Available online at www.eccomasproceedia.org
Eccomas Proceedia EUROGEN (2023) 59-72

ECCOMAS
Proceedia

EUROGEN 2023
15th ECCOMAS Thematic Conference on
Evolutionary and Deterministic Methods for Design, Optimization and Control
N. Gauger, K. Giannakoglou, M. Papadrakakis, J. Periaux (eds.)
Chania, Crete Greece, 1–3 June 2023

TOWARDS AERODYNAMIC SHAPE OPTIMISATION BY MANIFOLD LEARNING AND NEURAL NETWORKS

Rodrigo Castellanos[1], Javier Nieto-Centenero[1,3], Alejandro Gorgues[1,4], Stefano Discetti[2], Andrea Ianiro[2] and Esther Andrés[1]

[1]Theoretical and Computational Aerodynamics Branch, Flight Physics Department
Instituto Nacional de Técnica Aeroespacial (INTA)
Ctra. Ajalvir Km.4 Torrejón de Ardoz 28850, España
e-mail: {rcasgar, jniecen, gorguesva, eandper}@inta.es

[2]Department of Aerospace Engineering
Universidad Carlos III de Madrid
Avda. Universidad 30, Leganés 28911, España
e-mail: {sdiscett, aianiro}@ing.uc3m.es

[3] Escuela Técnica superior de Ingeniería Aeronáutica y del Espacio (ETSIAE)
Universidad Politécnica de Madrid (UPM)
28223 Madrid, España.

[4] Escuela Politécnica Superior
Universidad de Alcalá de Henares (UAH)
28805 Alcalá de Henares, España.

Abstract. *Multidisciplinary design and optimisation in aeronautics is a time-consuming process that requires extensive exploration of various configurations using costly simulations. In this study, we propose the development of a data-driven surrogate model to serve as a predictive tool during the preliminary design process. We utilise the Isometric Feature Mapping (Isomap) technique to identify a low-dimensional manifold where training dataset information is encoded. Our algorithm is fed with distributions of pressure coefficients on 3D wings, featured by varying geometry and under different flight conditions. We create a two-block decoder through a Deep Neural Network (DNN), which maps the design parameters to the manifold space. Additionally, we use back-mapping based on the k-nearest neighbour to reconstruct the pressure coefficient distributions. This technique, called Isomap+DNN, is compared to direct interpolation in the design parameter space. Our results indicate that, although the performance of the two methods is similar, Isomap+DNN is superior in predicting configurations with stronger compressibility effects.*

Keywords: Surrogate model; machine learning; reduced-order model; Isomap; deep neural network; aerodynamics.

Peer-review under responsibility of the organizing committee of EUROGEN 2023.
doi: 10.7712/140123.10190.18887

1 INTRODUCTION

The contemporary aerodynamic design heavily relies on computational fluid dynamics (CFD) simulations, which are computationally expensive and time-consuming. Currently, preliminary and intermediate stages of technological development rely on simplified models that provide a satisfactory quality of aerodynamic data while costing significantly less than high-accuracy methods, such as direct numerical simulations (DNS). Recent advancements in the field of reduced-order modelling (ROM) and machine learning (ML) have introduced novel models that aim to facilitate a more efficient design process based on surrogates to reduce the costs associated with expensive computational resources [1].

The optimisation of aerodynamic design using surrogate models relies on iterative model refinement [2]. These surrogate-based optimisation techniques have proven useful in a variety of aerodynamic shape optimisation applications [3]. The effectiveness of surrogate-based optimisation is strongly related to the performance of the surrogate model, which has been enhanced through the use of ML. Additionally, ML facilitated dimensionality reduction in shape design variables through modal shape parameterisation. Recently, a surrogate-based optimisation method utilising a reduced set of wing modes was demonstrated to be nearly as efficient and effective as CFD-based optimisation utilising all design variables [4]. This study focuses on the development of efficient and robust surrogate models combining adequate dimensionality-reduction techniques and regression models.

Notwithstanding the intrinsic nonlinear nature of transonic phenomena, these flows are characterised by recurring flow patterns and physical features that can be acquired through simulations or experimental data. Proper Orthogonal Decomposition (POD) [5], also known in the field of statistics as Principal Component Analysis (PCA) [6], is one of the most popular methods in both fluid mechanics and aeronautics for dimensionality reduction. Its applicability to surrogate models for aerodynamic data has been investigated by integrating it with an interpolator in the low-dimensional space (i.e. the basis composed of the POD modes) [7]. Within the field of statistical learning [8], a wide portfolio of dimensionality-reduction algorithms is available. The ability to compute low-dimensional data representations presents an opportunity to identify embedded manifolds of the flows under study. In this investigation, we employ a nonlinear manifold learner and integrate it with a regression model to predict transonic three-dimensional flows. Notably, manifold learning encompasses techniques such as locally linear embedding (LLE) [9], multi-dimensional scaling (MDS) [10] or isometric feature mapping (Isomap) [11]. The latter is the method considered in this study. The primary objective of manifold learning is to identify the surface (i.e., manifold) upon which the data resides or which closely approximates the dataset. The manifold serves as a geometric representation of the intrinsic relationships that exist between snapshots. Intriguingly, the evolution of various types of high-dimensional data can be simplified within the manifold, allowing for the identification of the principal characteristics underlying such evolution.

In contrast to conventional linear dimensionality-reduction techniques, such as POD, Isomap assumes that the high-dimensional data lie on a low-dimensional manifold. Isomap identifies the nonlinear features underlying complex natural observations by computing geodetic distances. While Isomap has been applied in various fields, its use in fluid mechanics and aeronautics has been limited, with few studies identifying manifolds from flow-visualization data [12], in the combustion field [13], and, recently, to understand the physics in shear flows [14]. Franz et al. [15] developed *Isomap+I*, a parametric ROM to predict shock waves on a 3D wing in the transonic regime. This methodology represents the state-of-the-art application of Isomap

to surrogate modelling in aerodynamics. A recent study on aerodynamic shape optimisation of a hypersonic vehicle [16] demonstrated that Isomap and LLE could lead in the regions near the shock waves to more accurate predictions than POD-based models. Similarly, our recent study on transonic wings [17] highlighted the superior performance of a surrogate model combining Isomap and Deep Neural Networks (DNN) to predict pressure distributions compared to POD-based methods or direct DNN regression without dimensionality reduction.

This contribution presents a proposal for a surrogate model aimed at optimizing the aerodynamic shape of transonic wings. The model is based on manifold learning applied to aerodynamic data obtained from CFD simulations. The database spans different flight conditions (Mach number and angle of attack) and geometry (through sweep angle, thickness, and span relative to a baseline configuration). To extract the nonlinear features of the data, the Isomap dimensionality reduction technique is combined with a Deep Neural Network (DNN) [18], which has demonstrated extraordinary performance in fitting data. Previous studies have demonstrated the applicability of artificial neural networks in the design process of airfoils [19] and nonlinear system identification techniques, frequently used to model fluid systems [20]. Additionally, surrogate models using artificial neural networks have been previously applied to predict the global force coefficient of wings [21]. Several algorithms have been proven sturdy for aerodynamic data prediction [22]; however, the combination of dimensionality reduction techniques with DNNs remains relatively unexplored. The surrogate models are developed in the low-data limit, where only a limited number of simulations are feasible, such as in the early design phase. The performance of each method is assessed using a database of CFD simulations of a 3D wing in the transonic regime, with the strengths and weaknesses of each approach described.

The paper is structured as follows: First, the methodology is presented in Section 2, starting with a brief theoretical background of the Isomap method, and is followed by a concise description of the Deep Neural Network regression model. Thereafter, the performance of the proposed Isomap+DNN model combining flight conditions and geometric variations as design features is evaluated. A set of results is presented in Section 3, for customary cases and finally, the conclusions are drawn in Section 4.

2 DEVELOPMENT OF THE SURROGATE MODEL

This section presents a detailed account of the methodology and database employed in this work. It commences with an outline of the database and the geometry of the model under consideration. Subsequently, the Isomap low-dimensional embedding is presented with a detailed analysis of the relevance of each feature. Additionally, a brief description of the used Deep Neural Network is provided.

2.1 Database and test case

The selected test case is a database of CFD simulations of the XRF1 model, which is a test case provided by Airbus to show the application of different technologies to a long-range wide-body aircraft. The simulations build upon a simplified version of the XRF1 geometry, excluding the engine and both the horizontal and vertical tailplanes, while allowing several geometric variations with respect to the baseline design (from now on denoted by the subscript '0'). The considered geometric variations are expressed in terms of the wing-span ratio b/b_0, increment of sweep angle $\Delta\Lambda$ $(= \Lambda - \Lambda_0)$, and the thickness ratio of the wing t/t_0. A total of 60 different wing geometries are simulated for 49 combinations of Mach number M and angle of attack α spanning over the design region of the flight envelope. A summary of the flow

Flight condition		Geometric features	
α	2°, 3°, 4°, 5°, 6°, 7°, 8°	b/b_0	0.9, 1.0, 1.1, 1.2
M	0.70, 0.74, 0.78, 0.80, 0.82, 0.84, 0.86	$\Delta\Lambda$	-5°, -2°, 0°, +2°, +5°
		t/t_0	0.9, 1.0, 1.1

Table 1: Definition of the different flight conditions and geometric features used to build the database. A total of 49 flight conditions and 60 wing geometries, summing up 2940 simulations.

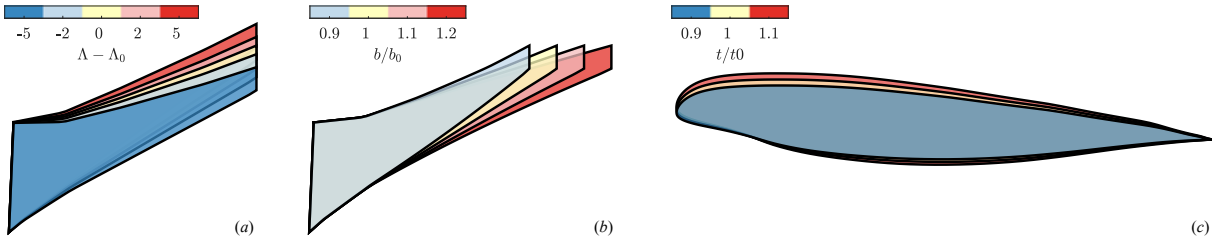

Figure 1: Schematic representation of the different geometric parameters used to generate the database. The geometric parameters describe an alteration of the nominal geometry, denoted with subscript 0 (a) Increment of sweep angle $\Delta\Lambda = \Lambda - \Lambda_0$ in degrees, (b) scaling of the wing span b/b_0, and (c) scaling of thickness ratio t/t_0.

and geometric features and their range is provided in table 1 and the geometric variations are sketched in figure 1. Additionally, it is noted that the load factor ($n_z = L/W$, being L the lift and W the weight of the aircraft) is fixed at $n_z = 1$ for all the cases.

The aerodynamic data is obtained by *RapidCFD* computations based on the BLWF code [23] developed by researchers at the Central Aerohydrodynamic Institute (TsAGI) [24]. This code solves a boundary-value problem for the full velocity-potential equation and viscosity is considered in the boundary layer approximation with a fixed position of the laminar-to-turbulent transition. The method simulates the occurrence of local supersonic areas and shock waves, provides computations of flows with small separation zones and is reliably verified. A total of 2940 simulations are performed combining the flight conditions and geometric alterations described in table 1. For each of these cases, the distribution of the pressure coefficient C_p on the wing surface is computed, which is interpolated in a common structured grid featured by 6519 points, covering 41 wing sections along the span with 159 grid points each.

The study is performed in the *low-data limit*, emulating a preliminary design phase in which either CFD simulations or wind-tunnel experiments are not affordable. Consequently, the 2940 cases composing the database split into two sets: the *training* set (30% of the database, 882 cases) used to compute the Isomap embedding and to train the DNN; and the *test* set with all the remaining cases (70% of the database, 2058 cases) to challenge the trained surrogate model, assessing its capability to interpolate C_p within the whole design space. The cases within the *test* set are never considered for either computing the low-dimensional embedding or training the DNN coefficient. For now on, four cases within the *test* set are chosen for visualisation purposes. The so-called *visualisation* cases are chosen due to the presence of shock waves, which are particularly challenging for the surrogate models.

It has to be noted that the optimisation of the aerodynamic shape based on the proposed geometric features is considerably affected by the flight condition under consideration. From a qualitative analysis of the CFD data, it is already possible to assess that the tuple (M, α) drives the pressure distribution over the wing surface, with smaller effects coming from the geometric variations, as later discussed. The pressures on the upper surface have drastic changes due to shock wave patterns as M increases, imposing a challenge for the surrogate model due to the

high nonlinear dependencies of the dataset.

2.2 Manifold Learning via Isomap

Isometric feature Mapping (Isomap) is a manifold learner with nonlinear features. The purpose of this investigation is to employ Isomap for embedding CFD data into a low-dimensional space [11]. This encoding procedure is entirely data-driven, and its objective is to unveil a latent low-dimensional space that facilitates the correlation of the new coordinates with the main features of the data. In prior studies [15, 17], Isomap has been utilised in surrogate models to carry out regression within the low-dimensional subspace, outperforming other linear techniques such as POD.

Consider N data samples in the high-dimensional space $\mathbb{R}^P$. Let $\mathbf{X} \in \mathbb{R}^{P\times N}$ be the data matrix containing the high-dimensional data and $\mathbf{x}^i \in \mathbb{R}^P$ be each of its rows, for $i = 1, \ldots, N$. The dataset in $\mathbf{X}$ is in our case characterised by the presence of shock waves, and being able to extract a meaningful small number of coordinates that captures the main characteristics of the flow is challenging. Isomap computes the low-dimensional embedding of the data points that best preserve the geodesic distances measured in the high-dimensional input space. This algorithm first relies on a conventional k-nearest neighbour (kNN) search to compute the matrix of Euclidean distances $d_{\mathbf{X}}(i, j)$ between data points $\mathbf{x}_i$ and $\mathbf{x}_j$ for all $i, j = 1, \ldots, N$ to identify the k closest observations to $\mathbf{x}_i$. Then the neighbouring graph G over these data points is computed such that two nodes i and j are connected by an edge of weight $d_{\mathbf{X}}(i, j)$ if they are neighbours. Given the pairs of vertices within G, Floyd's algorithm [25] is invoked to calculate the shortest paths between them, creating the matrix $\mathbf{D}_G$. Finally, the low-dimensional embedding is obtained $\mathbf{\Gamma} \in \mathbb{R}^{N\times p}$, $p << P$, using a classical Multi-Dimensional Scaling (MDS) [10] on the matrix of shortest path distances $\mathbf{D}_G$, so that the Euclidean pairwise distances resemble those in the neighbouring graph $d_{\mathbf{G}}(i, j)$.

From C_p to the manifold: Isomap encoder

The CFD data is provided in the form of N data vectors, containing the pressure coefficient on the wing surface, being N the number of cases within the training set. Each C_p distribution is an observation (point) in the high-dimensional space $\mathbb{R}^P$, being P the number of gridpoints. The definition of the number of neighbours k to compute the matrix G, and the dimension of the embedding p are selected based on the *residual variance* (RV) [11]. This metric is the ratio of the residual sum of squares to the total sum of squares based on the matrix of Euclidean distances between each pair of points in the low-dimensional embedding $\mathbf{D}_{\mathbf{\Gamma}}$ and the shortest distance matrix $\mathbf{D}_G$. Since the value of RV quantifies the information that remains unexplained by the low-dimensional embedding of the original data, the objective is to minimise it for a given p and k. Whereas the appropriate p mainly depends on the data and its features, the choice of k must guarantee the preservation of the structure and convexity of the manifold. The approach by Samko et al. [26] was followed to define a valid range of exploration for k, and the dimension of the embedding is explored in the range $p \in [2, 5]$ based on the RV value.

In a preliminary analysis, it was decided to fix the number of Isomap coordinates to the number of design features used to build the database, which includes the flight condition and the geometric variations ($p = 5$). The *correlation heatmap* between the high-dimensional features and the Isomap coordinates is included in table 2. Additionally, the relevance of the high-dimensional features and the Isomap coordinates on the CFD data is assessed by computing the correlation for the average and root mean square of the pressure coefficient ($\overline{C}_p$ and $C_{p_{rms}}$,

	M	α	$\Delta\Lambda$	t/t_0	b/b_0	$\overline{C}_p$	$C_{p_{rms}}$
γ_1	-0.75	-0.38	0.11	-0.09	-0.17	0.69	-0.21
γ_2	-0.01	0.60	0.04	-0.04	-0.11	-0.50	0.77
γ_3	0.13	0.29	-0.03	0.01	-0.01	-0.29	0.22
γ_4	-0.13	0.22	0.02	-0.02	0.00	-0.17	0.26
γ_5	-0.06	-0.23	-0.09	0.07	0.28	0.12	0.00
$\overline{C}_p$	-0.33	-0.92	0.07	-0.08	-0.08		
$C_{p_{rms}}$	-0.12	0.93	-0.02	0.02	0.03		

Table 2: Correlation heatmap between aerodynamic and geometric features, and the isomap variables.

respectively). Based on the correlation heatmap, it is observed that the main feature affecting the pressure coefficient is α, followed by M. The correlation with the geometric-variation features is very low, showing a minor relevance. Interestingly, the first and second Isomap coordinates, $\gamma_1, \gamma_2 \in \Gamma$, show a strong relationship with M and α, respectively, being also the most related to C_p.

The conclusions from the correlation heatmap could be directly related to the geometric alterations and their effect on the effective Mach number and angle of attack. The modification of the sweep angle, thickness ratio and span ratio led to effects similar to a modification of the local flight condition (M, α). For instance, the increase of the sweep angle ($\Delta\Lambda > 0$) has the effect of reducing the local M, leading to a positive correlation with γ_1, differently for M which has a negative correlation with γ_1. Similarly, thickening the wing ($t/t_0 > 1$) leads to an increase of the suction peak, promoting stronger shock waves for transonic flight conditions, which explains the weak negative correlation with γ_1. On the other side, elongating the wing ($b/b_0 > 1$) reduces the induced angle of attack and the parasite drag, in agreement with a negative correlation with γ_2, contrarily to α which has a positive correlation with γ_2. Additionally, increasing b also augments compressibility effects, due to a decrease in the sweep nearby the tip of the wing, as shown in figure 1.

Based on the previous analysis, it seems that the Isomap embedding gathers the majority of the contribution from the five design features into two variables. Hence, the selection of $p = 2$ is a compromise of minimising the residual variance while keeping the least amount of dimensions in the embedding. The representation of the pressure coefficient C_p in the low-dimensional embedding is depicted in figure 2 for three different sets of data: the full database, 70% of the database, and 30% of the database (the latest case being the one used for performance estimation in the remainder of the paper). The representation of the data in the bi-dimensional embedding depicts a topologically closed surface with a strong relationship with the problem features as already evident in table 2. The colour scale in the figure shows the distribution of the data for different M, illustrating a strong correlation with γ_1. Similarly, the size of the symbols in figure 2 illustrates the value of the angle of attack, which grows in a pseudo-radial direction within the (γ_1, γ_2) plane. It is to be noted that the shape and size of the manifold are not altered by reducing the amount of data under consideration, meaning that Isomap captures the intrinsic features of the C_p distributions with few samples. The manifold in figure 2c depicts the low-dimensional representation of the training set that will be used to develop the surrogate models.

From the manifold to C_p: k-NN decoder

The Isomap algorithm identifies an appropriate mapping function that projects high-dimensional data onto a low-dimensional embedding for a fixed and finite set of provided data samples. How-

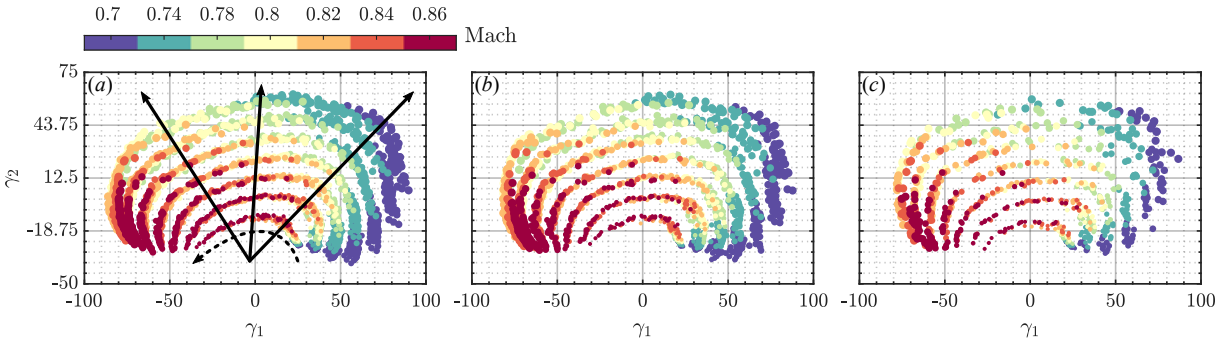

Figure 2: Isomap embedding. Representation of the pressure coefficient C_p in the low-dimensional ($\mathbb{R}^2$) embedding computed with Isomap for three subsets of the database: (a) Full database; (b) 70% of the database; (c) 30% of the database. In (a), the directions for the growing angle of attack (——) and Mach number (– –) are highlighted. The colour scale represents the Mach number while the symbol size represents the attack angle.

ever, this method does not furnish a decoding mechanism for the approximated reduced-order solution. Consequently, a backmapping procedure from the embedding to the high-dimensional solution space is necessary. The idea for the backmapping is based on LLE [9] and its different versions within the literature [15, 14, 27].

Suppose that $\mathbf{y}^i = [\gamma_1, \ldots, \gamma_p] \in \mathbb{R}^p$ is the low-dimensional counterpart of the C_p snapshots, i.e. $\mathbf{x}^i \in \mathbb{R}^P$. For an arbitrary point $\mathbf{y}$, let $\mathcal{Y} = [\mathbf{y}_1, \ldots, \mathbf{y}_\kappa]$ denote the κ nearest neighbours within the embedding space, which correspond isometrically to the nearest neighbours $\mathcal{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_\kappa]$ on the high-dimensional space given the geometry preservation feature granted by the Isomap algorithm. A purely data-driven approach, based on the combination of the κ nearest neighbours, is proposed so that $f : \mathbb{R}^p \longrightarrow \mathbb{R}^P$ is defined as the unknown backmapping function. The high-order reconstruction $\mathbf{x}$ is derived as a first-order Taylor expansion starting from the nearest neighbour to be mapped back to the original space, $\mathbf{x} = \mathbf{x}_1 + (\mathbf{y} - \mathbf{y}_1)\nabla f(\mathbf{y}_1)^\top$ [14]. The gradient tensor $\nabla f(\mathbf{y}_1)$ is estimated assuming an orthogonal projection of the $\kappa - 1$ directions provided by the $\kappa-$Nearest Neighbors in $\mathbb{R}^P$ to those in $\mathbb{R}^p$, namely

$$\Delta\mathcal{X} = \begin{bmatrix} \mathbf{x}_2 - \mathbf{x}_1 \\ \ldots \\ \mathbf{x}_\kappa - \mathbf{x}_1 \end{bmatrix} \simeq \begin{bmatrix} \mathbf{y}_2 - \mathbf{y}_1 \\ \ldots \\ \mathbf{y}_\kappa - \mathbf{y}_1 \end{bmatrix} \nabla f\left(\mathbf{y}_1\right)^\top = \Delta\mathcal{Y}\nabla f\left(\mathbf{y}_1\right)^\top \tag{1}$$

Using the least squares minimisation, the gradient tensor is approximated as

$$\nabla f\left(\mathbf{y}_1\right)^\top = \left(\Delta\mathcal{Y}^\top\Delta\mathcal{Y}\right)^{-1}\Delta\mathcal{Y}^\top\Delta\mathcal{X}. \tag{2}$$

The performance of the backmapping function $f : \mathbb{R}^p \longrightarrow \mathbb{R}^P$ is affected by an appropriate determination of the number of κ neighbours used to compute the gradient tensor. This decision is based on the distribution of the data within the embedding and the number of Isomap variables under consideration. For this study, good performance of the backmapping is achieved for $\kappa = 5$.

2.3 Nonlinear regression with deep neural networks

The core of the surrogate model is the *regressor*, which serves the purpose of predicting data for conditions outside of the training set. In this study, DNNs are proposed to exploit their capability of universal approximators to establish the nonlinear relation between design parameters and pressure distributions.

The study aims to compare the performance of a direct-DNN regression model to an Isomap+DNN surrogate. Hence, two different DNNs are constructed to cope with both ap-

Aerodynamic and Geometrical features					Direct Interpolation		Isomap Interpolation	
M	α	$\Lambda - \Lambda_0$	t/t_0	b/b_0	R^2	MSE ($\times 10^2$)	R^2	MSE ($\times 10^2$)
Average on test dataset					0.974	1.236	0.949	2.329
0.70	4	0	0.9	0.9	0.926	2.622	0.974	0.936
0.74	8	-5	0.9	1.1	0.975	0.733	0.974	0.756
0.80	5	5	1.1	1.2	0.979	2.123	0.990	1.036
0.82	3	-5	1.1	1.1	0.961	2.164	0.972	1.530

Table 3: Performance metrics of the surrogate model: the correlation coefficient R^2 and the Mean Square Error MSE. Results are provided for the selected cases and the average of the testing set.

proaches. The chosen DNN architecture for both cases is a multilayer perceptron (MLP) made up of different linear layers with the ReLU activation function. The number of coefficients defining the MLP is kept similar for the two models to perform a fair comparison. In both cases, the input layer is defined by 5 neurons, one for each of the design features, including the flight condition and the geometric variations $(M, \alpha, \Delta\Lambda, t/t_0, b/b_0)$. Regarding the output layer, the direct-DNN model provides the vector for the pressure distribution with a size equal to the number of gridpoints, whereas the Isomap+DNN model predicts the low-dimensional representation, i.e. a vector containing (γ_1, γ_2). For the former, the MLP is made out of 3 hidden layers that progressively grow by a factor of 2^4 with respect to the previous, leading to an architecture of $[5, 32, 256, 2048, 6519]$. For the latter, the architecture features an expansion followed by compression with a similar rationale to the previous, yielding an architecture of $[5, 32, 256, 2048, 256, 32, 2]$.

Despite the apparent simplicity of the MLP architecture, it was tested to perform at least as well as for deeper configurations with more hidden layers and a less-abrupt increment in the number of neurons per layer. Additionally, a dropout operator with a 25% probability is included after the largest hidden layer in both cases to avoid overfitting. Regarding the loss function, the classical Mean Squared Error (MSE) is chosen. Finally, the optimiser is set to be the Adam Optimiser [28], a stochastic gradient descent method that is based on adaptive estimation of first-order and second-order moments. In each epoch, the loss is computed and a gradient descent optimisation is computed aiming at minimising it. Both models converged in less than 10^4 epochs without exhibiting a divergence between training and testing loss values.

3 PREDICTION OF C_p FOR NEW DESIGN CONDITIONS

The proposed surrogate models are evaluated in this section. First, the performance of the DNN for the Isomap+DNN model is evaluated, assessing its prediction in the low-dimensional embedding. Then, the direct-DNN and Isomap+DNN models are compared in terms of prediction accuracy of the pressure coefficient distribution for configurations outside of the training set. For the comparison, the results are depicted for the *visualisation* cases described in §2.1, which are summarised in table 3 together with the metrics for each case and the average values for the *test* set. For now on, those variables with a tilde symbol on top ($\tilde{\cdot}$) refer to predictions.

Prediction in the low-dimensional embedding

The performance of the MLP regressor used in the Isomap+DNN model is evaluated. The prediction of each point within the test dataset is depicted in figure 3. The MLP regressor keeps

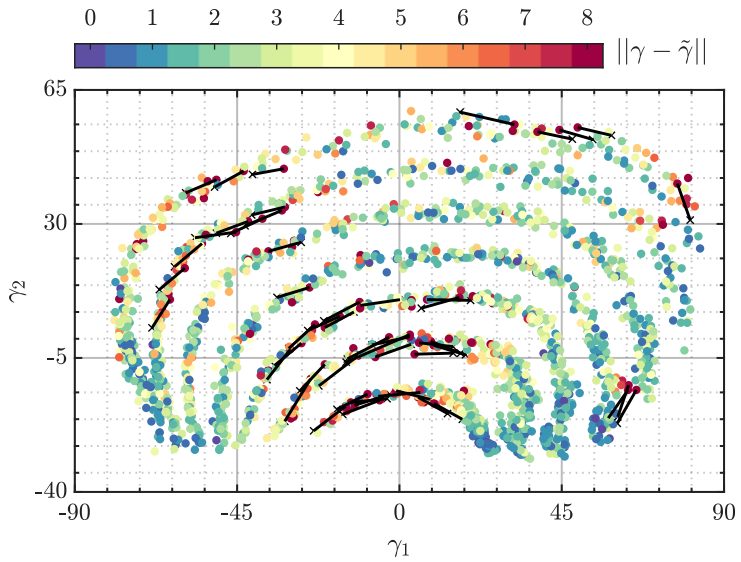

Figure 3: Prediction of the DNN in the low-dimensional embedding. The colour scale represents the displacement error of the prediction $||\gamma - \tilde{\gamma}||$, being $||\cdot||$ the Frobenius norm operator. The displacement from the real value (×) to the prediction by the DNN is shown for the 50 cases with the highest error.

the shape, distribution and size of the manifold. All mapped vectors fall inside the manifold surface. The colour scale refers to the distance between prediction $\tilde{\boldsymbol{\gamma}}$ and the real low-dimensional representation of the case under evaluation $\boldsymbol{\gamma}$, computed with the Frobenius norm, i.e. $||\tilde{\boldsymbol{\gamma}} - \boldsymbol{\gamma}||$. The displacement vector $(\tilde{\boldsymbol{\gamma}} - \boldsymbol{\gamma})$ is only shown for the 50 cases with the highest errors for the sake of simplicity.

The predicted test cases preserve in the embedding the stratified arrangement of data, in which each layer correspond to an angle of attack. Interestingly, only a few cases seem to be misallocated in the wrong layer, as is the case for low M and low α. Consequently, the cases with the greatest errors are mostly displaced along the constant-α layers, which implies that the error is prevalently related to the effect of M on the C_p. The wider picture of the error map within the embedding in figure 3 confirms that the MLP regressor performs reasonably well for all the regions within the manifold. It has to be noted that the displacement of the prediction is less than 5% of the manifold span for the worst cases, which can be considered an acceptable prediction error.

The comparison of direct-DNN and Isomap+DNN surrogates in estimating the pressure coefficient distributions is illustrated for the *visualisation* cases in figure 4 and figure 5. These results evaluate the performance of the MLP regressor for the direct-DNN alternative, and the performance of the combined MLP regressor and kNN backmapping for the Isomap+DNN approach. In figure 4, the data is presented in the form of chordwise C_p distributions for three arbitrary locations along the wing span, $\eta = [0.1, 0.5, 0.9]$ being η the spanwise coordinate normalised with the wing span b. On the other hand, the prediction error $\Delta C_p = C_p - \tilde{C}_p$ is depicted in figure 5 for the upper wing surface together with the regression plot.

Focusing on *Case 1*, it is already possible to assess the main weaknesses and strengths of the proposed methods. This case corresponds to a flight condition with a low Mach number and medium-high angle of attack $(M, \alpha) = (0.7, 4°)$. In such conditions, there is no shock wave on the upper surface of the wing, being the suction peak at the leading edge the main

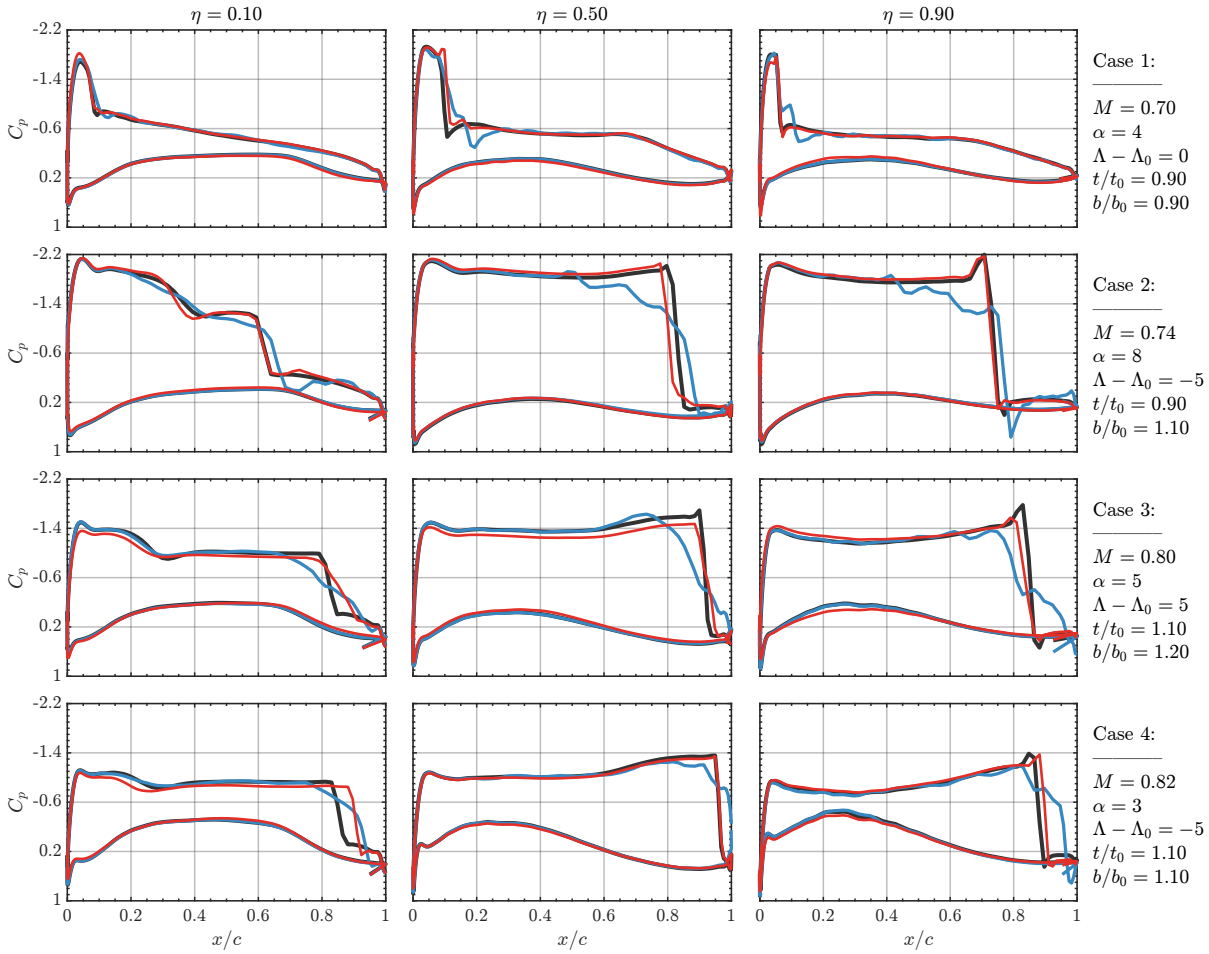

Figure 4: C_p chordwise profiles for CFD simulation(——), direct DNN regression(——) and Isomap+DNN regression (——). Profiles are shown for the visualisation cases at three locations along the wing span.

challenge for the surrogate model as shown in figure 5. Indeed, both models perform reasonably well for $\eta = 0.1$, in which the peak is mild at the root of the wing. However, the suction peak becomes more abrupt in the mid-span section and close to the tip, becoming hard to predict for the direct-DNN approach (see figure 4). Despite the nonlinear nature of the DNN, the direct regression jeopardises the prediction performance of the model in the presence of drastic changes of C_p. On the other hand, the Isomap+DNN model outstands in this matter. The former approach filters out the information within the C_p distribution since the prediction $\tilde{C}_p$ is derived from scratch by the DNN. Conversely, regardless of the regression performance in the low-dimensional embedding, the prediction by the Isomap+DNN model is achieved after the combination of existing data samples through the backmapping function, which gives more guarantees that the physical behaviour is preserved when interpolating.

The aforementioned attributes for each model are also present for *Cases 2 - 4*, in which the suction peak becomes irrelevant compared to the presence of a shock wave close to the trailing edge as illustrated in figure 5. We observe that the overall performance of Isomap+DNN in identifying the shock wave, its positions, and following the C_p through it is remarkably better than for the direct-DNN approach based on the C_p profiles in figure 4.

Despite the good performance of Isomap+DNN on the identification of abrupt alterations in C_p, its overall regression capabilities are slightly lower than for the direct-DNN approach, as extracted from the global metrics in table 3. The proposed Isomap+DNN model is a good

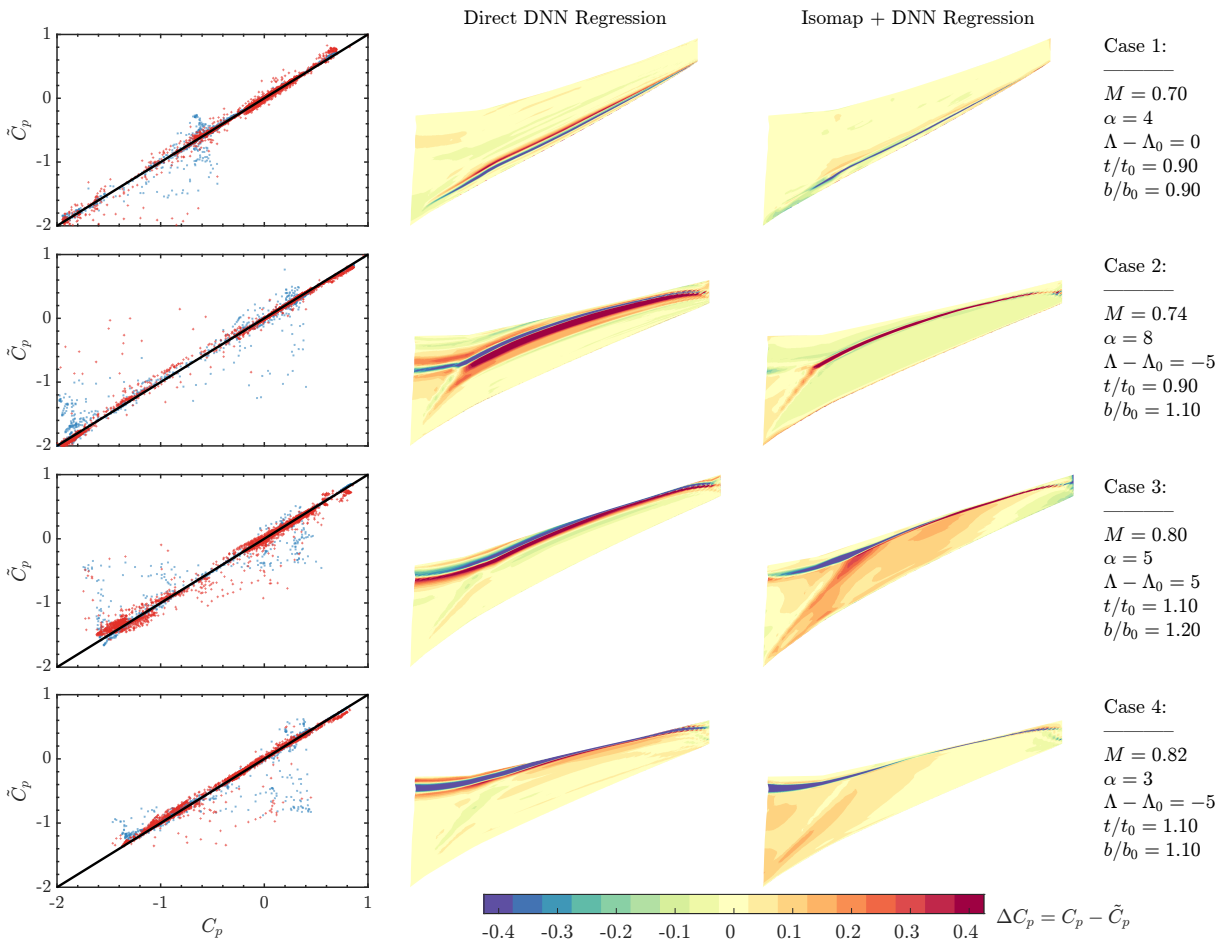

Figure 5: Comparison of direct-DNN and Isomap+DNN regression models for the selected cases. (Left column) regression plot for direct DNN (×) and Isomap+DNN (+) regression. Deviation of the prediction with respect to the simulation data for direct DNN (mid column) and Isomap+DNN (right column) regression.

alternative to retain complex physics within the data, however, the simplicity of the regression model in the low-dimensional embedding prevents the model to perform as desired when the C_p distribution is smooth. For all the visualisation cases, direct-DNN prediction better attach to the CFD data everywhere else but the shock wave region. As an example, the metrics for *Case 2* in table 3 conclude a slightly better performance of the direct-DNN approach even though there are strong deviations in the vicinity of the shock wave as demonstrated by the C_p distribution over the upper surface in figure 5. Similarly, *Case 3* and *Case 4* appear to provide better metrics for the Isomap+DNN although the prediction far from the shock wave is underestimated in both cases as shown in figure 5. Consequently, and considering that many cases in the test set do not exhibit such abrupt phenomena, the overall performance of the direct-DNN tends to be superior, while Isomap+DNN is generally better in the detection and interpolation of highly nonlinear phenomena.

As a final note, an assessment of the computational demands of each model is necessary. Isomap, which serves as a dimensionality reduction algorithm, effectively accomplishes its objective of simplifying the problem from a data management perspective. This is evidenced by the size of the Isomap+DNN model relative to the direct interpolation models described in the previous section. Isomap decreases the size of the output data array from the regression model from 6519 elements composing the C_p distribution to a mere $p = 2$ elements, resulting

in a reduction factor of roughly 3000. This reduction has a discernible effect on the size of the regression models as well as on the computational costs of training and prediction. These results are highly desirable for surrogate-based optimisation, where the surrogate model is incrementally refined as new information becomes available.

4 CONCLUSIONS

The development of surrogate models in the low-data limit for the prediction of pressure-coefficient distribution on transonic wings is investigated. The proposed surrogate combines a reduction-dimensionality and a nonlinear regressor to develop a robust and affordable yet performing predictive tool for its implementation in the preliminary design phase of multidisciplinary design and optimisation processes. The test case is the wing of the XRF1 model simulated for different flight conditions and geometries. The resultant database contains C_p data from CFD simulations based on the Mach number, angle of attack, sweep angle, thickness and span as design parameters. We propose a surrogate model referred to as Isomap+DNN combining a manifold learner based on Isomap to reduce the dimensionality of the C_p data with a multilayer-perceptron regressor that maps the design parameters to the manifold space. Additionally, we use a backmapping based on k-nearest neighbours to reconstruct the C_p from the low dimensional data.

The Isomap embedding exhibits a strong relation with M and α. The surface of the manifold is stratified with each layer corresponding to a specific α. The effect of the geometric variations is ingrained in the embedding as a modification of the local flight condition. The MLP regressor performs reasonably well when predicting the low-dimensional representation of C_p in the manifold space. The displacement error of the prediction does not show a clear trend based on the design parameters, with less than 5% error compared to the manifold size.

The prediction of C_p is very robust for both direct-DNN and Isomap+DNN surrogates, with slight superiority of the direct approach from a global perspective. However, the Isomap+DNN model is superior in identifying compressibility effects, such as shock waves or strong suction peaks. Despite the nonlinear nature of the DNN, the performance of direct-DNN is compromised in the presence of significant changes in the C_p distribution. This approach filters out the information contained within the C_p distribution since the prediction is obtained solely from the DNN based on the input parameters. In contrast, the Isomap+DNN model's prediction is achieved by combining existing data samples through the backmapping function, providing greater assurance that the physical nature is maintained during interpolation. Despite the regression performance in the low-dimensional embedding, the Isomap+DNN model is more reliable than the direct-DNN approach.

The Isomap+DNN model has demonstrated superior performance in predicting configurations with stronger compressibility effects, thereby creating an opportunity to develop robust and cost-effective surrogate models for predicting aerodynamic data on transonic wings. This approach significantly outperforms the direct-DNN method in identifying the shock wave and its positions and tracking the C_p through it. It is envisioned that combining both techniques could produce a synergistic surrogate model with strengths in both predicting the overall C_p distribution and identifying abrupt phenomena resulting from compressibility effects.

ACKNOWLEDGEMENTS

The authors would like to thank Airbus for providing the database for the XRF1 test case.

CODE AVAILABILITY

The Python library used in this research for the Isomap+DNN regression model is publicly published: `https://github.com/TACOMA-INTA/tacoma-lib`.

REFERENCES

[1] S. L. Brunton, B. R. Noack, and P. Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52(1):477–508, 2020.

[2] A. I.J. Forrester and A. J. Keane. Recent advances in surrogate-based optimization. *Progress in Aerospace Sciences*, 45(1):50–79, 2009.

[3] J. Li, X. Du, and J. R.R.A. Martins. Machine learning in aerodynamic shape optimization. *Progress in Aerospace Sciences*, 134:100849, 2022.

[4] J. Li and M. Zhang. Adjoint-free aerodynamic shape optimization of the common research model wing. *AIAA Journal*, 59(6):1990–2000, 2021.

[5] L. Sirovich. Turbulence and the dynamics of coherent structures. i, ii, iii. *Quarterly of applied mathematics*, 45(3):561–590, 1987.

[6] Jonathon Shlens. A tutorial on principal component analysis, 2014.

[7] H. V. Ly and H. T. Tran. Modeling and control of physical processes using proper orthogonal decomposition. *Mathematical and Computer Modelling*, 33(1):223–236, 2001. Computation and control VI proceedings of the sixth Bozeman conference.

[8] J. Franklin. The elements of statistical learning: data mining, inference and prediction. *The Mathematical Intelligencer*, 27(2):83–85, 2005.

[9] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.

[10] W. S. Torgerson. Multidimensional scaling: I. theory and method. *Psychometrika*, 17(4):401–419, 1952.

[11] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[12] F. Tauro, S. Grimaldi, and M. Porfiri. Unraveling flow patterns through nonlinear manifold learning. *PLOS One*, 9(3):1–6, 03 2014.

[13] G. Bansal, A. A. Mascarenhas, and J. H. Chen. Identification of intrinsic low dimensional manifolds in turbulent combustion using an isomap based technique. Technical report, Sandia National Lab.(SNL-CA), Livermore, CA (United States), 2011.

[14] E. Farzamnik, A. Ianiro, S. Discetti, N. Deng, K. Oberleithner, B.R. Noack, and V. Guerrero. From snapshots to manifolds – a tale of shear flows. *Journal of Fluid Mechanics*, 955:A34, 2023.

[15] T. Franz, R. Zimmermann, S. Görtz, and N. Karcher. Interpolation-based reduced-order modelling for steady transonic flows via manifold learning. *International Journal of Computational Fluid Dynamics*, 28(3-4):106–121, 2014.

[16] K. Decker, H. D. Schwartz, and D. Mavris. *Dimensionality Reduction Techniques Applied to the Design of Hypersonic Aerial Systems*.

[17] R. Castellanos, J. Bowen-Varela, A. Gorgues, and E. Andrés. An assessment of reduced-order and machine-leaning models for steady transonic flow prediction on wings. In *33th Congress of the International Council of the Aeronautical Sciences*, pages 03.1 – Aerodynamics – CFD Methods and Validation. ICAS, 2022.

[18] A.K. Jain, J. Mao, and K.M. Mohiuddin. Artificial neural networks: a tutorial. *Computer*, 29(3):31–44, 1996.

[19] X. Du, P. He, and J. R. R. A. Martins. Rapid airfoil design optimization via neural networks-based parameterization and surrogate modeling. *Aerospace Science and Technology*, 113:106701, 2021.

[20] B. Glaz, L. Liu, and P. P. Friedmann. Reduced-Order Nonlinear Unsteady Aerodynamic Modeling Using a Surrogate-Based Recurrence Framework. *AIAA Journal*, 48(10):2418–2429, 2010.

[21] J. Li and M. Zhang. Data-based approach for wing shape design optimization. *Aerospace Science and Technology*, 112:106639, 2021.

[22] E. Andrés-Pérez and C. Paulete-Periáñez. On the application of surrogate regression models for aerodynamic coefficient prediction. *Complex & Intelligent Systems*, 7(4):1991–2021, 2021.

[23] V.E. Kovalev and O.V. Karas. Calcul de l'écoulement transsonique autour d'une configuration aile-plus-fuselage compte tenu des effects visqueux et d'une région décollée mince. *La Recherche. Aérospatiale*, (1):23–38, 1994.

[24] K Zhang and M. Hepperle. Evaluation of the BLWF Code - A Tool for the Aerodynamic Analysis of Transonic Transport Aircraft COnfigurations. Technical report, Institut für AeroDynamik und Strömungstechnik, Juli 2010.

[25] R. W. Floyd. Algorithm 97: shortest path. *Communications of the ACM*, 5(6):345, 1962.

[26] O. Samko, A.D. Marshall, and P.L. Rosin. Selection of the optimal parameter value for the isomap algorithm. *Pattern Recognition Letters*, 27(9):968–979, 2006.

[27] L. K. Saul and S. T Roweis. Think globally, fit locally: unsupervised learning of low dimensional manifolds. *Journal of machine learning research*, 4(Jun):119–155, 2003.

[28] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2014.