

---

# POLYNOMIAL RIDGE FLOWFIELD ESTIMATION

---

A PREPRINT

A. Scillitoe<sup>\*1</sup>, P. Seshadri<sup>1,2</sup>, C. Y. Wong<sup>3</sup>, and A. Duncan<sup>1,2</sup>

<sup>1</sup>Data-Centric Engineering, The Alan Turing Institute.

<sup>2</sup>Department of Mathematics (Statistics Section), Imperial College London.

<sup>3</sup>Department of Engineering, University of Cambridge.

July 15, 2021

## ABSTRACT

Computational fluid dynamics plays a key role in the design process across many industries. Recently, there has been increasing interest in data-driven methods, in order to exploit the large volume of data generated by such computations. This paper introduces the idea of using spatially correlated polynomial ridge functions for rapid flowfield estimation. Dimension reducing ridge functions are obtained for numerous points within training flowfields. The functions can then be used to predict flow variables for new, previously unseen, flowfields. Their dimension reducing nature alleviates the problems associated with visualising high dimensional datasets, enabling improved understanding of design spaces and potentially providing valuable physical insights.

The proposed framework is computationally efficient; consisting of either readily parallelisable tasks, or linear algebra operations. To further reduce the computational cost, ridge functions need only be computed at only a small number of subsampled locations. The flow physics encoded within covariance matrices obtained from the training flowfields can then be used to predict flow quantities, conditional upon those predicted by the ridge functions at the sampled points.

To demonstrate the efficacy of the framework, the incompressible flow around an ensemble of aerofoils is used as a test case. On unseen aerofoils the ridge functions' predictive accuracy is found to be reasonably competitive with a state-of-the-art convolutional neural network (CNN). The local ridge functions can also be reused to obtain surrogate models for integral quantities such as a loss coefficient, which is advantageous in situations where long-term storage of the CFD data is problematic. Finally, use of the ridge framework with varying boundary conditions is demonstrated on a three dimensional transonic wing flow.

**Keywords** Flowfield Estimation · Dimension Reduction · Machine Learning · Computational Fluid Dynamics

## 1 Introduction

The industrial penetration of Reynolds Averaged Navier-Stokes (RANS) has gone far beyond its mainstay of aerospace engineering. Today it is being heavily used in automotive design [1], building design [2], ventilation system design [3], marine propeller design [4] and more recently in the understanding of oral transmissions [5, 6]. In using a turbulence model, RANS avoids the need to resolve pertinent length- and time-scales of turbulence; significantly slashing the cost of generating a flowfield. That said, owing to the size and increasingly realistic mesh topologies, these RANS computations can themselves take several hours on modern clusters. Additionally, in both design and analysis contexts, there is a need to assess multiple geometry and boundary conditions, adding to the computational cost.

It is therefore unsurprising that many have sought to develop ways of (i) efficiently representing a flowfield and (ii) approximating it for a new set of boundary conditions or geometry definition. Falling under the remit of sparse reconstruction, there exists a variety of approaches for reconstructing full unsteady flowfield information from a limited

---

<sup>\*</sup>Corresponding author; Email: ascillitoe@turing.ac.uk, Web: ascillitoe.com.

subset of time-dependent measurements. A suitable basis must first be chosen, onto which the sparse measurements are projected in order to obtain a full reconstruction. Often the chosen basis is linear, and involves a direct computation of the singular value decomposition. For example; Manohar et al. [7] use proper orthogonal decomposition to reconstruct full vorticity fields of laminar flow over a cylinder, while Discetti et al. [8] reconstruct turbulent velocity fields from sparse sensors. More recently, supervised machine learning methods have been used to provide a non-linear framework, with Erichson et al. [9] using a shallow neural network to reconstruct the laminar flow over a cylinder.

Meanwhile, in a range of fields, including uncertainty quantification, design optimisation, and sensitivity analysis, lower fidelity surrogate models (or emulators) are constructed from existing full-field simulation data. Deep neural network architectures, routinely used in data mining, have been used with considerable success as a function approximation technique for high-dimensional physics derived datasets [10, 11]. To increase accuracy, a network can be extended to arbitrary depth and width, however this leads to many weights and a very large amount of training data being required to alleviate overfitting. Recently, *convolutional neural networks* (CNN's) have been shown to be a promising approach to ameliorate this problem for CFD applications. CNN's use convolutional layers to take advantage of local spatial coherence in the input. These layers, combined with successive spatial resizing of the input data, can significantly reduce the number of weights compared to fully connected neural networks. Guo et al. [12] introduced the idea of using a CNN to learn a mapping between an object's geometric representation and the flowfield around it. Bhatnagar et al. [13] and Thuerey et al. [14] recently built upon this work by introducing boundary conditions as an additional input. Such approaches offer accurate and fast data-driven flowfield predictions, allowing for near-immediate feedback for real-time design iterations. Following a different strategy, Tompson et al. [15] replace the pressure projection step in a CFD solver with a CNN in order to accelerate simulations, whilst Sekar et al. [16] use a CNN for the inverse design of an aerofoil, and Jin et al. [17] use a CNN to estimate velocity fields from surface pressure distributions.

Although applications of deep learning networks to real-world problems have become ubiquitous, our understanding of why they are so effective is lacking [18]. Techniques for interpretation are available [19], but they may not be particularly accessible for those without some specialised deep learning knowledge. Moreover, for problems with large numbers of features, widely adopted approaches such as those based on Shapley values [20] quickly become intractable. Consequently, in many fields of engineering, there is still a demand for more readily interpretable surrogate models, even if these bring a slight trade-off in predictive accuracy. One approach to easing the challenge of interpretation, is to reduce the dimensionality of the problem. By reducing the dimensionality of the input space, data-driven ridge functions [21] facilitate straightforward visualisation and understanding, whilst at the same time reducing data requirements. del Rosario et al. [22] show how ridge functions can offer important insights during the conceptual design phase, and in [23], ridge functions aid exploration of a stagnation temperature probe's design space. Wong et al. [24] introduce the idea of using dimension reduction to estimate flow quantities, embedding polynomial ridge functions on an aerofoil surface to estimate the pressure distribution. In this paper, we develop spatially correlated ridge functions, with the target output being flowfield variables such as pressure, velocity or even turbulent variables.

To demonstrate the utility of the proposed framework, we begin by exploring the prediction of flowfields around a set of aerofoils. The design of efficient aerofoils is important since their use is ubiquitous across many industries; for example aircraft wings and controls surfaces, car spoilers, and the gas turbine blades found in aircraft engines and power generation turbines. To provide a point of reference, we compare and contrast the proposed approach to a state-of-the-art flowfield prediction framework based on a convolutional neural network. In addition to examining the predictive accuracy, we explore how the dimension reducing nature of the learned model aids understanding of the existing training dataset. To assess the limitations of the proposed approach, we then tackle the transonic flow over a three dimensional wing, with the inflow boundary conditions varied. We examine how the underlying polynomials perform when extrapolating away from the training data, and discuss approaches to handling this situation. For a hands-on exploration of the framework, example code and an interactive web app are available at [github.com/ascillitoe/flowfield\\_approx](https://github.com/ascillitoe/flowfield_approx).

## 2 Mathematical foundations

In this section, we detail some of the mathematical ideas that underpin this paper.

### 2.1 Flowfield representation

It will be useful to represent a steady flowfield as a set of related scalar-field quantities. These include, but are not limited to, static pressure  $p$ , density  $\rho$ , a velocity component i.e.  $v_x$ , and even the turbulent viscosity  $\nu_t$ . These fluid properties are assumed known for any particular discretised location  $\mathbf{s} \in \mathbb{R}^3$  within the three-dimensional flow domain  $\mathcal{S}$ . We consider  $\mathbf{s}$  to be synonymous with the number of nodes in the flowfield, where  $\mathbf{s}_1$  corresponds to the first coordinate and  $\mathbf{s}_N$  corresponds to the last.

Ignoring model parameters and limiters embedded within any modern flow-solver, one can group the parameters influencing a flowfield into (i) those that govern the mesh and its associated geometry, and (ii) those that set the boundary conditions. For simplicity, we condense the former into a vector of shape parameters  $\mathbf{x} \in \mathbb{R}^d$  arising from an appropriately chosen design space  $\mathcal{X}$ , e.g., Hicks-Henne bump functions, free-form deformation amplitude, or the inputs to an autoencoder that characterises the geometry [25]. We characterise the boundary conditions by a vector  $\mathbf{v} \in \mathbb{R}^b$  where  $b$  represents the number of input boundary conditions, e.g., inlet stagnation pressure, exit static pressure, isothermal walls, etc.

Then, following the statement at the start of this section, we can represent the steady flowfield as a set of related vectors

$$\mathcal{F} = \left\{ \begin{bmatrix} p(\mathbf{x}, \mathbf{v}; \mathbf{s}_1) \\ \vdots \\ p(\mathbf{x}, \mathbf{v}; \mathbf{s}_N) \end{bmatrix}, \begin{bmatrix} \rho(\mathbf{x}, \mathbf{v}; \mathbf{s}_1) \\ \vdots \\ \rho(\mathbf{x}, \mathbf{v}; \mathbf{s}_N) \end{bmatrix}, \begin{bmatrix} v_x(\mathbf{x}, \mathbf{v}; \mathbf{s}_1) \\ \vdots \\ v_x(\mathbf{x}, \mathbf{v}; \mathbf{s}_N) \end{bmatrix}, \dots, \begin{bmatrix} \nu_t(\mathbf{x}, \boldsymbol{\zeta}; \mathbf{s}_1) \\ \vdots \\ \nu_t(\mathbf{x}, \boldsymbol{\zeta}; \mathbf{s}_N) \end{bmatrix} \right\}. \quad (1)$$

For generality, we refer to any scalar field quantity by  $f$  and thus each of the vectors in  $\mathcal{F}$  is generalized via

$$\mathbf{f} = \begin{bmatrix} f(\mathbf{x}, \mathbf{v}; \mathbf{s}_1) \\ \vdots \\ f(\mathbf{x}, \mathbf{v}; \mathbf{s}_N) \end{bmatrix}. \quad (2)$$

Each element in (2) is a scalar-valued function depending only on  $\mathbf{x}$  and  $\mathbf{v}$ . Identifying a suitable surrogate model for  $\mathbf{f}$  is challenging because it effectively has  $\mathbb{R}^d \times \mathbb{R}^b \times \mathbb{R}^N$  degrees of freedom, leading to an insuperable number of model evaluations. In this paper, we demonstrate how subspace-based dimension reduction approaches can be used to thwart the prohibitive cost associated with constructing such a surrogate. In what follows, we present a class of techniques for approximating each element. Without loss in generality, we assume the boundary conditions are fixed, such that each scalar-valued function depends only on  $\mathbf{x}$ . However, in Section 5 we shall consider an alternative case where  $\mathbf{v}$  is varied whilst  $\mathbf{x}$  is fixed.

## 2.2 Ridge approximations

Central to this paper is the notion that a scalar-valued high-dimensional function, such as the static pressure at a given computational node  $\mathbf{s}_i$ , can be approximated by a low-dimensional function defined over a sub-manifold

$$f(\mathbf{x}; \mathbf{s}_i) \approx p_i(\mathbf{W}_i^T \mathbf{x}; \mathbf{s}_i), \quad (3)$$

where  $i = (1, \dots, N)$ . Here  $\mathbf{W}_i \in \mathbb{R}^{d \times n}$  is an orthonormal matrix where  $n \ll d$ , and  $p_i$  is a function in  $\mathbb{R}^n$ . We define the approximation in (3) as a ridge approximation, an approximate form of what Pinkus [21] terms a generalised ridge function. The intuitive computational advantage behind ridge approximations is that instead of estimating a function in  $\mathbb{R}^d$ , we approximate it in  $\mathbb{R}^n$ , which is intrinsically more data efficient. Numerous approaches are available in literature for identifying such structure using only data-driven techniques. These require input-output data pairs of the form  $\{\mathbf{x}_m, f(\mathbf{x}_m; \mathbf{s}_i)\}_m^M$ , where  $M$  represents a suitable number of design of experiments required to approximate the function in  $\mathbb{R}^n$ . Approaches rooted in classical regression include sliced inverse regression [26], principal Hessian directions [27], contour regression [28], sliced average variance estimation [29], and minimum average variance estimation [30]. More recent approaches include the works of Constantine et al. [31] and Hokanson and Constantine [32] that assume  $p_k$  is a polynomial, and explore optimisation techniques over the Grassmann manifold to identify both the coefficients of the polynomial and the dimension reducing subspace  $\mathbf{W}$ . Algorithms for identifying such ridge structure exist even for Gaussian process models (see [33] and [34]).

To provide an intuitive understanding of what these methods can offer, in Figure 1 we apply subspace-based dimension reduction to a dataset obtained from [35] and publicly available in [github.com/equadratures/data-sets](https://github.com/equadratures/data-sets). When a suitable subspace  $\mathbf{W}$  is found, the  $M$  number of input-output training data pairs  $\{\mathbf{x}_m, f_m\}_{m=1}^M$  collapse onto a ridge over the subspace, where they can be used to obtain a functional approximation (a polynomial in this case). Plots of function values over this subspace are called *sufficient summary plots* [31], and are useful in visualising the behaviour of functions in high dimensions.

## 2.3 Polynomial ridge approximations

Let us assume access to a training dataset  $\{\mathbf{X}, \mathbf{F}\}$ , where

$$\mathbf{X} = \begin{bmatrix} | & & | \\ \mathbf{x}_1 & \dots & \mathbf{x}_M \\ | & & | \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} - & \mathbf{f}_1^T & - \\ \vdots & & \\ - & \mathbf{f}_N^T & - \end{bmatrix} = \begin{bmatrix} | & & | \\ \tilde{\mathbf{f}}_1 & \dots & \tilde{\mathbf{f}}_M \\ | & & | \end{bmatrix} \quad (4)$$

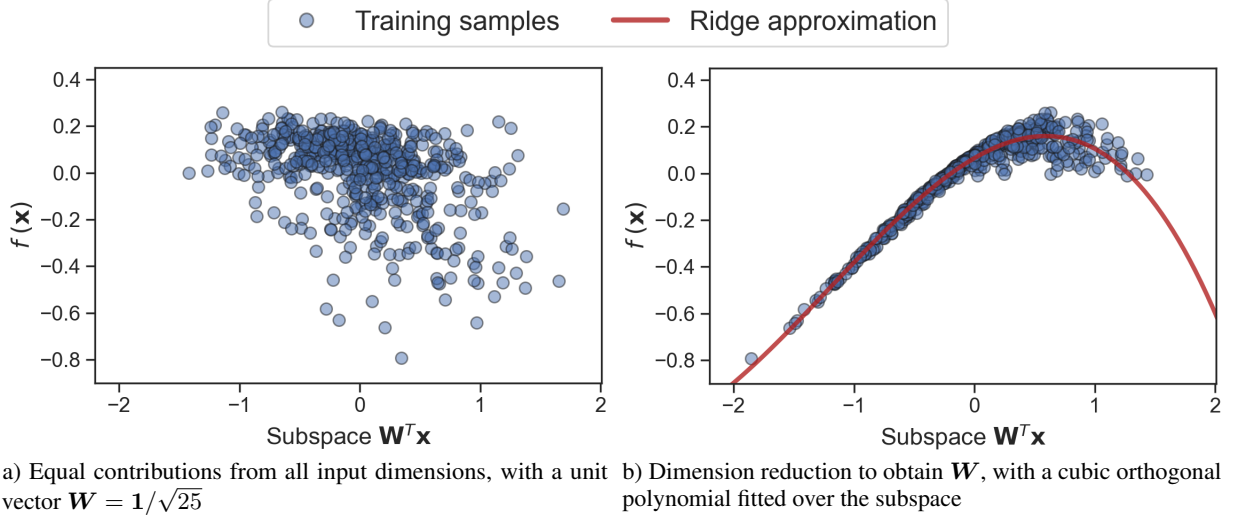


Figure 1: Subspace-based dimension reduction applied to a turbo-machinery problem[35]. The output scalar-valued quantity of interest  $f(\mathbf{x})$ , the efficiency of a turbo-machinery blade, is a function of the  $d = 25$  design parameters in  $\mathbf{x} \in \mathbb{R}^{25}$ . The quantity  $f$  is projected over two different subspaces  $\mathbf{W}^T \mathbf{x}$ . In a),  $\mathbf{W}$  is chosen so that the subspace is formed from an equal contribution of all 25 parameters, whilst in b), a single linear combination of the 25 parameters which best describes the variation in  $f$  is obtained.

with  $\mathbf{X} \in \mathbb{R}^{d \times M}$  and  $\mathbf{F} \in \mathbb{R}^{N \times M}$ . As before  $N$  represents the number of spatial nodes;  $d$  the dimension associated with the design space  $\mathbf{x} \in \mathcal{X}$ , and  $M$  the number of distinct flow-fields, each corresponding to a new design. At an isolated node  $s_i$  across all the flow-fields, we wish to identify a polynomial ridge approximation of the form given in (3). To elaborate on the form of this ridge approximation, we rewrite it as

$$f(\mathbf{x}; s_i) \approx \mathbf{P}(\mathbf{W}_i^T \mathbf{x}^*) \boldsymbol{\alpha}_i \quad (5)$$

where

$$\mathbf{P}(\mathbf{W}_i^T \mathbf{X}) = \begin{bmatrix} \phi_1(\mathbf{W}_i^T \mathbf{x}_1) & \dots & \phi_L(\mathbf{W}_i^T \mathbf{x}_1) \\ \vdots & \vdots & \vdots \\ \phi_1(\mathbf{W}_i^T \mathbf{x}_M) & \dots & \phi_L(\mathbf{W}_i^T \mathbf{x}_M) \end{bmatrix} \quad (6)$$

represents a Vandermonde-type matrix with  $n$ -variate polynomial basis terms  $\{\phi_1, \dots, \phi_L\}$  with a cardinality of  $L$ , with unknown polynomial coefficients  $\boldsymbol{\alpha}_i \in \mathbb{R}^L$ . The cardinality here is defined as the total number of basis terms, and will vary depending on the highest degree in along each dimension and the nature of the interactive polynomial terms between different dimensions. For further details on the construction of multivariate polynomials, we refer the interest reader to Section 1.1 in [36]. The intention behind phrasing our objective as (7), is that for a new design  $\mathbf{x}^*$ , we can estimate the scalar field value at any node  $s_i$  using the computed polynomial ridge approximation. To ensure solutions to this problem are feasible, it will be useful to ensure that the polynomial matrix  $\mathbf{P}$  is well conditioned. This can be aided by adopting orthogonal polynomial basis terms—i.e., Legendre, Hermite or Jacobi—as these have demonstrably lower condition numbers than their monomial counterparts. In this paper, we utilise Legendre polynomials. This has the additional advantage of being able to rapidly yield moments and sensitivities (see [37]).

The ridge approximation in (5) can be obtained by solving the non-linear least squares problem

$$\underset{\mathbf{W}_i \in \mathbb{R}^{d \times n}, \boldsymbol{\alpha}_i \in \mathbb{R}^L}{\text{minimise}} \quad \|\mathbf{f}_i - \mathbf{P}(\mathbf{W}_i^T \mathbf{X}) \boldsymbol{\alpha}_i\|_2^2, \quad (7)$$

which can be simplified to an optimisation problem over the Grassmann manifold via the variable projection method[38]

$$\underset{\mathbf{W}_i \in \mathbb{R}^{d \times n}}{\text{minimise}} \quad \|\mathbf{f}_i - \mathbf{P}(\mathbf{W}_i^T \mathbf{X}) \mathbf{P}(\mathbf{W}_i^T \mathbf{X})^\dagger \mathbf{f}_i\|_2^2 \quad (8)$$

$$\Rightarrow \underset{\mathbf{W}_i \in \mathbb{R}^{d \times n}}{\text{minimise}} \quad \left\| \left( \mathbf{I} - \mathbf{P}(\mathbf{W}_i^T \mathbf{X}) \mathbf{P}(\mathbf{W}_i^T \mathbf{X})^\dagger \right) \mathbf{f}_i \right\|_2^2 \quad (9)$$

$$\Rightarrow \underset{\mathbf{W}_i \in \mathbb{R}^{d \times n}}{\text{minimise}} \quad \left\| \mathbf{O}_{\mathbf{P}(\mathbf{W}_i^T \mathbf{X})}^\perp \mathbf{f}_i \right\|_2^2 \quad (10)$$

where the superscript  $\dagger$  denotes the matrix pseudoinverse, and  $\mathbf{O}^\perp$  is the orthogonal projector onto the complement of the column space of  $\mathbf{P}$ . Gradients for this objective can be readily computed and used in a Gauss-Newton algorithm over the Grassman manifold, as demonstrated in [32]. The computation of the coefficients  $\alpha_i$  trivially follows. This enables us to compute the scalar field quantity at a given node  $i$  for a new design  $\mathbf{x}^*$  with  $\mathbf{P}(\mathbf{W}_i^T \mathbf{x}^*) \alpha_i$  and by extension the scalar field across all  $N$  nodes via

$$\begin{bmatrix} f(\mathbf{x}^*; \mathbf{s}_1) \\ \vdots \\ f(\mathbf{x}^*; \mathbf{s}_N) \end{bmatrix} \approx \begin{bmatrix} \mathbf{P}(\mathbf{W}_1^T \mathbf{x}^*) \alpha_1 \\ \vdots \\ \mathbf{P}(\mathbf{W}_N^T \mathbf{x}^*) \alpha_N \end{bmatrix}, \quad (11)$$

which by construction constraints all subspaces  $\mathbf{W}_1, \dots, \mathbf{W}_N$  to have the same dimension. While the calculation of polynomial ridge approximations across all  $N$  nodes is an embarrassingly parallel operation—as each the input-output data pairs at each node are treated independently—it will be useful to leverage the spatial correlations in a given scalar field to reduce the number of times (10) has to be solved.

## 2.4 Spatial correlations

Subspace-based dimension reduction approaches and the related active subspaces [39] have enjoyed tremendous computational success over the past few years—ushering in a shift from prior sparse grids and full-space design of experiment approaches. As a general observation, within computational fluid dynamics applications, these methods have been primarily applied on integral quantities of scalar field outputs, i.e., efficiency, lift, drag, etc. Wong et al. [24] recognised this, and posited that approximating the composite scalar fields of such integral quantities using dimension reducing subspaces may not only offer greater insight into significant flow features, but may even further reduce the number of simulations required to do so. One important idea born from their work is that the ridge approximations of nodes that are adjacent to each other in a flowfield are likely to be similar in the way they depend on input parameters. In locally smooth regions of the flowfield, nearby nodes are predominantly affected by a similar subset of inputs, resulting in similarity in their ridge subspaces. This notion is quantified via the subspace distance, and in Wong et al. [24] the authors demonstrate that this property enables efficient compression and recovery of flow fields. In what follows, we build upon this idea using a different formalism, leveraging some of our preliminary work in [40].

Using the data available in  $\mathbf{F}$ , we construct a sample covariance matrix  $\mathbf{K} \in \mathbb{R}^{N \times N}$  of the form

$$\mathbf{K} = \frac{1}{M} \sum_{m=1}^M \left( \tilde{\mathbf{f}}_m - \mathbb{E}[\tilde{\mathbf{f}}_m] \right) \left( \tilde{\mathbf{f}}_m - \mathbb{E}[\tilde{\mathbf{f}}_m] \right)^T, \quad (12)$$

where  $\mathbb{E}[\tilde{\mathbf{f}}_m]$  denotes the mean of the  $m$ -th scalar field across all the  $N$  nodes. It is important to note that this covariance matrix captures the spatial correlations across all the nodes for a given scalar field quantity. Its diagonal represents the sample variances in the scalar field at each of the  $N$  nodes.

We partition the scalar field across all  $N$  nodes into two subsets: a smaller subset  $\mathcal{J}$  with  $J$  nodes ( $\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_J$ ), and a subset  $\mathcal{R}$  containing the remaining  $N - J$  nodes ( $\mathbf{s}'_1, \dots, \mathbf{s}'_{N-J}$ ) with  $J \ll N$ . We assume that the scalar field for any design  $\mathbf{x}^*$  across all  $N$  nodes can be expressed as a sample from the multivariate normal distribution  $\mathcal{N}(\mathbb{E}[\mathbf{f}], \mathbf{K})$ , where the first argument represents the mean and the second the covariance. The covariance matrix in (12) can be partitioned into  $(N - J) \times (N - J)$  and  $J \times J$  blocks

$$\mathbf{K} = \begin{bmatrix} \mathbf{K}_{11} & \mathbf{K}_{12} \\ \mathbf{K}_{12}^T & \mathbf{K}_{22} \end{bmatrix}, \quad \text{where } \mathbf{K}_{11} \in \mathbb{R}^{(N-J) \times (N-J)} \text{ and } \mathbf{K}_{22} \in \mathbb{R}^{J \times J}. \quad (13)$$

As the scalar fields arising from this partition will have Gaussian marginals, we can use the Schur complement to approximate the scalar field at the nodes in  $\mathcal{R}$  by computing the ridge approximations at the nodes in the smaller set  $\mathcal{J}$  only

$$\begin{bmatrix} f(\mathbf{x}^*; \mathbf{s}'_1) \\ \vdots \\ f(\mathbf{x}^*; \mathbf{s}'_{N-J}) \end{bmatrix} \approx \mathbf{K}_{12} \mathbf{K}_{22}^{-1} \left( \begin{bmatrix} \mathbf{P}(\mathbf{W}_1^T \mathbf{x}^*; \hat{\mathbf{s}}_1) \alpha_1 \\ \vdots \\ \mathbf{P}(\mathbf{W}_J^T \mathbf{x}^*; \hat{\mathbf{s}}_J) \alpha_J \end{bmatrix} \right). \quad (14)$$

In a nutshell, this implies that we only need to solve (10)  $J$  times rather than that  $N$ . However, this approach relies on the assumption that a suitable  $\mathbf{W}_i$  is found at each point in  $\mathcal{J}$ , such that the projection  $\mathbf{W}_i^T \mathbf{x}$  at each point in the subset explains a large percentage of the total variance at that point. Without this assumption, the resulting correlation matrix  $\mathbf{K}$  will become increasingly diagonal in higher dimensions, and we will be unable to recover the scalar field at the remaining nodes in  $\mathcal{R}$ . To alleviate this issue, poor quality ridge approximations are removed from the smaller subset

$\mathcal{J}$ , and added back into  $\mathcal{R}$ . Various heuristics can be used to assess the quality of ridge approximations. In this paper, we use the training  $R^2$  scores, as discussed in Section 4.5. Finally, for the stability of this procedure it is important that the inverse of  $\mathbf{K}_{22}$  is computed carefully. In the present work, this is done via Cholesky decomposition, with a small constant added to the diagonal to prevent singular matrices.

### 3 Computational setup

In this section we outline the two test cases examined in this paper, in addition to summarising the implementation of the ridge approximation framework and the CNN it will be compared to.

#### 3.1 Test Cases

Within an engineering context, flowfield predictions around a given object are often required at a range of operating conditions. Whereas in other cases, flowfield predictions may be required at a given operating condition, whilst the object itself is altered. To examine both scenarios we use two test cases; 1) a two dimensional subsonic aerofoil flow, with a large number of different aerofoil designs; and 2) a three dimensional transonic wing flow, with the freestream conditions varied.

##### 3.1.1 Exploring the design of a subsonic aerofoil

We start with the well known NACA0012 aerofoil, discretised with a  $449 \times 129$  curvilinear C-grid<sup>2</sup>, and deform it to obtain an ensemble of aerofoil designs. The aerofoil surface is deformed using  $d = 50$  Hicks-Henne bump functions [42]

$$s_2(s_1) = s_{2,base}(s_1) + \sum_{j=1}^d \beta_j b_j(s_1), \quad (15)$$

where  $s_{2,base}$  are the  $s_2$  coordinates of the baseline aerofoil,  $b_j$  is the  $j^{th}$  bump function, and  $s_1$  and  $s_2$  are normalised by the aerofoil's axial chord length  $C_1$ . The bump amplitudes  $[\beta_1, \dots, \beta_d]$  are then stored within the input vector  $\mathbf{x}_m \in \mathbb{R}^d$  for each  $m^{th}$  design.




Figure 2: Deformations made to the NACA0012 aerofoil. Fifty randomly selected deformed designs are shown.

To generate a dataset for training and testing, we create a ( $M = 2000$ )-point design of experiments (DoE) with uniformly distributed Monte Carlo samples for  $\mathbf{x}$ . A random sample of the resulting aerofoil designs, as well as the full range of deformations, is shown in Figure 2. Flowfields are simulated for each design using the incompressible solver of the SU2 CFD code [43]. The commonly used Spalart–Allmaras RANS model[44] is used to represent the effects of turbulence, with the freestream turbulence viscosity ratio set to  $(\nu_t/\nu)_\infty = 5$ . The freestream velocity magnitude is set to  $U_\infty = 1$  m/s, and the reference static pressure is set to zero. The laminar viscosity  $\nu_\infty$  is then set to give a Reynolds number of  $Re = U_\infty C_1/\nu_\infty = 6 \times 10^6$ . Each design is run at three angles of incidence  $\alpha_\infty = 0^\circ, 10^\circ$  and  $15^\circ$ , leading to a dataset consisting of 6000 flowfields in total.

The grid deformations performed when perturbing the baseline aerofoil to reach each new design mean that, even away from the aerofoil surface, grid points are at a slightly different location for each design. Since we wish to learn a functional mapping for the flowfield variables at fixed points in space, we re-sample<sup>3</sup> each flowfield onto the  $90 \times 318$  Cartesian grid shown in Figure 3, and remove points lying inside the solid region of the aerofoil. We choose a Cartesian grid here since we are interested in full flowfield visualisations, but the approach is equally applicable to individual points, surfaces (see Ref. [24]) or planes.

<sup>2</sup>This grid is used as a verification case by the AIAA Fluid Dynamics Technical Committee Turbulence Model Benchmarking Working Group (TMBWG) and is available from [turbmodels.larc.nasa.gov/naca0012\\_grids.html](http://turbmodels.larc.nasa.gov/naca0012_grids.html). Diskin et al. [41] show lift and drag coefficients to be sufficiently grid independent at the  $449 \times 129$  grid resolution.

<sup>3</sup>Resampling is performed with the pyvista python library [45], which uses linear interpolation for resampling.




Figure 3: Zoomed-in view of the discrete curvilinear C-grid representation of a deformed aerofoil (blue points), and the Cartesian grid (black) the flow variables are resampled onto. The aerofoil boundary is shown in yellow.

After the above procedure we are left with  $N = 90 \times 318 = 28.6 \times 10^3$  output values  $f_{i,m}$  for each input  $\mathbf{x}_m$ , forming  $M$  input/output pairs  $(\mathbf{x}_m, \mathbf{f}_m)_{m=1}^M$ , for each of the three angles of incidence investigated. The scalar  $f_{i,m}$  is the field variable at the  $i^{th}$  grid point for the  $m^{th}$  design, where we take the static pressure  $p$ , velocities  $u, v$ , and turbulent viscosity  $\nu_t$  as the field variables to predict. We normalise the aforementioned variables by taking the static pressure coefficient  $C_p = (p - p_\infty)/(p_{0,\infty} - p_\infty)$ , axial velocity ratios  $u/U_\infty, v/U_\infty$ , and turbulent viscosity ratio  $\nu_t/\nu_\infty$ . The  $M = 2000$  designs are further split up into  $M_{train} = 500$  training designs and  $M_{test} = 1500$  testing designs.

### 3.1.2 Varying freestream conditions for a transonic wing

To explore predictions with varying operating conditions, we take the compressible inviscid flow around the well known ONERA M6 transonic wing test case [46]. The salient details of this case are presented in Figure 4. It is discretised by a  $N = 108.4 \times 10^3$  point unstructured grid, with no boundary layer grid necessary due to the inviscid nature of the flow. Freestream boundaries are placed 10.5 root chord length's away from the wing, with the freestream Mach number and angle of incidence set at  $Ma_\infty = 0.8395$  and  $\alpha_\infty = 3.06^\circ$  for the baseline case. A span-wise symmetry boundary is then placed at  $s_3 = 0$ . The flows are simulated using the compressible inviscid solver of the SU2 CFD code[43], and in Figure 4 the resulting Mach contours are shown on the surface of the wing for the baseline case. The flowfield here bears a close resemblance to the real viscous flow, with a characteristic “lambda shock” pattern visible on the upper surface of the wing.

Shown in the inset figure in Figure 4 is the design of experiments used to explore the freestream conditions. For the training data, operating points for  $Ma_\infty$  and  $\alpha_\infty$  are obtained by applying six-point Gauss-Lobatto quadrature rules over the intervals  $Ma_\infty \in [0.6, 1.02]$  and  $\alpha_\infty \in [0, 6]$ , leading to a total of  $M_{train} = 36$  training flows. For testing,  $M_{test} = 25$  points are sampled from the multivariate normal distribution  $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ , with  $\boldsymbol{\mu} = [0.8, 3]^T$  and  $\boldsymbol{\sigma} = [0.14, 3]^T$ . Unlike in test case 1 the geometry is fixed here, therefore there is no need to resample the flowfields onto a base grid. For the scalar field variable  $f_{i,m}$  we consider the local Mach number  $Ma$ .

## 3.2 Polynomial ridge function implementation

To obtain dimension reducing ridges, we use the *equadratures* [47] python library which contains an implementation of the polynomial variable projection method described in Section 2.3. Since each ridge function is independent of one another, we parallelise the task in an embarrassingly parallel fashion, with computation of individual ridge approximations performed by individual Python worker processes. The code is parallelised using the *joblib* library with the *loky* backend. Each parallel Python process is run concurrently on a separate computational core on a F72s\_v2 virtual machine (36 physical CPU's, 144 GiB memory) on the Microsoft Azure cloud computing service. For the  $N = 28.6 \times 10^3$  points in test case 1, with 500 training designs, training takes approximately 20 minutes per field variable ( $\approx 24$  ridge approximations per second). Note that training times can be substantially reduced by taking advantage of spatial correlation in the manner described in Section 4.5.

The flowfield reconstruction procedure described in Section 4.5 is entirely made up of linear algebra operations, thus it is straightforward to implement efficiently with existing linear algebra libraries. For this we use the numpy [48] python library with Intel® MKL. To limit computation and storage requirements, we compute only the  $\mathbf{K}_{12}$  and  $\mathbf{K}_{22}$  blocks of






Figure 4: Mach number contours on the suction surface of the ONERA M6 transonic wing, under the baseline freestream conditions of  $Ma_\infty = 0.8395$  and  $\alpha_\infty = 3.06^\circ$ . The unstructured grid is shown on the  $s_3 = 0$  symmetry plane, and the  $s_3/S = 0.2$  slice is highlighted. The design of experiments (DoE) used to generate the different freestream conditions for training and test data is also shown in the inset figure.

the covariance matrix in (13), since only these blocks are required by the Schur complement step in (14). If further cost savings are required, low-rank approximations for  $\mathbf{K}$  can be obtained, as is done in [40].

### 3.3 Convolutional neural network implementation

Thurey et al. [14] explored a number of CNN architectures for flowfield predictions, and found the U-Net architecture [49] to be the most successful. In this paper we implement a modified version of the framework proposed by Thurey et al. [14], shown in Figure 5. The U-Net architecture consists of an encoder, which progressively down-samples the  $128 \times 128 \times 4$  input data with strided convolutions. The four input channels consist of a boolean mask to define the aerofoil geometry, and three uniform input channels defining the freestream conditions  $Re$ ,  $\alpha_\infty$ , and  $(\nu_t/\nu)_\infty$ . This allows the network to extract increasingly large-scale and abstract information as the number of feature channels grows, until we are left with a  $1 \times 1 \times 256$  derived feature vector. The decoder then does the opposite, with depooling layers reducing the number of features while increasing the spatial resolution. Eventually we are left with four  $128 \times 128$  output channels, consisting of flowfield estimates for  $C_p$ ,  $u/U_\infty$ ,  $v/U_\infty$  and  $\nu_t/\nu_\infty$ . Skip connections help the network to consider low-level input information during the reconstruction of the solution in the decoding layers. The weights  $w_{ij}^{(l)}$  (and biases  $b_i^{(l)}$ ) of the neural network are obtained by minimising the loss function

$$\begin{aligned} \mathcal{L} = & \frac{1}{M} \frac{1}{N} \sum_{m=1}^M \sum_{i=1}^N \phi_{hub} \left( \hat{C}_{p,i,m} - C_{p,i,m} \right) + \phi_{hub} \left( \frac{\hat{u}_{i,m}}{U_\infty} - \frac{u_{i,m}}{U_\infty} \right) \\ & + \phi_{hub} \left( \frac{\hat{v}_{i,m}}{U_\infty} - \frac{v_{i,m}}{U_\infty} \right) + \phi_{hub} \left( \frac{\hat{\nu}_{ti,m}}{\nu_\infty} - \frac{\nu_{ti,m}}{\nu_\infty} \right), \end{aligned} \quad (16)$$

with  $\hat{\cdot}$  denoting the CNN prediction for the field variable, and  $\phi_{hub}$  denoting the Huber loss function [50]. As seen in (16), the error is averaged over all  $M$  number of training designs and  $N$  number of grid points. This loss function is differentiable, enabling it to be back-propagated into the network in order to compute the weight gradient  $\nabla_{\mathbf{w}} \mathcal{L}$ , which is important for the optimisation (*learning*) process.

The complexity of the above CNN is altered in this paper by varying the number of channels (and therefore the number of parameters), with the 488k parameter CNN in Figure 5 found to offer the best compromise between training times and accuracy. All CNN's detailed in this paper are implemented in the PyTorch python library, and trained on a GTX970 1664 core NVIDIA GPU. Training times range from 50 seconds to 80 minutes, depending on the complexity of the network and the number of training designs. Further details are included in Appendix A.






Figure 5: The 488k parameter U-net convolutional neural network architecture used for flowfield predictions. Arrows indicate the direction of forward operations. The network is fully convolutional with 14 layers. Further details of the architecture are given in Appendix A.




Figure 6: Normalised axial velocity,  $u/U_\infty$ , for a deformed aerofoil (from the test set) at an angle of incidence of  $\alpha_\infty = 10^\circ$ . Colour contours show predictions from the second degree polynomial ridge approximations, isolines show the true CFD solution. Sufficient summary plots at four locations are shown.

#### 4 Test case 1: Design of a Subsonic Aerofoil

As an initial demonstration, ridge approximations are obtained for the static pressure coefficient  $C_p = (p_{0\infty} - p_0)/(p_{0\infty} - p_\infty)$ , normalised axial velocity  $u/U_\infty$ , and turbulent viscosity ratio  $\nu_t/\nu$ , with 500 aerofoil designs used for training. Flowfield predictions for a new aerofoil design can then be obtained by transforming the new design vector  $\mathbf{x}$  to the reduced dimensional space, and evaluating the polynomial ridge approximation  $p_i(\mathbf{W}_i^T \mathbf{x}; \mathbf{s}_i)$  at each point  $i = 1, \dots, N$  in the flowfield. In Figure 6, the predicted normalised  $u$  velocity is shown for a randomly selected design from the test set. Qualitatively, the predictions are in close agreement to the *true* velocity field from the CFD solution, shown by the black iso-lines. The design shown here was not used during training, hence the ridge approximations appear to be able to *learn* enough of the flow physics to be able to make predictions for the new design.

In addition to making flowfield predictions, a second key aspect of the ridge approximations are *sufficient summary plots*, seen in the insets within Figure 6. Applied to aerofoil design, they provide visualisation of how the quantity of interest, in this case the  $u$  velocity at selected locations, is altered as the shape is changed. This is made possible since the sufficient summary plot shows the  $u$  velocity over the reduced dimensional space  $\mathbf{W}^T \mathbf{x} \in \mathbb{R}$ , instead of the full space  $\mathbf{x} \in \mathbb{R}^{50}$ . The low scatter of train and test points in the summary plots highlights that low dimensional structure has been successfully found.

#### 4.1 Accuracy of Ridge Approximations

To explore how the accuracy of the proposed ridge approximation framework compares to a state-of-the-art convolutional neural network (CNN), ridge approximations and the 488k parameter U-Net in Figure 5 are trained on  $C_p$  and  $u/U_\infty$  around the aerofoil at two angles of incidence,  $\alpha_\infty = 0^\circ$  and  $\alpha_\infty = 10^\circ$ . To enable fair accuracy comparisons the ridge approximations are obtained on the same  $128^2$  grid used for the CNN instead of the finer  $90 \times 318$  grid used elsewhere. The normalised mean absolute error of predictions

$$MAE_\sigma = \frac{1}{N} \sum_{i=1}^N \left[ \frac{\frac{1}{M} \sum_{m=1}^M |\hat{f}_i(\mathbf{x}_m) - f_i(\mathbf{x}_m)|}{\sigma[f_i(\mathbf{x}_m)]} \right], \quad (17)$$

is then measured for the training and test designs. At each point, the absolute error is averaged for  $M$  number of designs, then normalised by the standard deviation of the true data at that point. The result is averaged over all  $N$  number of points. This is done whilst varying the number of training designs, in order to obtain the error curves in Figure 7. Both the ridge approximation and CNN prediction errors exhibit the same general trends; for low numbers of training designs the test errors are considerably higher than the training errors, suggesting a lack of generalisation (i.e. the models are overfitting to the training data). However the test errors are significantly decreased as the number of training designs is increased. Perhaps surprisingly, despite the data hungry reputation of deep neural networks, the CNN framework is able to achieve low accuracies with a relatively small number of training designs ( $MAE_\sigma < 10\%$  with  $M_{train} < 100$ ). As noted by Ronneberger et al. [49], this is a key benefit of the U-Net and other CNN architectures; the strong use of data augmentation via convolutions allows the available training data to be used more efficiently (compared to a fully connected deep neural network).

Encouragingly, for the  $\alpha_\infty = 0^\circ$  case (Figs 7a and 7b) the ridge approximations are able to achieve mean test errors which are competitive with the CNN. At the higher angle of incidence of  $\alpha_\infty = 10^\circ$  (Figs 7c and 7d) the flowfield is more complex, with stronger streamline curvature leading to larger pressure gradients and greater non-linear behaviour. This case is more challenging for the ridge approximations, with test errors of  $MAE_\sigma \approx 4\%$  compared to the CNN test errors of  $MAE_\sigma \approx 1\%$ . However, accuracies of under 4% would be acceptable for many preliminary design applications. Additionally, as the MAE fields in Figure 8 show, the ridge approximations' prediction errors are actually lower than the CNN errors in many regions of the flow. Compared to the rather randomly distributed CNN errors, the ridge approximations' errors exhibit a smoother spatial distribution, with higher errors near to the aerofoil surface and in the wake.

In the above results, ridge approximations are computed using first degree polynomials. In Table 1, results with higher degree polynomials are presented. Clearly, increasing the polynomial degree,  $p$ , allows for lower test errors to be obtained (although excessively high-degree polynomials will suffer from over-fitting). However, increasing  $p$  also increases the cardinality of the polynomial; if all interaction terms are included, and the polynomial has degree  $p$  in all  $n$  directions, its cardinality is given by  $L = (p + 1)^n$ . More degrees of freedom result in higher variances of the resulting regression estimators, necessitating more training samples to mitigate this. The exponential dependence on polynomial degree implies that higher degree polynomials are significantly more costly when in higher dimensions. This highlights a key benefit of obtaining polynomials in a dimension reducing subspace of dimension  $n$ , where  $n \ll d$ . Various approaches exist to mitigate this exponential scaling[36]. However, even then, it is still desirable to reduce the number of dimensions. As a compromise between accuracy and computational cost, second degree polynomials are used in the remainder of the results presented for the present test case.

#### 4.2 Design Space Exploration

The dimension reducing nature of the ridge approximations is particularly useful for design space exploration. To demonstrate, ridge approximations for the turbulent viscosity ratio  $\nu_t/\nu_\infty$  are obtained using all 500 training designs. In Figure 9 a sufficient summary plot for a point in the aerofoil's wake is shown. The turbulent viscosity ratio in the wake is of interest because high levels of turbulence in wakes can be a leading source of loss in many flows [51]. The sufficient summary plot is examined, and a design at either end of the  $\mathbf{W}^T \mathbf{x}$  subspace is selected. Contours of the turbulent viscosity ratio for the two designs are shown in Figure 9. Through this approach, the end user can explore how

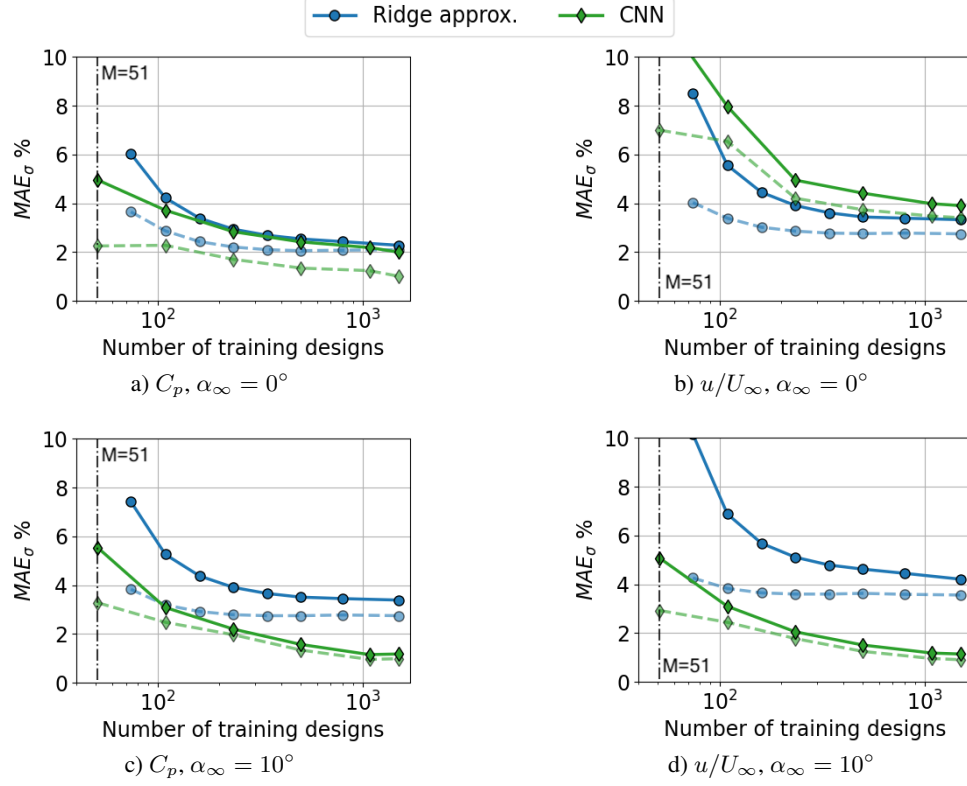


Figure 7: Normalised MAE of  $C_p$  and  $u/U_\infty$  predictions for first degree polynomial ridge approximations and convolutional neural network at two angles of incidence. Solid lines = test error, dashed lines = training error. Results are averaged over three randomly selected train/test splits.

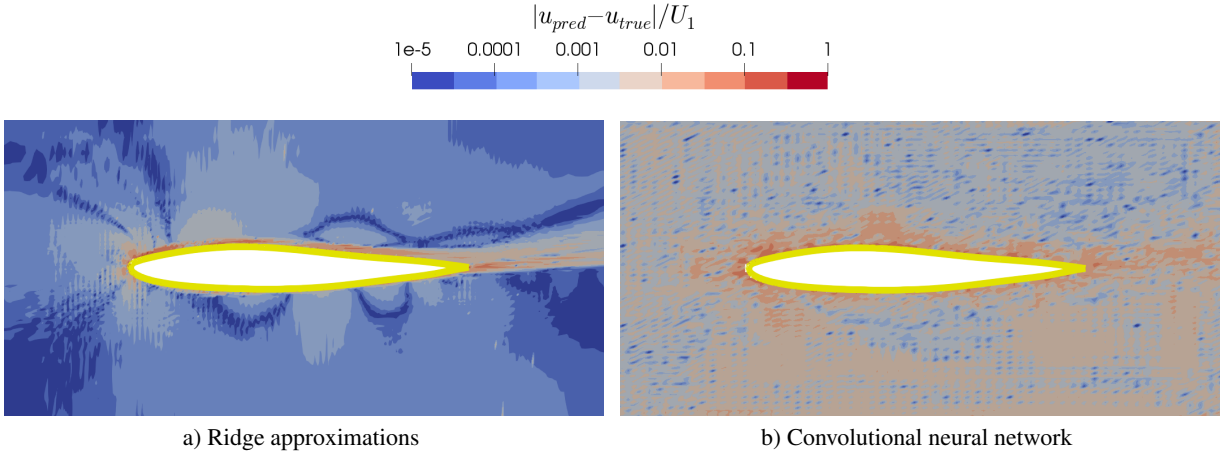


Figure 8: Comparison of absolute error in predictions of normalised axial velocity from the first degree ridge approximations and convolutional neural network, for an aerofoil from the test set at  $\alpha_\infty = 10^\circ$ .

flowfield variables at select locations in the flowfield change as they traverse the design space. Such an activity would be far more challenging when negotiating the original design space  $\mathbf{x} \in \mathbb{R}^d$ . Effectively, dimension reduction distils the design space into critical directions along which quantities of interest exhibit maximum variance, which enables efficient design space exploration.

Table 1: Normalised MAE of ridge approximations with various polynomial degrees, for the  $C_p$  and  $u/U_\infty$  fields at an angle of incidence of  $\alpha_\infty = 10^\circ$ .

Polynomial degree, $p$	Training $MAE_\sigma$ (%)		Test $MAE_\sigma$ (%)	
	$C_p$	$u/U_\infty$	$C_p$	$u/U_\infty$
1	2.8	3.5	3.6	4.5
2	2.1	2.8	2.9	3.7
3	1.6	2.0	2.7	3.6




Figure 9: Turbulent viscosity ratio,  $\nu_t/\nu_\infty$ , for two deformed aerofoils (from the test set) at an angle of incidence of  $\alpha_\infty = 10^\circ$ . Colour contours show predictions from the ridge approximations, isolines show the true CFD solutions. The two designs are at either end of the sufficient summary plot for a point in the wake.

In addition to exploring directions of maximum variance, one can also explore directions along which quantities of interest exhibit minimum variance. Following Section 2.2, the input design vector can be decomposed as

$$\begin{aligned} \mathbf{x} &= \mathbf{W}\mathbf{W}^T\mathbf{x} + \mathbf{V}\mathbf{V}^T\mathbf{x} \\ &= \mathbf{W}\mathbf{w} + \mathbf{V}\mathbf{v} \end{aligned} \quad (18)$$

where  $\mathbf{w} = \mathbf{W}^T\mathbf{x}$  is the active coordinate and  $\mathbf{v} = \mathbf{V}^T\mathbf{x}$  the inactive coordinate. The inactive subspace  $\mathbf{V}$  permits identification of designs which are invariant with regards to the quantity of interest. The task is to find input vectors  $\mathbf{x}$  with a fixed  $\mathbf{w}$  but different  $\mathbf{v}$ , whilst obeying the constraint  $-1 \leq \mathbf{x} \leq 1$ , for which we use a hit-and-run algorithm similar to that implemented in [52]. In Figure 10, this is done for the ridge approximations at the locations labelled a) and b) in Figure 6. The resulting *design envelopes* inform us how the aerofoil can be deformed without the quantity of interest  $u/u_\infty$  at these two locations being affected. Such information can be used to form manufacturing tolerances, or to decide whether to scrap components, as is done in [23, 53, 54].




Figure 10: Design envelopes obtained by sampling the inactive subspaces of the ridge approximations for  $u/u_\infty$  at locations a) and b) in Figure 6. 1000 samples are generated using a hit-and-run algorithm [52], with the grey lines corresponding to the blade profiles of 10 randomly selected samples.

### 4.3 Physical Insight

In addition to exploring the design space, ridge functions can also provide important physical insights. As an example, we consider the stagnation pressure around the aerofoil. In an ideal (inviscid) incompressible flow, the stagnation pressure,  $p_0 = p + 0.5\rho U^2$ , is constant along streamlines<sup>4</sup>. Streamline curvature results in a reversible exchange between the static pressure  $p$  and dynamic pressure  $0.5\rho U^2$ . However, in a real flow, irreversible processes such as

<sup>4</sup>The stagnation pressure is constant everywhere if the flow is inviscid and irrotational.

viscous effects cause a loss in  $p_0$ , which it is important to control. Take the case of an aircraft; high losses over the wing are synonymous with high drag [55], while in the engines, high losses lead to decreased component efficiencies, both of which lead to a higher fuel consumption. The pressure loss at a given location  $s$  can be quantified through the loss coefficient

$$Y_p(s) = \frac{p_{0\infty} - p_0(s)}{0.5\rho_\infty U_\infty^2}, \quad (19)$$

where  $p_{0\infty}$  is the freestream stagnation pressure, and  $Y_p = 0$  indicates no loss in stagnation pressure, while  $Y_p = 1$  indicates a total loss in stagnation pressure.

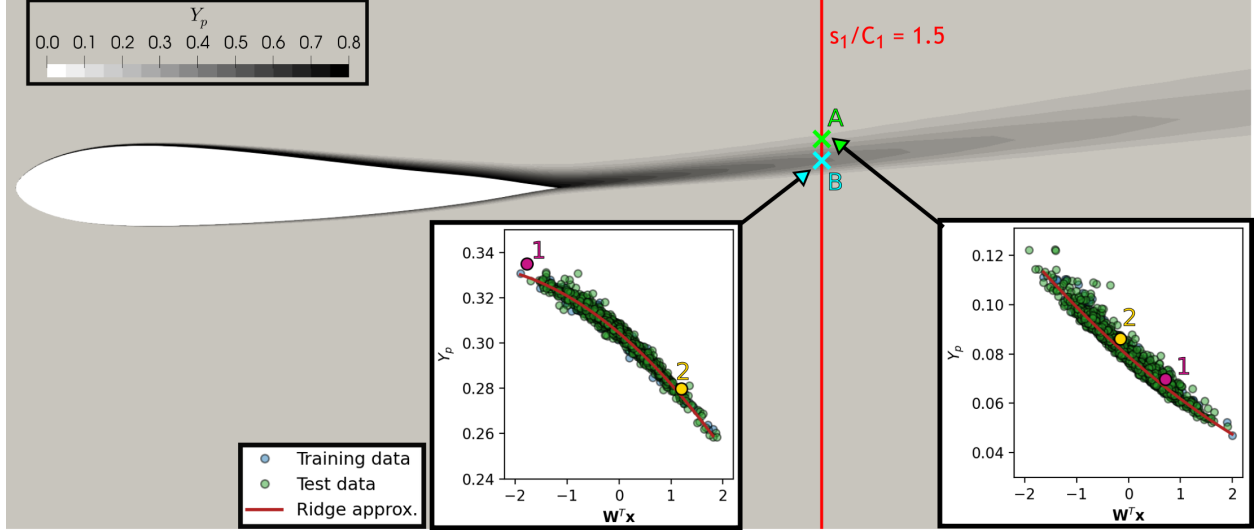


Figure 11: Contours of loss coefficient,  $Y_p$ , for a deformed aerofoil from the test set (“Design 1”) at an angle of incidence of  $\alpha_\infty = 10^\circ$ . A line at 50% axial chord downstream of the trailing edge is highlighted, and sufficient summary plots are shown for two points within the wake, on this line. Design 1, and a second design, are highlighted in the summary plots.

The loss coefficient  $Y_p$  is calculated from the CFD data for 500 training designs, and the approximations for a randomly selected test design are visualised in Figure 11. In Figure 12a, profiles of  $Y_p$  across the  $s_1/C_1 = 1.5$  plane are plotted for the two designs labelled 1 and 2 in Figure 11. For both designs, the high loss in the wake (around  $0 \leq s_2/C_1 \leq 0.15$ ) is accurately predicted by the ridge functions.

The sufficient summary plots in Figure 11 show how  $Y_p$  at probe locations A and B change as the design space is traversed. Furthermore, the dimension reducing subspace vector  $\mathbf{W}$  can be examined. Recalling that in the present framework  $\mathbf{x}$  contains the Hicks-Henne bump amplitudes, the elements of  $\mathbf{W}$  show us how the bump amplitudes affect the reduced coordinates  $\mathbf{W}^T \mathbf{x}$ . In Figure 12a, the scaled elements of  $\mathbf{W}$  for the  $Y_p$  ridge functions at A and B are projected onto the baseline aerofoil geometry<sup>5</sup>. The weights offer important physical insights here; they show us how the aerofoil surfaces must be deformed in order to affect the loss coefficient  $Y_p$  at probes A and B.

#### 4.4 Obtaining Global Quantities

Understanding the sensitivity of local flow variables to a design can be invaluable to a designer. But, sometimes, a designer may also be interested in global quantities of interest. For example, in the present case, the mass averaged loss across the  $s_1/C_1 = 1.5$  line in Figure 11 is an important quantity, and is given by

$$\bar{Y}_p = \frac{\int u(s_1, s_2) Y_p(s_1, s_2) ds_2}{\int u(s_1, s_2) ds_2}, \quad (20)$$

where  $u(s_1, s_2)$  and  $Y_p(s_1, s_2)$  denote the local axial velocity and loss coefficient respectively. The mass averaged loss represents the total pressure loss caused by the aerofoil. A surrogate model for  $\bar{Y}_p$  could be obtained by preprocessing the CFD data, and then computing a ridge approximation for the precomputed  $\bar{Y}_p$  values. However, alternatively, the local ridge functions themselves can be integrated following the approach outlined in Appendix B. A dimension

<sup>5</sup>Each element of the subspace weight vector is plotted at the surface location of its associated bump.

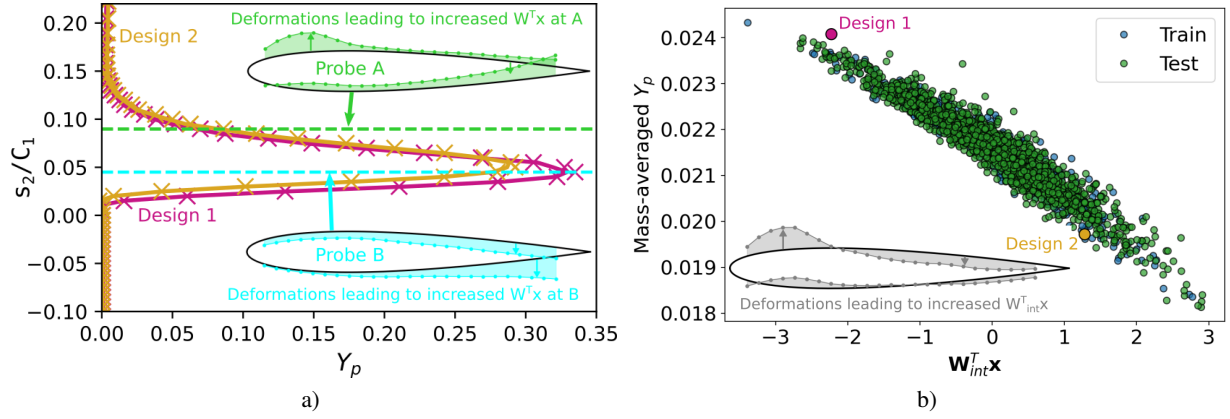


Figure 12: a) Profiles of loss coefficient  $Y_p$  across  $s_1/C_1 = 1.5$ , and their ridge approximations, for the two designs highlighted in Figure 11. Also shown are the  $s_2$  locations of probes A and B from Figure 11 (dashed lines), along with the projections of  $\mathbf{W}$  at these two locations onto the baseline aerofoil. b) Sufficient summary plot for the mass-averaged loss coefficient across the  $s_1/C_1 = 1.5$  line, with designs 1 and 2 highlighted, and the associated subspace vector  $\mathbf{W}_{int}$  projected onto the baseline aerofoil.

reducing subspace for  $\bar{Y}_p$ ,  $\mathbf{W}_{int}$ , is obtained directly from the ridge approximations for  $u(s_1, s_2)$  and  $Y_p(s_1, s_2)$ . The resulting one dimensional sufficient summary plot for  $\bar{Y}_p$  is shown in Figure 12b. It is apparent that designs 1 and 2 lie at either end of the design space with regards to their mass averaged loss, indicating that design 2 is objectively the better design (with regards to loss). Similarly to before, in Figure 12b the elements of  $\mathbf{W}_{int}$  are projected onto the baseline aerofoil. This shows that the total loss can be reduced by deforming the suction surface outwards near the leading edge, and inwards near to the trailing edge.

To understand why design 2 has a lower loss, a second law analysis can be performed. Herwig and Schmandt [56] show that the overall entropy generation rate

$$S' = \frac{\rho(\nu + \nu_t)}{T} \left( \frac{\partial u_i}{\partial s_j} + \frac{\partial u_j}{\partial s_i} \right) \frac{\partial u_i}{\partial s_j}, \quad (21)$$

can be used to understand the sources of drag or pressure loss within a flow. The contours of  $S'$  in Figure 13 show that, at  $\alpha = 10^\circ$ , the majority of the entropy generation comes from the suction surface boundary layer and the wake region<sup>6</sup>. In Figure 13 the delta between the  $y$ -averaged entropy generation for design 2 and the baseline design ( $\mathbf{W}^T \mathbf{x} = 0$ ) is shown. Deforming the suction surface outwards near the leading edge, and flattening the aft portion (see Fig. 12b), appears to increase entropy generation over the first quarter of the suction surface. But, this is offset by a reduction in entropy generation in the wake and leading edge stagnation region, leading to a net decrease in the entropy generation. Whilst an experienced aerodynamicist may have foreseen this conclusion, the present example serves to illustrate the potential of an easily interpretable surrogate model which can be readily integrated to obtain global quantities.

#### 4.5 Exploiting Spatial Correlation

Up to this point, flowfield predictions have been made using ridge approximations at every grid point. For larger grids, such as three dimensional ones, obtaining  $N$  number of ridge approximations might be computationally intractable. Taking a similar approach to how the CNN takes advantage of spatial coherence in the flowfield, the method outlined in Section 2.4 exploits spatial correlations to reduce the number of ridge approximations required. In Figure 14, ridge approximations are obtained for  $J = 1000$  randomly subsampled points (a sampling rate of  $J/N = 3.5\%$ ). Since the baseline mesh has a high mesh density near the aerofoil (where the gradients are high), the subsampled points are also clustered near to the aerofoil, which is desirable. The covariance matrix  $\mathbf{K}$ , computed on the training data, is then used to *upsample* the ridge approximations at the subsampled points back to the remaining  $N - J$  points.

The colouring of the subsampled points in Figure 14 shows that there are a small number of points with poor quality ridge approximations. Preliminary work[40] showed that these poor quality approximations can cause spurious noise in the flowfield estimates when upsampled. As discussed in Section 2.4, the training  $R^2$  score can be used as a heuristic to

<sup>6</sup>Decomposing Equation 21 into entropy generation due to laminar viscosity  $\nu$ , and that due to turbulent viscosity  $\nu_t$ , shows that laminar irreversibilities dominate over the aerofoil surfaces, while turbulent mixing dominates in the wake.






Figure 13: Contours of total entropy generation rate,  $S'$ , for design 2 (see Figs. 11- 12), and axial profile of the difference in y-averaged generation rate between design 2 and the baseline design. The areas under the profile at various locations are also shown. The units of  $S'$  are  $Wm^{-3}K^{-1}$ .




Figure 14: Flowfield estimate of axial velocity,  $u/U_\infty$ , obtained by upsampling ridge approximations at  $J = 1000$  randomly subsampled points. The aerofoil is from the test set, and the angle of incidence is  $\alpha_\infty = 10^\circ$ . Colour contours show upsampled ridge approximations, while the iso-lines show the true CFD solution. The subsampled points are coloured by the mean  $R^2$  score of their ridge approximations over the training data.

measure the quality of each ridge. Ridges with  $R^2 < 0.6$  are ignored during upsampling, with their grid points added back into the  $N - J$  set of points. As seen in Figure 14, this strategy results in qualitatively accurate flowfield predictions using ridge approximations at only 3.5% of the grid points, reducing the training time by a factor of approximately 29.

## 5 Test case 2: Varying Freestream Conditions for a Transonic Wing

Following on from the success of the first test case, we now turn to the transonic wing test case described in Section 3.1.2 in order to expose the limitations of polynomial ridge functions. Here the geometry, represented by  $\mathbf{x} \in \mathbb{R}^b$ , is fixed.



However, the boundary conditions, represented by  $\mathbf{v} \in \mathbb{R}^b$ , are varied. The transonic nature of the flow, and the possibility of extrapolating beyond the training dataset, make this test case particularly challenging. The input vector  $\mathbf{v}$  only consists of two dimensions in this case,  $Ma_\infty$  and  $\alpha_\infty$ . We could once again obtain one dimensional ridge functions. However, two dimensional ridge functions are still simple to visualise. Considering a ridge function of the form

$$f(\mathbf{v}) \approx g(\mathbf{W}^T \mathbf{v}) \quad \text{where} \quad \mathbf{W} \in \mathbb{R}^{b \times n} \quad \text{and} \quad \mathbf{v} \in \mathbb{R}^b, \quad (22)$$

we set  $\mathbf{W} = \mathbf{I}$  (the identity matrix), resulting in two dimensional ridge functions ( $n = b$ ).

### 5.1 Interpolation versus extrapolation

In this case, the behaviour of the ridge functions as we extrapolate away from the training boundary conditions is important. To explore this, polynomial ridge functions are trained on the  $M_{train} = 36$  point DoE shown in the inset of Figure 4. The transonic nature of these flows results in a noticeable step change in the flow behaviour as the freestream Mach number increases, and at least fifth degree polynomials are found to be necessary to properly capture the surface Mach number's response. The two dimensional ridge functions are used to estimate the Mach number on the wing surface at the  $s_3/S = 0.2$  span-wise slice highlighted in Figure 4, with the estimates for two different DoE points from the *test* set shown in Figures 15a and 15b. The corresponding two dimensional ridges for the locations labelled A and B are then visualised in Figures 15c and 15d. From these visualisations it is clear that DoE point 1 lies within the training distribution, hence the its predictions (Fig. 15a) can be considered to be interpolations. In this case, the predicted Mach number distribution agrees well with the true distribution.

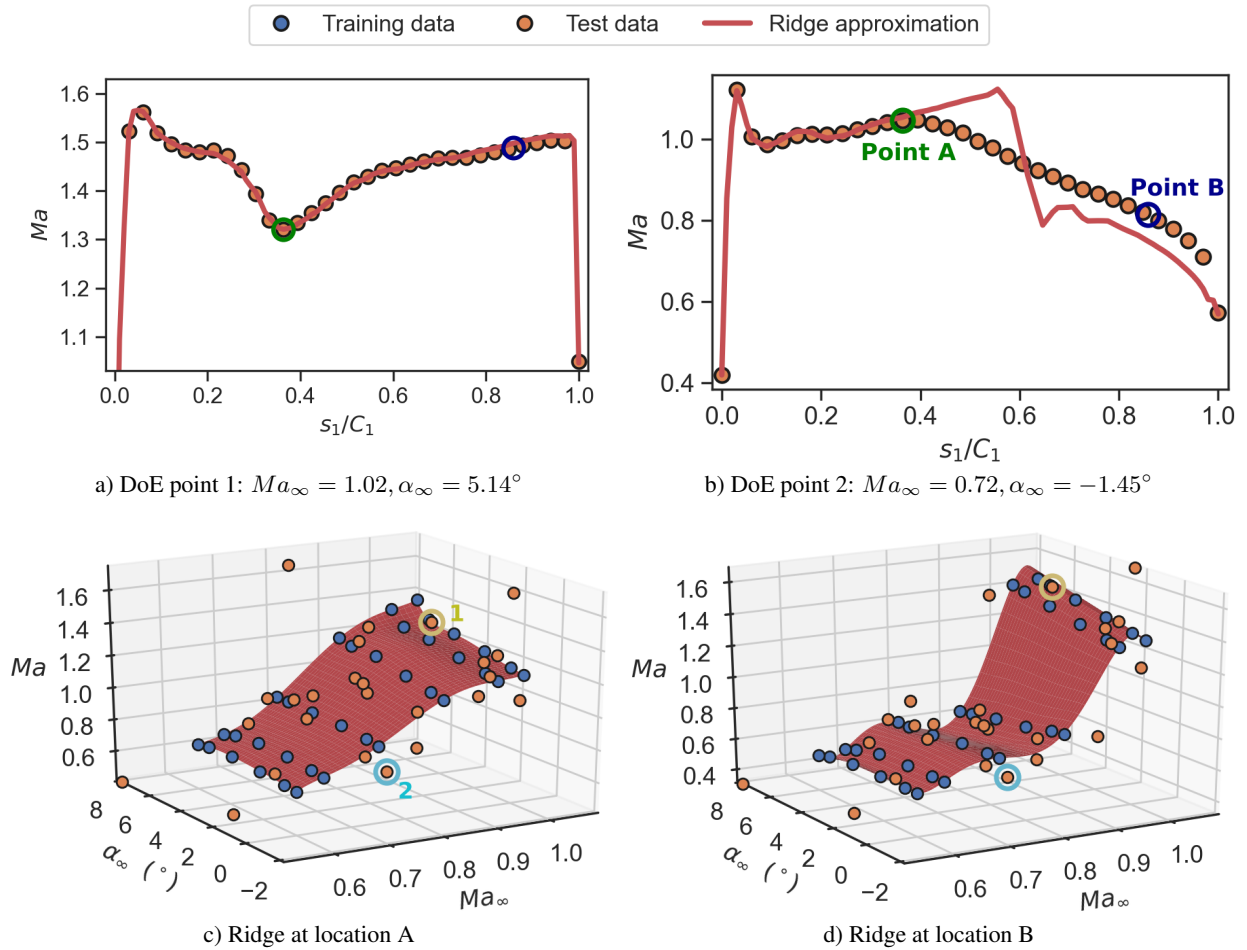


Figure 15: Polynomial ridge approximations for Mach number on the surface of the transonic wing at the  $s_3/S = 0.2$  span-wise slice. The predicted Mach distributions for two DoE points from the *test* set are shown in a) and b), while the corresponding ridges for two  $s_1/C_1$  locations are visualised in c) and d). The two DoE points are also labelled in c) and d).

Moving on to DoE point 2, here the ridge approximations must extrapolate away from the flow conditions seen during training. Unsurprisingly, the predicted distribution in Figure 15b agrees less well with the true distribution in this case, especially on the aft portion of the wing. This highlights a note of caution regarding polynomial ridge functions; as with many function approximation techniques, caution must be taken when using the underlying polynomials to extrapolate beyond the training data. If such a task is routinely necessary, it is possible to replace the polynomials with other classes of models with more stable behaviour under extrapolation, such as piecewise polynomials (splines) or kernel regressors. However, we shall leave this as a topic of future work.

## 5.2 Three-dimensional flowfield predictions

When extending the framework to three dimensional flowfields, computational efficiency is crucial. To achieve a surrogate model for the full flowfield the subsampling and upsampling strategy is used once again. Ridge functions are fitted at the  $J = 5000$  randomly sampled points (a sampling rate of 4.6%), with a covariance matrix used to upsample the ridge approximations to full three dimensional flowfield estimates. A slice of a flowfield estimate for a *test* DoE point is shown in Figure 16. Comparing the colour contours and iso-lines, the prediction appears to be reasonable, with only a slight disagreement on the aft section of the suction surface.

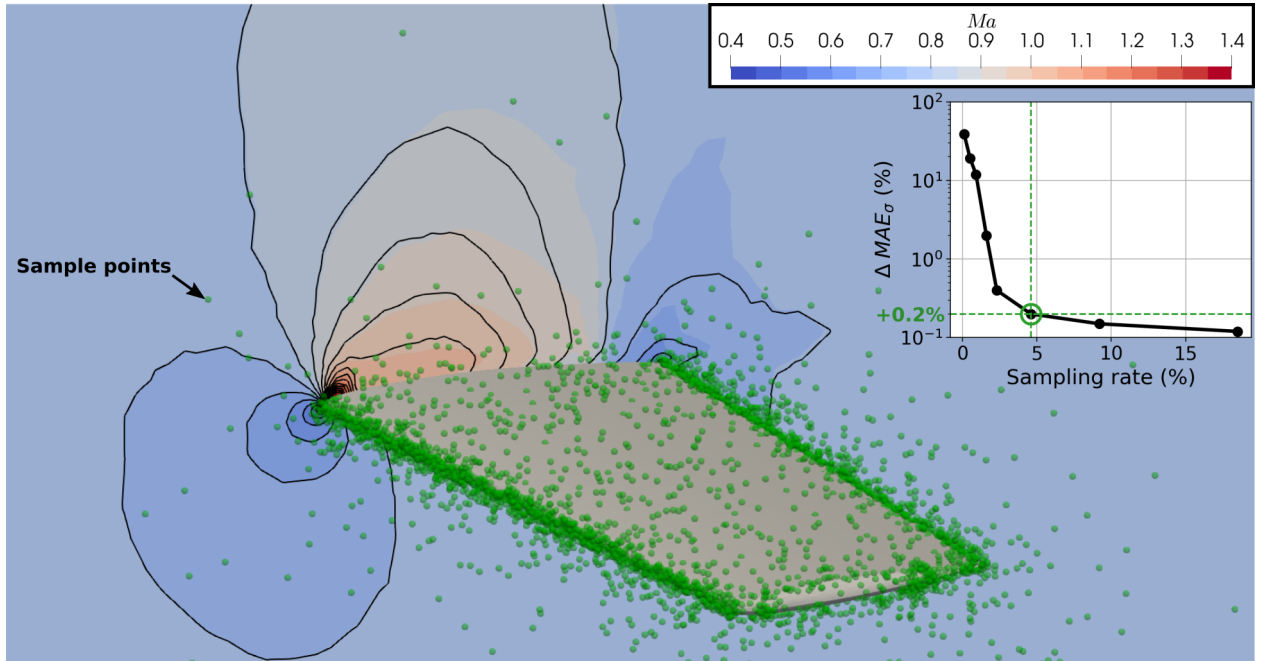


Figure 16: A span-wise slice ( $s_3/S = 0.2$ ) of the upsampled transonic wing flowfield estimates of Mach number, for a DoE point from the *test* set, with  $Ma_\infty = 0.77$  and  $\alpha_\infty = 4.6^\circ$ . The  $J = 5000$  randomly subsampled points are shown in green. The inset figure shows the difference in normalised MAE between the subsampled-upsampled estimates, and the estimates with no subsampling ( $J = N$ ), for different sampling rates.

The above result suggests the subsampling strategy is effective for three dimensional flows in addition to two dimensional ones. The inset figure within Figure 16 shows the effect of sampling rate on the prediction error. For the present case, sampling rates of lower than  $\approx 4\%$  lead to significantly increased errors. Nevertheless, sampling rates of 4-10% still allow for significant computational cost savings.

## 6 Conclusions

This paper exposes the idea of using spatially correlated ridge functions to rapidly estimate flowfields. The resulting data-driven framework could be trained on existing CFD data where available, or it could be integrated within a wider design of experiment, with new CFD data generated specifically for training. Once trained, the framework provides rapid flowfield predictions, which can be used for design space exploration, design optimisation tasks, performance predictions, or even to guide CFD mesh design.

On the subsonic aerofoil test case, comparisons with the state-of-the-art convolutional neural network (CNN) suggest that the ridge function framework is able to achieve competitive predictive accuracy on unseen aerofoil designs. In addition to serving as a surrogate model, the learned ridge functions can aid understanding of the existing training data. *Sufficient summary plots* can be viewed for any given point. The reduced dimensional nature of these plots lends itself to easy visualisation, allowing for easy comparison between designs, and new physical insights. Furthermore, the local ridge functions can be integrated to obtain new ridge functions for integral quantities such as a loss or drag coefficient. This allows new interpretable surrogate models for integral quantities to be generated on demand, without the need for further pre-processing of the CFD data. Such capability is advantageous in scenarios where the long-term storage of the full CFD dataset is problematic. The ridge functions can also provide *insensitivity* information, informing us how a design can be altered without influencing the flow at a given location.

The independent nature of each ridge function means that their training can be viewed as an embarrassingly parallel task. This makes the ridge function framework trivial to implement in a parallel fashion, allowing for excellent scaling with problem size. However, for larger problems, computing ridge functions for every single grid point is undesirable. Instead, ridge functions are obtained at a much smaller number of randomly subsampled points. The flow physics encoded within covariance matrices, computed from the training data, can be used to upsample the ridge functions' predictions back to the rest of the flowfield. The computation of the covariance matrices, and the subsequent use of them, involves a variety of linear algebra operations which can be implemented efficiently. For the three dimensional transonic wing test case, sampling rates around 5% are achieved with minimal additional errors, allowing for significant reductions in training time.

The transonic wing test case demonstrates how ridge functions can be used to estimate flowfields with varying boundary conditions. In such cases, there is a danger of extrapolating too far from the training data, and the underlying polynomials' predictions can not be trusted here. If extrapolation is important, there is scope to replace the polynomials used in this work with alternative models. In addition to exploring this further, a worthwhile area of future work would be to combine the approaches used in the two test cases, incorporating geometric and boundary condition input parameters together.

## Acknowledgments

This research was supported in part through computational resources provided by The Alan Turing Institute and with the help of a generous gift from Microsoft Corporation. The authors were supported by The UKRI Strategic Priorities Fund under the EPSRC Grant EP/T001569/1, particularly the "Digital twins for complex systems engineering" theme within that grant and The Alan Turing Institute, and by the Lloyd's Register Foundation-Alan Turing Institute programme on Data-Centric Engineering under the LRF grant G0095.

## A Convolutional Neural Network Setup

This appendix provides more details on the convolutional neural network setup in Figure 5. This network is a modified form of that proposed by Thurey et al. [14], which is based upon the U-Net architecture first proposed by Ronneberger et al. [49]. Since the boundary conditions are scalar values, it might appear wasteful to repeat the values over the entirety of the  $128 \times 128 \times 3$  boundary condition input channels, and some CNN flowfield prediction frameworks such as those of Bhatnagar et al. [13] inject the scalar values into the feature vector instead. However, Thurey et al. [14] claim that, while the network would eventually propagate the boundary information via the convolutional bow tie structure, specifying the redundant boundary condition information with skip connections allows for a more efficient training process.

The general design features of the convolutional blocks from Thurey et al. [14] are retained; in the encoder leaky ReLU activations functions are chosen to avoid the *dying ReLU problem*[57], whilst standard ReLU functions are selected for the decoder, both with a slope of 0.2. Convolutional filter kernels of various sizes and strides are then used for encoding and decoding (See Fig. 5), with nearest neighbour upsampling followed by a regular convolution for the decoder blocks. To help mitigate overfitting, batch normalisation is used in all the blocks except for the first and last ones (1c and 1d), in addition to having a slight dropout rate of 0.01 for all layers.

### A.1 Learning process

To obtain the the weights (and biases) of the network, the learning process uses the *Adam optimiser*[58] to minimise the Huber loss function given in (16), with the loss averaged over all  $M_{train}$  number of training designs and  $N = 128^2$

grid points. A learning rate of  $4 \times 10^{-4}$  is used, which is then reduced by a factor of ten after 500 epochs, and a batch size of ten is chosen.




Figure 17: Convergence history for the 488k parameter CNN trained on the  $\alpha_\infty = 10^\circ$  dataset, with  $M_{train} = 500$ .

The learning curve for the 488k parameter network is plotted in Figure 17; the convergence behaviour appears to be satisfactory, and the small difference between train and test loss suggests the network isn't suffering from excessive overfitting. Similar behaviour is observed for the other network sizes explored in this paper.

## A.2 Effect of network size

To allow for a fair comparison between the ridge function framework and the CNN, the performance for various size/complexity networks is examined in order to find an optimal size. The network's size is altered by changing the number of channels, with the number of channels in the  $i^{th}$  layer given by  $C_i = 2^e c_i$ . The constants  $c_i$  are defined in Table 2, and the network's size is varied by adjusting the exponent  $e$ , with values of  $e = \{2, 3, 4, 5\}$  tested. The variation in performance

Table 2: Multiplier  $c$  used to set the number of channels in each layer of the convolutional neural network. The number of channels in the  $i^{th}$  layer is given by  $C_i = 2^e c_i$ , where the exponent  $e$  is a scalar parameter used to set the model complexity.

Encoder	Layer	1c	2c	3c	4c	5c	6c	7c
	$c_i$	1	2	2	4	8	8	8
Decoder	Layer	1d	2d	3d	4d	5d	6d	7d
	$c_i$	8	16	16	8	4	4	4

The normalised MAE, defined in (17), is measured on the training and test sets for the four different network sizes, with the results plotted in Figure 18. Generally, the larger networks achieve lower predictive errors. However, this is at the cost of longer training times (Fig. 18d), with the 1.94M parameter network taking up to 80 minutes to train. The larger networks also exhibit more stable predictions, with less variance in their MAE scores across the three train/test folds. To achieve a reasonable compromise between accuracy and training cost, the 488k parameter is chosen as the benchmark for comparisons elsewhere in this paper. A similar pattern between MAE, training times, and the number of training designs is also observed. The chosen value of  $M_{train} = 500$  is seen to offer a good compromise between low MAE scores and short training time.

## B Integrating Ridge Functions

This appendix presents a procedure for identifying a dimension reducing subspace for the mass-averaged pressure loss coefficient  $\bar{Y}_p$ , defined in (20). A more general treatment on the subspaces of an *integral of a scalar field quantity* can be found in Wong et al. [24]. For simplicity, we stick to a two dimensional domain, but the approach is readily generalisable to three dimensions. Recalling (20), we wish to obtain the mass-averaged loss coefficient  $\bar{Y}_p$ , at a given




Figure 18: Prediction errors and training time for different size convolutional neural networks, for the  $\alpha_\infty = 10^\circ$  dataset. Solid lines indicate test errors, and dashed lines training errors. Results are averaged over three train/test folds, and the error bars denote the standard deviation in  $MAE_\sigma$  across these folds. The size of the network is altered by adjusting the number of channels via the exponent  $e$ .

axial chord location  $C_x$

$$\bar{Y}_p = \frac{1}{\dot{m}} \int_{C_1} u(s_1, s_2) Y_p(s_1, s_2) ds_2, \quad (23)$$

where the mass flow rate across the plane,  $\dot{m} = \int_{C_1} u(s_1, s_2) ds_2$ , is taken to be a constant here<sup>7</sup>. The local axial velocity and loss coefficient are approximated with two ridge approximations

$$\begin{aligned} u(s_1, s_2) &\approx g_{s_1, s_2}(\mathbf{W}_{s_1, s_2}^T \mathbf{x}) \\ Y_p(s_1, s_2) &\approx h_{x, y}(\mathbf{U}_{s_1, s_2}^T \mathbf{x}). \end{aligned} \quad (24)$$

Here the input vector  $\mathbf{x} \in \mathbb{R}^d$  has its usual meaning, whilst  $\mathbf{W}^T \in \mathbb{R}^{n_w \times d}$  and  $\mathbf{U}^T \in \mathbb{R}^{n_u \times d}$  denote the subspaces for axial velocity and loss coefficient respectively, where naturally  $n_w \ll d$  and  $n_u \ll d$ . Substituting (24) into (23) yields

$$\bar{Y}_p = \rho_\infty \int_{C_1} g_{s_1, s_2}(\mathbf{W}_{s_1, s_2}^T \mathbf{x}) h_{s_1, s_2}(\mathbf{U}_{s_1, s_2}^T \mathbf{x}) ds_2. \quad (25)$$

To identify the dimension reducing subspace for  $\bar{Y}_p$  with respect to its constituent input parameters  $\mathbf{x}$ , we need to compute the *averaged outer product of its gradient* [39] – a covariance matrix given by

$$\mathbf{C} = \int_{\mathbb{R}^d} (\nabla_{\mathbf{x}} \bar{Y}_p(\mathbf{x})) (\nabla_{\mathbf{x}} \bar{Y}_p(\mathbf{x}))^T \omega(\mathbf{x}) d\mathbf{x}, \quad (26)$$

where  $\omega(\mathbf{x})$  denotes the distribution associated with the input parameters  $\mathbf{x}$ . Assuming the existence of a design of experiment with  $M$  distinct CFD evaluations for different input parametrisations, and taking  $\omega(\mathbf{x})$  to be a uniform distribution here, we can approximate  $\mathbf{C}$  with its finite sample estimate

$$\mathbf{C} \approx \frac{1}{M} \sum_{m=1}^M (\nabla_{\mathbf{x}} \bar{Y}_p(\mathbf{x}_m)) (\nabla_{\mathbf{x}} \bar{Y}_p(\mathbf{x}_m))^T. \quad (27)$$

<sup>7</sup>The standard deviation of  $\dot{m}$  across the  $C_1$  plane is less than 0.2% of the mean.

The gradient evaluations can be written as

$$\begin{aligned}\nabla_{\mathbf{x}} \bar{Y}_p(\mathbf{x}_m) &= \rho_\infty \int_{\mathbb{R}^d} \frac{\partial}{\partial \mathbf{x}_m} [g_{s_1, s_2}(\mathbf{W}_{s_1, s_2}^T \mathbf{x}_m) h_{s_1, s_2}(\mathbf{U}_{s_1, s_2}^T \mathbf{x}_m)] d\mathbf{s}_2 \\ &= \rho_\infty \int_{\mathbb{R}^d} [\mathbf{W}_{s_1, s_2} \nabla_{\mathbf{w}} g_{s_1, s_2}(\mathbf{w}_m) h_{s_1, s_2}(\mathbf{u}_m) + \mathbf{U}_{s_1, s_2} \nabla_{\mathbf{u}} h_{s_1, s_2}(\mathbf{u}_m) g_{s_1, s_2}(\mathbf{w}_m)] d\mathbf{s}_2,\end{aligned}\quad (28)$$

where  $\mathbf{w}_m = \mathbf{W}^T \mathbf{x}_m$  and  $\mathbf{u}_m = \mathbf{U}^T \mathbf{x}_m$  are the reduced dimensional co-ordinates for the  $m^{\text{th}}$  design in the  $Y_p(s_1, s_2)$  and  $u(s_1, s_2)$  subspaces respectively. If we further assume that spatially, there are  $N$  nodes along which this integration has to be performed, in which case we have the set

$$\{g_i, h_i, \mathbf{W}_i, \mathbf{U}_i, \nabla_{\mathbf{w}} g_i, \nabla_{\mathbf{u}} h_i\}_{i=1}^N, \quad (29)$$

we approximate

$$\nabla_{\mathbf{x}} \bar{Y}_p(\mathbf{x}_m) \approx \rho_\infty \sum_{i=1}^N [\mathbf{W}_i \nabla_{\mathbf{w}} g_i h_i + \mathbf{U}_i \nabla_{\mathbf{u}} h_i g_i] \varphi_i \quad (30)$$

where  $\varphi_i$  denotes appropriately selected quadrature weights depending on the spatial location of the nodal centres. Plugging this into (27) yields

$$\begin{aligned}\mathbf{C} &\approx \mathbb{E} \left[ \left( \rho_\infty \sum_{i=1}^N [\mathbf{W}_i \nabla_{\mathbf{w}} g_i h_i + \mathbf{U}_i \nabla_{\mathbf{u}} h_i g_i] \varphi_i \right) \left( \rho_\infty \sum_{k=1}^N [\mathbf{W}_k \nabla_{\mathbf{w}} g_k h_k + \mathbf{U}_k \nabla_{\mathbf{u}} h_k g_k] \varphi_k \right)^T \right] \\ &= \rho_\infty^2 \sum_{i=1}^N \sum_{k=1}^N \varphi_i \varphi_k \mathbb{E} [h_i h_k \mathbf{W}_i \nabla_{\mathbf{w}} g_i \nabla_{\mathbf{w}} g_k^T \mathbf{W}_k^T] + \rho_\infty^2 \sum_{i=1}^N \sum_{k=1}^N \varphi_i \varphi_k \mathbb{E} [h_i g_k \mathbf{W}_i \nabla_{\mathbf{w}} g_i \nabla_{\mathbf{u}} h_k^T \mathbf{U}_k^T] + \\ &\quad \rho_\infty^2 \sum_{i=1}^N \sum_{k=1}^N \varphi_i \varphi_k \mathbb{E} [g_i h_k \mathbf{U}_i \nabla_{\mathbf{u}} h_i \nabla_{\mathbf{w}} g_k^T \mathbf{W}_k^T] + \rho_\infty^2 \sum_{i=1}^N \sum_{k=1}^N \varphi_i \varphi_k \mathbb{E} [g_i g_k \mathbf{U}_i \nabla_{\mathbf{u}} h_i \nabla_{\mathbf{u}} h_k^T \mathbf{U}_k^T]\end{aligned}\quad (31)$$

This equation is easy to implement, as long as the gradients  $\nabla_{\mathbf{w}} g$  and  $\nabla_{\mathbf{u}} h$  can be computed. When applying this approach in Section 4.3,  $g$  and  $h$  are taken to be the underlying polynomials of the ridge function approximations, rather than the posterior mean of the ridge kernels as is the case in the rest of the paper. This is done since the gradients of orthogonal polynomials are readily available. Finally, an eigendecomposition of  $\mathbf{C}$  is performed, and the leading eigenvectors with the largest eigenvalues are chosen to form  $\mathbf{W}_{\text{int}}$ , the dimension reducing subspace for  $\bar{Y}_p$ . In the present case, the first eigenvalue was found to contribute 99.3% of the trace. Therefore, a one dimensional dimension reducing subspace was deemed to be sufficient.

## References

- [1] T. Kobayashi and M. Tsubokura. CFD application in automotive industry. In *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, pages 285–295. Springer Berlin Heidelberg, 2009. doi: 10.1007/978-3-540-70805-6\_22. URL [https://doi.org/10.1007/978-3-540-70805-6\\_22](https://doi.org/10.1007/978-3-540-70805-6_22).
- [2] Fei Ding, Ahsan Kareem, and Jiawei Wan. Aerodynamic tailoring of structures using computational fluid dynamics. *Structural Engineering International*, 29(1):26–39, January 2019. doi: 10.1080/10168664.2018.1522936.
- [3] R. Ramponi and B. Blocken. CFD simulation of cross-ventilation for a generic isolated building: Impact of computational parameters. *Building and Environment*, 53:34–48, July 2012. doi: 10.1016/j.buildenv.2012.01.004. URL <https://doi.org/10.1016/j.buildenv.2012.01.004>.
- [4] Goutam Kumar Saha, Md. Hayatul Islam Maruf, and Md. Rakibul Hasan. Marine propeller modeling and performance analysis using CFD tools. In *8th BSME International Conference on Thermal Engineering*. AIP Publishing, 2019. doi: 10.1063/1.5115883. URL <https://doi.org/10.1063/1.5115883>.
- [5] Talib Dbouk and Dimitris Drikakis. On coughing and airborne droplet transmission to humans. *Physics of Fluids*, 32(5):053310, May 2020. doi: 10.1063/5.0011960. URL <https://doi.org/10.1063/5.0011960>.
- [6] Talib Dbouk and Dimitris Drikakis. On respiratory droplets and face masks. *Physics of Fluids*, 32(6):063303, June 2020. doi: 10.1063/5.0015044. URL <https://doi.org/10.1063/5.0015044>.
- [7] Krithika Manohar, Bingni W. Brunton, J. Nathan Kutz, and Steven L. Brunton. Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns. *IEEE Control Systems*, 38(3):63–86, June 2018. doi: 10.1109/mcs.2018.2810460. URL <https://doi.org/10.1109/mcs.2018.2810460>.

- [8] Stefano Discetti, Marco Raiola, and Andrea Ianiro. Estimation of time-resolved turbulent fields through correlation of non-time-resolved field measurements and time-resolved point measurements. *Experimental Thermal and Fluid Science*, 93:119–130, May 2018. doi: 10.1016/j.expthermflusci.2017.12.011. URL <https://doi.org/10.1016/j.expthermflusci.2017.12.011>.
- [9] N. Benjamin Erichson, Lionel Mathelin, Zhewei Yao, Steven L. Brunton, Michael W. Mahoney, and J. Nathan Kutz. Shallow neural networks for fluid flow reconstruction with limited sensors. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 476(2238):20200097, June 2020. doi: 10.1098/rspa.2020.0097. URL <https://doi.org/10.1098/rspa.2020.0097>.
- [10] Maziar Raissi and George Em Karniadakis. Hidden physics models: Machine learning of nonlinear partial differential equations. *Journal of Computational Physics*, 357:125–141, March 2018. doi: 10.1016/j.jcp.2017.11.039.
- [11] Lionel Agostini. Exploration and prediction of fluid dynamical systems using auto-encoder technology. *Physics of Fluids*, 32(6):067103, June 2020. doi: 10.1063/5.0012906. URL <https://doi.org/10.1063/5.0012906>.
- [12] Xiaoxiao Guo, Wei Li, and Francesco Iorio. Convolutional neural networks for steady flow approximation. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, August 2016. doi: 10.1145/2939672.2939738.
- [13] Saakaar Bhatnagar, Yaser Afshar, Shaowu Pan, Karthik Duraisamy, and Shailendra Kaushik. Prediction of aerodynamic flow fields using convolutional neural networks. *Comput. Mech.*, 2019. ISSN 01787675. doi: 10.1007/s00466-019-01740-0.
- [14] Nils Thuerey, Konstantin Weißenow, Lukas Prantl, and Xiangyu Hu. Deep Learning Methods for Reynolds-Averaged Navier–Stokes Simulations of Airfoil Flows. *AIAA J.*, 58(1):25–36, jan 2020. ISSN 0001-1452. doi: 10.2514/1.j058291.
- [15] Jonathan Tompson, Kristofer Schlachter, Pablo Sprechmann, and Ken Perlin. Accelerating Eulerian fluid simulation with convolutional networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3424–3433, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [16] Vinothkumar Sekar, Mengqi Zhang, Chang Shu, and Boo Cheong Khoo. Inverse design of airfoil using a deep convolutional neural network. *AIAA Journal*, 57(3):993–1003, March 2019. doi: 10.2514/1.j057894. URL <https://doi.org/10.2514/1.j057894>.
- [17] Xiaowei Jin, Peng Cheng, Wen-Li Chen, and Hui 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(4):047105, April 2018. doi: 10.1063/1.5024595. URL <https://doi.org/10.1063/1.5024595>.
- [18] Terrence J. Sejnowski. The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences*, page 201907373, January 2020. doi: 10.1073/pnas.1907373117. URL <https://doi.org/10.1073/pnas.1907373117>.
- [19] Q. Zhang, Y. N. Wu, and S. Zhu. Interpretable convolutional neural networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8827–8836, 2018.
- [20] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.
- [21] Allan Pinkus. *Ridge functions*, volume 205. Cambridge University Press, 2015.
- [22] Zachary del Rosario, Paul Constantine, and Gianluca Iaccarino. Developing design insight through active subspaces. In *19th AIAA Non-Deterministic Approaches Conference*. American Institute of Aeronautics and Astronautics, January 2017. doi: 10.2514/6.2017-1090. URL <https://doi.org/10.2514/6.2017-1090>.
- [23] Ashley D Scillitoe, Bryn Ubald, Pranay Seshadri, and Shahrokh Shahpar. Design space exploration of stagnation temperature probes via dimension reduction. In *Proc. ASME Turbo Expo*, 2020.
- [24] Chun Yui Wong, Pranay Seshadri, Geoffrey T Parks, and Mark Girolami. Embedded ridge approximations. *Computer Methods in Applied Mechanics and Engineering*, 372:113383, 2020.
- [25] Yuyang Wang, Kenji Shimada, and Amir Barati Farimani. Airfoil gan: Encoding and synthesizing airfoils for aerodynamic-aware shape optimization, 2021.
- [26] Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.



- [27] Ker-Chau Li. On principal hessian directions for data visualization and dimension reduction: Another application of stein’s lemma. *Journal of the American Statistical Association*, 87(420):1025–1039, 1992.
- [28] Bing Li, Hongyuan Zha, Francesca Chiaromonte, et al. Contour regression: a general approach to dimension reduction. *The Annals of Statistics*, 33(4):1580–1616, 2005.
- [29] R Dennis Cook and Sanford Weisberg. Sliced inverse regression for dimension reduction: Comment. *Journal of the American Statistical Association*, 86(414):328–332, 1991.
- [30] Yingcun Xia, Howell Tong, W. K. Li, and Li-Xing Zhu. An adaptive estimation of dimension reduction space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):363–410, 2002. ISSN 1467-9868. doi: 10.1111/1467-9868.03411. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/1467-9868.03411>.
- [31] Paul G Constantine, Armin Eftekhari, Jeffrey Hokanson, and Rachel A Ward. A near-stationary subspace for ridge approximation. *Computer Methods in Applied Mechanics and Engineering*, 326:402–421, 2017.
- [32] Jeffrey M Hokanson and Paul G Constantine. Data-driven polynomial ridge approximation using variable projection. *SIAM Journal on Scientific Computing*, 40(3):A1566–A1589, 2018.
- [33] Pranay Seshadri, Shaowu Yuchi, and Geoffrey T Parks. Dimension reduction via gaussian ridge functions. *SIAM/ASA Journal on Uncertainty Quantification*, 7(4):1301–1322, 2019.
- [34] Xiaoyu Liu and Serge Guillas. Dimension reduction for gaussian process emulation: An application to the influence of bathymetry on tsunami heights. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):787–812, 2017.
- [35] Pranay Seshadri, Shahrokh Shahpar, Paul Constantine, Geoffrey Parks, and Mike Adams. Turbomachinery active subspace performance maps. *Journal of Turbomachinery*, 140(4), 2018.
- [36] Pranay Seshadri, Akil Narayan, and Sankaran Mahadevan. Effectively subsampled quadratures for least squares polynomial approximations. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):1003–1023, 2017.
- [37] Chun Yui Wong, Pranay Seshadri, and Geoffrey Parks. Extremum sensitivity analysis with polynomial monte carlo filtering. *Reliability Engineering & System Safety*, page 107609, 2021.
- [38] Gene Golub and Victor Pereyra. Separable nonlinear least squares: the variable projection method and its applications. *Inverse problems*, 19(2):R1, 2003.
- [39] Paul G Constantine. *Active subspaces: Emerging ideas for dimension reduction in parameter studies*, volume 2. SIAM, 2015.
- [40] Ashley D Scillitoe, Pranay Seshadri, and Chun Yui Wong. Instantaneous flowfield estimation with gaussian ridges. In *AIAA Scitech 2021 Forum*, page 1138, 2021.
- [41] Boris Diskin, James L. Thomas, Christopher L. Rumsey, and Axel Schwöppe. Grid-convergence of reynolds-averaged navier–stokes solutions for benchmark flows in two dimensions. *AIAA Journal*, 54(9):2563–2588, September 2016. doi: 10.2514/1.j054555.
- [42] Raymond M. Hicks and Preston A. Henne. Wing design by numerical optimization. *Journal of Aircraft*, 15(7): 407–412, July 1978. doi: 10.2514/3.58379.
- [43] Thomas D. Economon, Francisco Palacios, Sean R. Copeland, Trent W. Lukaczyk, and Juan J. Alonso. SU2: An open-source suite for multiphysics simulation and design. *AIAA Journal*, 54(3):828–846, March 2016. doi: 10.2514/1.j053813.
- [44] P. Spalart and S. Allmaras. A one-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics, January 1992. doi: 10.2514/6.1992-439. URL <https://doi.org/10.2514/6.1992-439>.
- [45] C. Sullivan and Alexander Kaszynski. PyVista: 3d plotting and mesh analysis through a streamlined interface for the visualization toolkit (VTK). *Journal of Open Source Software*, 4(37):1450, May 2019. doi: 10.21105/joss.01450.
- [46] V. Schmitt and F. Charpin. Pressure distributions on the onera m6 wing at transonic mach numbers. In *Experimental data base for computer program assessment. Report of the Fluid Dynamics Panel, Working Group 04*. AGARD, 1979. ISBN 92-835-1323-1. doi: 10.1007/978-0-387-74759-0\_121.
- [47] Pranay Seshadri and Geoffrey Parks. Effective-Quadratures (EQ): Polynomials for Computational Engineering Studies. *J. Open Source Softw.*, 2(11):166, 2017. doi: 10.21105/joss.00166.

- [48] Charles R. Harris, K. Jarrod Millman, St'efan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fern'andez del R'io, Mark Wiebe, Pearu Peterson, Pierre G'erard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- [49] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Lecture Notes in Computer Science*, pages 234–241. Springer International Publishing, 2015. doi: 10.1007/978-3-319-24574-4\_28. URL [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28).
- [50] Peter J. Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, March 1964. doi: 10.1214/aoms/1177703732. URL <https://doi.org/10.1214/aoms/1177703732>.
- [51] Ashley D. Scillitoe, Paul G. Tucker, and Paolo Adami. Numerical Investigation of Three-Dimensional Separation in an Axial Flow Compressor: The Influence of Freestream Turbulence Intensity and Endwall Boundary Layer State. *J. Turbomach.*, 139(2):021011 1–10, oct 2016. ISSN 0889-504X. doi: 10.1115/1.4034797.
- [52] Paul Constantine, Ryan Howard, Andrew Glaws, Zachary Grey, Paul Diaz, and Leslie Fletcher. Python active-subspaces utility library. *Journal of Open Source Software*, 1(5):79, 2016. doi: 10.21105/joss.00079. URL <https://doi.org/10.21105/joss.00079>.
- [53] Chun Yui Wong, Pranay Seshadri, Ashley Scillitoe, Andrew B. Duncan, and Geoffrey Parks. Blade envelopes part i: Concept and methodology, 2021.
- [54] Chun Yui Wong, Pranay Seshadri, Ashley Scillitoe, Bryn Noel Ubald, Andrew B. Duncan, and Geoffrey Parks. Blade envelopes part ii: Multiple objectives and inverse design, 2020.
- [55] Wei Wang, Jun Wang, Hui Liu, and Bo yan Jiang. CFD prediction of airfoil drag in viscous flow using the entropy generation method. *Mathematical Problems in Engineering*, 2018:1–15, 2018. doi: 10.1155/2018/4347650. URL <https://doi.org/10.1155/2018/4347650>.
- [56] Heinz Herwig and Bastian Schmandt. How to determine losses in a flow field: A paradigm shift towards the second law analysis. *Entropy*, 16(6):2959–2989, May 2014. doi: 10.3390/e16062959. URL <https://doi.org/10.3390/e16062959>.
- [57] Lu Lu. Dying ReLU and initialization: Theory and numerical examples. *Communications in Computational Physics*, 28(5):1671–1706, June 2020. doi: 10.4208/cicp.oa-2020-0165. URL <https://doi.org/10.4208/cicp.oa-2020-0165>.
- [58] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014.