "Probabilistic transfer learning through ensemble probabilistic deep neural network," in AIAA SCITECH 2023 Forum (2023) p. 1479.
- network," in AIAA SCITECH 2023 Forum (2023) p. 1479.
 ⁴⁸S. Ghosh, P. Pandita, S. Atkinson, W. Subber, Y. Zhang, N. C. Kumar, S. Chakrabarti, and L. Wang, "Advances in bayesian probabilistic modeling for industrial applications," ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering 6, 030904 (2020).
- ⁴⁹J. Denker and Y. LeCun, "Transforming neural-net output levels to probability distributions," Advances in neural information processing systems 3 (1990).

- ⁵⁰D. J. MacKay, "A practical bayesian framework for backpropagation networks," Neural computation 4, 448–472 (1992).
- ⁵¹C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural network," in *International conference on machine learn-ing* (PMLR, 2015) pp. 1613–1622.
- ⁵² Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning* (PMLR, 2016) pp. 1050–1059.
- ⁵³O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18 (Springer, 2015) pp. 234–241.
- ⁵⁴A. Bhaduri, A. Gupta, and L. Graham-Brady, "Stress field prediction in fiber-reinforced composite materials using a deep learning approach," Composites Part B: Engineering 238, 109879 (2022).
- ⁵⁵A. Gupta, A. Bhaduri, and L. Graham-Brady, "Accelerated multiscale mechanics modeling in a deep learning framework," Mechanics of Materials, 104709 (2023).
- ⁵⁶G. R. Anderson, M. Nemec, and M. J. Aftosmis, "Aerodynamic Shape Optimization Benchmarks with Error Control and Automatic Parameterization," in 53rd AIAA Aerospace Sciences Meeting (American Institute of Aeronautics and Astronautics, Reston, Virginia, 2015) pp. 1–18.
- ⁵⁷C. Lee, D. Koo, K. Telidetzki, H. Buckley, H. Gagnon, and D. W. Zingg, "Aerodynamic shape optimization of benchmark problems using jetstream," in 53rd AIAA Aerospace Sciences Meeting (2015) p. 0262.
- ⁵⁸B. M. Kulfan, "Universal Parametric Geometry Representation Method," Journal of Aircraft 45, 142–158 (2008).
- ⁵⁹D. A. Masters, D. J. Poole, N. J. Taylor, T. Rendall, and C. B. Allen, "Impact of shape parameterisation on aerodynamic optimisation of benchmark problem," in 54th AIAA Aerospace Sciences Meeting (2016) p. 1544.
 ⁶⁰R. M. Hicks and P. A. Henne, "Wing Design By Numerical Optimization,"
- ⁶⁰R. M. Hicks and P. A. Henne, "Wing Design By Numerical Optimization," Journal of Aircraft 15, 407–412 (1978).
- ⁶¹D. A. Masters, N. J. Taylor, T. C. Rendall, C. B. Allen, and D. J. Poole, "Geometric comparison of aerofoil shape parameterization methods," AIAA Journal 55, 1575–1589 (2017).
- ⁶²G. K. W. Kenway, G. J. Kennedy, and J. R. R. A. Martins, "Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Adjoint Derivative Computations," AIAA Journal **52**, 935–951 (2014).
- ⁶³B. Mufti, M. Chen, C. Perron, and D. N. Mavris, "A Multi-Fidelity Approximation of the Active Subspace Method for Surrogate Models with High-Dimensional Inputs," in *AIAA AVIATION 2022 Forum* (American Institute of Aeronautics and Astronautics Inc, AIAA, 2022).
- ⁶⁴B. Mufti, C. Perron, R. Gautier, and D. N. Mavris, "Design Space Reduction using Multi-Fidelity Model-Based Active Subspaces," (American Institute of Aeronautics and Astronautics (AIAA), 2023).
- ⁶⁵T. D. Economon, F. Palacios, S. R. Copeland, T. W. Lukaczyk, and J. J. Alonso, "Su2: An open-source suite for multiphysics simulation and design," Aiaa Journal 54, 828–846 (2016).
- ⁶⁶N. Iyengar, D. Rajaram, K. Decker, C. Perron, and D. Mavris, "Nonlinear Reduced Order Modeling using Domain Decomposition," AIAA Science and Technology Forum and Exposition, AIAA SciTech Forum 2022 (2022), 10.2514/6.2022-1250.