Expert Systems With Applications 255 (2024) 124726

Contents lists available at ScienceDirect

Expert Systems With Applications

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

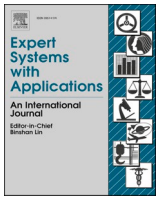

Parsimonious airfoil Parameterisation: A deep learning framework with Bidirectional LSTM and Gaussian Mixture models

Vincent le Roux [a,*], Marelie H. Davel [a,b,c], Johan Bosman [a]

[a] Faculty of Engineering, North-West University, South Africa
[b] Centre for Artificial Intelligence Research (CAIR), South Africa
[c] National Institute for Theoretical and Computational Sciences (NITheCS), South Africa

ARTICLE INFO

Keywords:
Airfoil
Parameterisation
Optimisation
Deep learning

ABSTRACT

The choice of airfoil parameterisation method significantly influences the overall wing optimisation performance by affecting the flexibility and computational efficiency of the process. Ideally, one should be able to intuitively constrain airfoil shape and structural characteristics as input to the optimisation process. Current parameterisation techniques lack the flexibility to generate airfoils efficiently by specifying parsimonious shape and structural features. To address this limitation, a deep learning framework is proposed, enabling conditional airfoil generation from an airfoil's shape and structural feature definition.

Specifically, we demonstrate the application of Bidirectional Long Short Term Memory models and Bayesian Gaussian Mixture models to derive airfoil coordinates from a compact set of shape and structural characteristics that we define. The proposed framework is shown to achieve favorable airfoil performance optimisation due to improved exploration and exploitation of the design space, compared to traditional approaches. Overall, the proposed optimisation framework is able to realise a 9.04% performance improvement over an airfoil design optimised with traditional parameterisation techniques.

1. Introduction

In a typical wing optimisation process, the objective is to maximise some performance metric, such as the lift-to-drag ratio, over the entire flight envelope. Here, optimisation typically involves adjusting the wing geometry (station chord lengths and wingspan increments) and the airfoil profiles present at each wing station. Multiple stations are located along the span of a wing, potentially with distinct airfoils, each represented by 100 to 500 coordinate points Mukesh et al., 2014). Due to this complexity, optimising in the high-dimensional airfoil coordinate space (specifically x and y) becomes computationally expensive.

Optimising in these high dimensions is also unintuitive and limits the designer in the constraints that can be set on the shape characteristics of each airfoil (based on prior experience or external objectives).

To mitigate this, aerodynamicists have resorted to the act of airfoil parameterisation (Lu et al., 2018). Airfoil parameterisation can be described as the process of representing the high-dimensional airfoil shape space in a lowerdimensional parametric space. Parameterising the individual airfoils in the 2D space enables more efficient and intuitive optimisation in the integrated 3D space. However, the choice of parameterisation method used to generate the airfoil shapes used at the stations along the wingspan has a significant impact on the flexibility of this optimisation process (Samareh, 2001).

The two categories of airfoil parameterisation methods are deformative and constructive methods (Lu et al., 2018; Masters et al., 2017). Deformative methods take an existing airfoil which is then deformed to create a new shape. Constructive methods represent an airfoil shape that is purely based on a series of parameters specified. Whether deformative or constructive, Padulo et al. (2009) define the desirable properties of an airfoil parameterisation method as: (1) Completeness, (2) Orthogonality, (3) Flawlessness, (4) Parsimony, and (5) Intuitiveness. Having an airfoil parameterisation technique that adheres to all five of these criteria means efficient, accurate, and intuitive airfoil generation will result, from a limited set of parameters, whilst ensuring sufficient local and global control during the shape generation process.

Several methods exist for airfoil shape parameterisation, including B-splines (Piegl & Tiller, 1995), Hicks-Henne (Hicks & Henne, 1978), Class-Shape Transformation (CST) (Kulfan, 2008), Singular Value Decomposition (SVD) (Toal et al., 2010), and PARSEC-11 (Sobieczky, 1999). Although widely used, each of the mentioned methods fails to

* Corresponding author at: Faculty of Engineering, North-West University, Kikuyu Balwin Properties, Unit 262, Gauteng, Midrand, 2090, South Africa.
E-mail addresses: 25023438@student.g.nwu.ac.za (V. le Roux), marelie.davel@nwu.ac.za (M.H. Davel), johan.bosman@nwu.ac.za (J. Bosman).

https://doi.org/10.1016/j.eswa.2024.124726
Received 12 December 2023; Received in revised form 23 June 2024; Accepted 4 July 2024
Available online 10 July 2024

comply with all five criteria of Padulo et al. (2009). Specifically, Hicks-Henne functions, B-splines, CST, and SVD lack parsimony, intuitiveness, and flawlessness. Although PARSEC-11 has intuitiveness and parsimony implicit in the parameter definitions, it lacks the capability to represent any airfoil (completeness) and is also prone to generating impracticalor non-airfoil shapes (flawlessness). To address the completeness issue, Nemati, and Jahangirian extended the PARSEC-11 method by adding the capability of explicit leading and trailing edge control (Nemati & Jahangirian, 2020). This meant that the leading edge could be described with six parameters and the trailing edge with 11 parameters. Although their solution meant more flexibility in shape representation, it did not account for additional intuitiveness or flawlessness. To enhance the intuitiveness of the PARSEC-11 method, Christie et al. (2019) combined the method with the CST parameterisation method to create the intuitive CST (iCST) method. The iCST methodology involves transforming the conventional CST parameterisation technique into a comprehensive set of intuitive parameters through the application of a transformation matrix. Although the iCST enables basic intuitive constraints on the airfoil generation process it is complex to implement and a change in any given curve parameter prompts a global, rather than local, change to the airfoil shape – which is not desirable during airfoil performance optimisation.

With the advent of big data and deep learning, more sophisticated methods were developed in an attempt to create a Padulo-complete airfoil generation method, that is, an airfoil generation method that strictly adheres to all five criteria set out by Padulo et al. (2009). In this endeavor, various deep learning methods, such as Generative Adversarial Networks (GANs) Goodfellow et al., 2014) and Variational Auto Encoders (VAEs) (Kingma & Welling, 2019), have been used in approaches to airfoil inverse design and airfoil shape generation.

Recently, Kou et al. (2023) used VAEs to enhance airfoil parameterisation whilst using fewer optimisation variables. With their proposed parameterisation method they were able to achieve a 3% noise reduction when comparing optimal airfoil geometry trailing edge noise. The work of Wang et al. (2022) showed that generative models such as Wasserstein GANs in combination with Conditional VAEs can be used to generate target wall Mach distributions for the inverse design that matches specified features, such as the location of suction peaks, shocks and aft loading. The work of Achour et al. (2020) saw the development of a Conditional GAN capable of predicting an airfoil profile given the maxi-mum lift-to-drag ratio and total airfoil area requirements. Efforts by Sekar et al. (2019) showed that deep Convolutional Neural Networks could also be utilised to obtain the airfoil shape from the desired pressure distributions. Recently, Wada et al. (2024) showed that utilising physics-guided training during GAN development yields a model that generalised well in the prediction of a smooth airfoil profile – for a given set of desired performance characteristics. Also focusing on GANs applied to airfoil generation, Chen & Fuge (2018) introduced the BezierGAN architecture. This architecture enabled smooth synthesis of aerodynamic shapes from a set of uninterpretable low-dimensional latent vectors.

Another way to parameterise airfoils is by means of design morphing. This methodology entails gradually transforming shapes that are homeomorphic, meaning they have the same underlying topological structure, into one another. This process creates a continuous and unrestricted design space for exploration. A Design-by-Morphing (DbM) framework was developed in Sheikh et al. (2023), where it was shown how this framework can reconstruct the UIUC airfoil data base with a greater than 99% accuracy and how DbM be applied to find novel airfoil designs during airfoil performance optimisation.

Although each of the mentioned studies were successful in their individual objectives, limitations of the current approaches to airfoil parameterisation include: (1) lack of explicit global and local control over the airfoil profile characteristics and (2) only partial adherence to Padulo's criteria, i.e. no current method is Padulo complete. The main goal of this work is to address these limitations. To this end, a novel conditional airfoil generation framework is developed and a Bidirectional Long Short Term Memory (BLSTM) architecture in combination with a Bayesian Gaussian Mixture Model (BGMM) architecture is used to enable efficient mapping from the low-dimensional parametric space to the high-dimensional coordinate space. The main contribution of this research is therefore the development of an airfoil parameterisation framework that adheres to as many of Padulo's criteria as possible while enabling the generation of airfoil profiles with explicit control over the global and local profile characteristics.

The remainder of this article can be summarised as follow. In Section 2 the descriptive and intuitive airfoil parameters used in this research are introduced. In Section 3 a thorough overview of the data mining process and the BGMM and BLSTM are given. Section 4 sees the development of the proposed airfoil generation framework. In Section 5 the proposed airfoil generation framework is applied to constrained airfoil performance optimisation and the results are discussed in Section 6. Finally, in Section 7, a concise discussion is given on the findings of this research including limitations of the proposed framework and potential areas of further research to extend the utility, efficiency, and robustness of the proposed framework.

2. Airfoil parameterisation

In this section, we describe the parsimonious shape and structural variables selected for use in the proposed framework. Currently, there is no available consensus on what the most important parsimonious airfoil characteristics are when designing airfoils. The current, widely adopted airfoil parametrisation technique is the PARSEC (Sobieczky, 1999) method. Although this method gives 11 parsimonious airfoil shape features, it lacks providing information about:

1. The camber distribution.
2. Structural characteristics.
3. Maximum camber location and magnitude.
4. Magnitude and location of the maximum airfoil thickness.
5. Curvature and location of the maximum lower surface.

Masters et al. (2017) determined that at least 25 design variables are needed to cover 80% of the design space regardless of the deformative or constructive airfoil parametrisation method used. This is significantly more than the 11 parameters of PARSEC or the number of control points typically used when optimising 2D airfoils with B-Spline or CST methods. In this research, a total of 23 parsimonious variables are chosen. The number of parameters are intentionally chosen to be close to the 25 parameter limit set by Masters in orderto maximise the coverage of the design space of the proposed airfoil generation technique.

In order to maximise airfoil coverage by the selected 23 parsimonious variables, the widely used PARSEC shape parameters, the mentioned features not covered by this method, and additional shape and structural features are numerically extracted from raw airfoil x and y coordinates. The airfoil parameterisation framework is designed such that it can utilise the x and y coordinates produced with any constructive or deformative method.

Fig. 1 shows the airfoil parsimonious shape parameters used in this research to describe an airfoil's profile and in Table 1. the detail of these 17 shape variables can be seen. As can be seen from this figure, all the shape features are chosen to be intuitive, thereby ensuring intuitive airfoil generation with the downstream developed airfoil generation model.

Apart from the 17 parsimonious shape variables mentioned, the profile of a given airfoil is also mapped to its corresponding structural characteristics. The following six structural properties are used: (1) airfoil area, (2) airfoil perimeter, (3) airfoil x centroid, (4) airfoil y centroid, (5) bending inertia about the y-axis, and (6) bending inertia about the x-axis. These additional structural variables are included, as this allows for the extension of the final system's flexibility in terms of

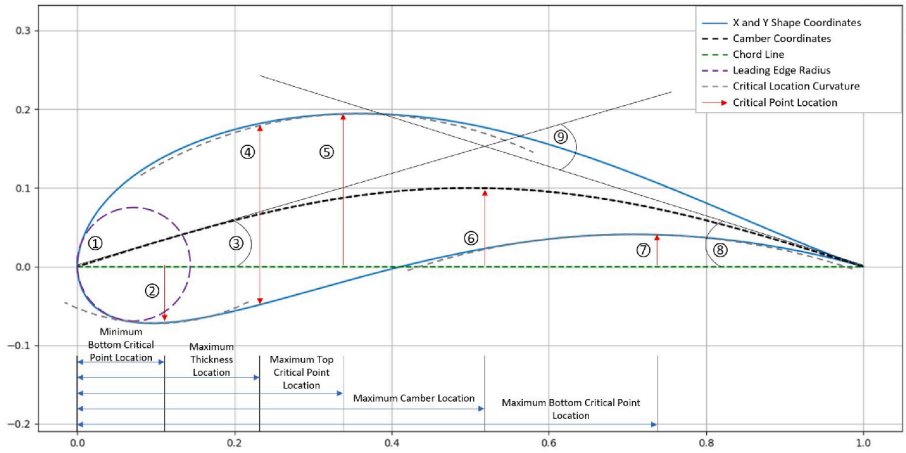

Fig. 1. Depiction of the chosen airfoil parsimonious shape variables used to describe a diverse range of airfoil profiles.

Table 1
The selected 17 airfoil parsimonious shape parameters, as depicted in Fig. 1.

Item	Number of features	Feature description
1	1	Nose radius of the leading edge
2	3	Magnitude of the minimum lower surface peak Location of the minimum lower surface peak Minimum lower surface peak curvature
3	1	Leading edge camber angle
4	2	Magnitude of the maximum airfoil thickness Location of the maximum airfoil thickness
5	3	Magnitude of the upper surface peak Location of the upper surface peak Upper surface peak curvature
6	2	Magnitude of the maximum airfoil camber Location of the maximum airfoil camber
7	3	Magnitude of the maximum lower surface peak Location of the maximum lower surface peak Maximum lower surface peak curvature
8	1	Trailing edge camber angle
9	1	Camber angle
Total		17

intuitive and explicit constraint setting. For example, if an airfoil with a low area is required, then all airfoils generated that do not adhere to the defined area constraint can be excluded from the optimisation procedure upfront.

3. Experimental setup

To develop a model that is capable of generating airfoils in terms of their 23 parsimonious variables (as defined in Section 2) a rich airfoil corpus is required. The development of such a corpus is described in Section 3.1. This is followed by an overview of the model development process utilising this corpus (Section 3.2).

3.1. Data mining

In the airfoil corpus, the lower-order dimensional representation of an airfoil's 23 parsimonious variables needs to be mapped to the corresponding higherorder profile definition in terms of x and y coordinates. Seeing that such a dataset is not publicly available at the time of conducting this research, a custom data mining pipeline is developed to mine these parsimonious variables and the corresponding airfoil profile x and y coordinates. Such a data mining pipeline requires the following four core components:

1. An airfoil generation technique: responsible for generating the airfoil profile and hence the x and y coordinate definition of an airfoil.
2. A sampling procedure: responsible for efficiently exploiting the input design space of the airfoil generation technique to generate diverse airfoil profiles that maximises coverage.
3. An airfoil parameterisation framework: responsible for calculating an airfoil's 23 parsimonious variables from its x and y coordinate definition.
4. A filtering procedure: responsible for filtering out any airfoil profiles that either represent non-airfoil shapes or do not fall within the desired profile distribution.

For the sake of brevity, the details of the data mining pipeline are not discussed here, however, a high-level overview is given. For the airfoil generation technique, the Simple Parametric Model (SPM) of (Ziemkiewicz, 2017) is used. This constructive method is chosen due to its simplicity in constructive airfoil generation and the diverse airfoil population it can describe – the data miningpipeline can, however, make use of any constructive or destructive airfoil generation technique to extend the airfoil diversity resulting from the data mining pipeline.

To achieve the goal of maximising the coverage of the input design space of the SPM while drawing as few samples as possible, Latin Hypercube Sampling (LHS) is used (Li et al., 2016). LHS is a stratified sampling method that divides each dimension of the input space into equal intervals and ensures that each interval has exactly one sample point. LHS can produce more uniform and representative samples than simple random sampling, and it can also reduce the number of samples needed to achieve a certain level of accuracy.

For airfoil filtering, a two-step process is taken. First, realistic airfoil shapes are ensured by means of filtering out airfoils with intersecting geometry, highly curved surfaces, and/or airfoils with multiple first derivative changes – indicating saw-tooth airfoil surfaces. Secondly, airfoils that do not have the desired parsimonious variable characteristics are removed from the database. This includes filtering out airfoils with too large of an area, the camber location sitting too close to the trailing edge, a high degree of camber, etc. (these requirements will depend on the specific goal and can be specified by the user). Once an airfoil is generated with this method, the 23 parsimonious variables are numerically calculated by making use of its x and y coordinate profile definition.

To develop the required airfoil corpus, the data mining pipeline is executed to mine parsimonious variable and airfoil profile data for a total of 18,614 airfoils. This dataset is referred to as the airfoil shape dataset and the schema of this dataset is as depicted in Table 2. It is important to note that, although the SPM's six parameters are logged during the data mining process, and hence reported in the schema in Table 2. these parameters are discarded during the model development as they do not add any additional information or parsimony to the

Table 2
Data schema for the airfoil shape dataset.

Variable	Dimensionality
Parsimonious variables	(23) 17 shape variable 6 structural variables
SPM coefficient	(6) B, T, P, E, C, and R coefficients
x and y coordinates	(160) 80 x and 80 y coordinates

overall shape data set.

Furthermore, this airfoil shape dataset is manipulated and reshaped as required by the BGMM and the BLSTM architectures. Specifically, for the BGMM, the model inputs and the model responses are set to be the same,

seeing that this is an unsupervised model and the task is to effectively separate the distance between the identified airfoil clusters. For the BLSTM model development, in order to ensure uniform sequential structure for each sample, the output data is ordered as follows: (1) the 80x and 80y coordinates are divided into 40 top and 40 bottom coordinates respectively; (2) the 40 top x and y coordinates are arranged from trailing edge to leading edge; and (3) the 40 bottom x and y coordinates are arranged from leading edge to trailing edge. In this research, 80x and 80y coordinates are used to describe the airfoil profiles mined with the SPM as this gives sufficient accuracy in capturing variation in fine local structure. However, this number can be increased or decreased based on the user's requirements.

For the development of both models, a training set consisting of 70% of the 18,614 airfoil samples, a testing set of 20%, and a validation set of 10% of the samples is generated by means of sampling without replacement.

3.2. Modelling overview

As introduced earlier, two architectures are used during airfoil generation, namely the BGMM and the BLSTM architectures. The architectures themselves and the rationale for their selection are described here.

3.2.1. BGMM

The BGMM is used to model the distribution of airfoils in the parsimonious variable space. The BGMM architecture is specifically chosen due to its computational efficiency for large data sets and its ability to accurately model the co-variance of one variable in relation to another. This characteristic means that a realistic parsimonious variable representation of an airfoil profile can be sampled from the trained BGMM parameters, as it captures the conditional variance, and hence the local sequential structure, between the shape and structural parameters. BGMMs also have the added advantage of conditional data sampling, meaning that the desired property of conditionally generating airfoil profiles with selected properties is satisfied.

The probability density function of a multivariate Gaussian distribution with K components, can be expressed as:

$$N\left(\vec{x}\middle|\vec{\mu}_i,\Sigma_i\right)=\frac{1}{\sqrt{(2\pi)^2|\Sigma_i|}}exp\left(-\frac{1}{2}\left(\vec{x}-\vec{\mu}_i\right)^T\Sigma_i^{-1}\left(\vec{x}-\vec{\mu}_i\right)\right) \quad (1)$$

With $\vec{\mu}_i$ and Σ_i the mean and covariance matrix for the i^{th} component. Here, $\vec{\mu}_i$ is a K dimensional vector and Σ_i is a $K\times K$ dimensional vector. If the i^{th} mixing component, or mixture weight, is expressed as ϕ_i with the constraint of $\sum_{i=1}^{K}\phi_i=1$ such that the total probability distribution normalises to 1, the GMM probability density is defined as a linear function of densities of all these K distributions:

$$p(\vec{x})=\sum_{i=1}^{K}\phi_i N\left(\vec{x}\middle|\vec{\mu}_i,\Sigma_i\right) \quad (2)$$

Using The Expectation Maximisation (EM) algorithm, as defined in McLachlan et al. (2004), the parameters of Equation (2) can be found for the K components – with the number of components K being a hyperparameter optimised for during model training. In this study, the BGMM as discussed in Lu (2021) is employed to ensure the optimal number of clusters are found. With this method, regardless of the number of components selected, the EM algorithm is embedded into a variational Bayesian estimation framework that allows inference of the effective number of components from the data.

3.2.2. BLSTM

The second model required should be able to generate the sequence of x and y coordinate pairs, describing the airfoil profile, such that there is a one-to-one mapping between the parsimonious variable domain and the x and y coordinate domain. There exist various sequence-to-sequence models capable of handling this problem, including RNNs (Sherstinsky, 2020) and LSTMs (Hochreiter & Schmidhuber, 1997). Both these models are capable of modelling data with different sequence lengths, i.e. map data from a lower dimension to a higher dimension and vice versa. It is, however, well documented that RNNs are subject to the vanishing gradient problem (Schmidt, 2019). This means RNNs perform poorly when processing long sequence lengths. Given the x and y coordinate dimensionality required to accurately describe an airfoil profile, RNNs are therefore not suitable for use in the proposed framework (referred to hereon out as the proposed framework).

LSTMs aim to solve the problem of the vanishing gradient by means of employing various memory units. In this memory unit, each unit contains three gates: (1) an input gate i_t, (2) a forget gate f_t, and (3) an output gate o_t. The operation of the network follows the following equations:

$$i_t=\text{sigm}(W_i x_t+U_i h_{t-1}+b_i) \quad (3)$$

$$f_t=\text{sigm}\left(W_f x_t+U_f h_{t-1}+b_f\right) \quad (4)$$

$$o_t=\text{sigm}(W_o x_t+U_o h_{t-1}+b_o) \quad (5)$$

$$\widetilde{c}=\tanh(W_c x_t+U_c h_{t-1}+b_c) \quad (6)$$

$$c_t=i_t\odot\widetilde{c}_t+f_t\odot c_{t-1} \quad (7)$$

$$h_t=o_t\odot\tanh(c_t) \quad (8)$$

With h_t the hidden state vector i.e the LSTM output, c_t the self-connected cell state vector, $\widetilde{c}$ the cell input activation vector, $\odot$ representing element-wise vector operation, b^* the respective gates bias term, W^* the input-to-hidden weights matrix, and U^* the state-to-state weights matrix.

Although LSTMs propose a solution to the vanishing gradient and longterm-dependence problems encountered in the traditional RNN architectures, the vanilla LSTM lacks the capability to use past and future context to perform inference. This is a desirable property for the proposed framework seeing that the mapping from the parsimonious variable domain to the x and y coordinate domain will be more accurate when the model is capable of reading the parsimonious vector sequence in both directions. This multi-directional processing capability is innate to the BLSTM architecture (Hamdi et al., 2022). Essentially, this configuration involves two LSTM networks: one functioning on the input sequence in its natural order, and the other operating on the sequence in reverse order. This means the BLSTM is capable of using both past and future context while at time of inference.

Given the data requirements of this architecture, the application of these models is limited to cases where the full sequence is available at

time of inference. This adds latency to the prediction process if the data is generated sequentially and therefore BLSTMs are not suitable for tasks such as real-time online speech recognition. Also, because the BLSTM has two unidirectional LSTM networks, they are more computationally expensive to train. However, given the modest size of the data set and the fact that the entire parsimonious variable vector is available at time of inference, the use of the BLSTM architecture in the proposed framework is appropriate.

4. Model development and validation

The development and validation of the proposed framework is discussed next. First, in Section 4.1, the architectures and respective training protocols of the BGMM and the BLSTM is discussed. This is followed by a case study showcasing the conditional airfoil generation capacity of the proposed framework in Section 4.2. Next, in Section 4.3, the robustness of the proposed framework is demonstrated in the reconstruction of an out-of-sample airfoil. A comparison is then drawn in Section 4.4 between the efficiency of the proposed frameworkand that of a traditional airfoil generation technique in targeted airfoil generation. Finally the extent to which the proposed framework adheres to Padulo's completeness criteria is discussed in Section 4.5.

4.1. Model development and training protocol

The inference workflow of the proposed framework can be seen in Fig. 2. As depicted in this figure, the inference process starts by sampling from the multi-variate Gaussian distributions fitted to the data during the training process. In this sampling process, the fitted Gaussian parameters and co-variance matrices are indexed for the airfoil cluster of interest. This is akin to conditioning the sampling of the airfoil to a particular airfoil profile. From this sampling, a 23 dimensional vector results, representing the 23 parsimonious variables of the conditioned airfoil. This 23 dimensional vector is then transformed appropriately for use in the BLSTM where the x and y coordinates of the airfoil profile are predicted.

A favourable characteristic of this architectural design is the controllability of the shape and structural features of the generated airfoil at two different model entry points: (1) on the conditioned airfoil sampled by using the desired fitted Gaussian parameters and co-variance matrices, and (2) the 23 parsimonious variables used as input to the BLSTM model. That is, when the 23parsimonious variables are not within the desired upper or lower bounds, another sample or set of samples can iteratively be drawn until desirable shape and structural characteristics result. The advantage here, above the traditional method, is robustness and efficiency seeing that the airfoil profile generation may only commence once desirable characteristics are encountered. This not only reduces computational complexity in airfoil generation but also in optimisation, seeing that unfavourable designs can be more heavily penalised in the objective function without the requirement for any expensive downstream objective function evaluations.

To train both the BGMM and the BLSTM, the respective training sets are used to optimise the model parameters and the respective validation sets are used to evaluate the model performance. In both training instances, the optimal model hyperparameters are found by using a random search strategy. Specifically, in training the BGMM, the objective is to find the set of hyperparameters that minimise the Bayesian Information Criterion (BIC). Here the hyperparameters of interest (and their respective search boundaries) are as follows:

1. $2 \leq$ number of components ≤ 500,
2. covariance type: random choice between full, tied, diag (diagonal), and spherical.
3. weight concentration prior type: random choice between Dirichlet process, and Dirichlet distribution, and
4. $0.001 \leq$ weight concentration prior ≤ 5000.

In the search for the optimal BGMM parameters, a total of 1000 random search experiments are executed, each with a unique set of hyperparameters. It is important to note that the number of components hyperparameter directly controls the number of distinct airfoil clusters found in the data sets and hence dictates the quality of the conditional airfoil generation process. The higher the quality of the cluster separation, the more distinct each cluster is from the remaining clusters in the set – while ensuring homogeneity within the same cluster set. It is therefore important to ensure that the optimised BGMM modelyields a low BIC thereby ensuring optimal cluster separation and hence enabling effective conditional airfoil generation. The BGMM architecture that yielded the lowest BIC (−877 517.22) had the following architecture: (1) 325 components, (2) each component utilised a full covariance matrix, (3) made use of a Dirichlet process for the weight concentration prior, and (4) a weight concentration prior of 0.01.

For the training of the BLSTM, the objective is to minimise the Root Mean Squared Error (RMSE) between the predicted airfoil x and y coordinates, and the actual airfoil x and y coordinates. In the BLSTM random search pipeline, a total of 100 experiments are executed with a hyperparameter configuration between the bounds defined in Table 3. In each BLSTM random search experiment, the learning curve is monitored to ensure low bias and variance and hence high generalisation results, whilst sufficient RMSE is achieved. The learning curve is therefore monitored to ensure that the difference between the training loss and validation loss is sufficiently small, both the training and validation loss have the same decreasing trend, and the validation loss does not increase whilst the training loss continues to decrease. Furthermore,

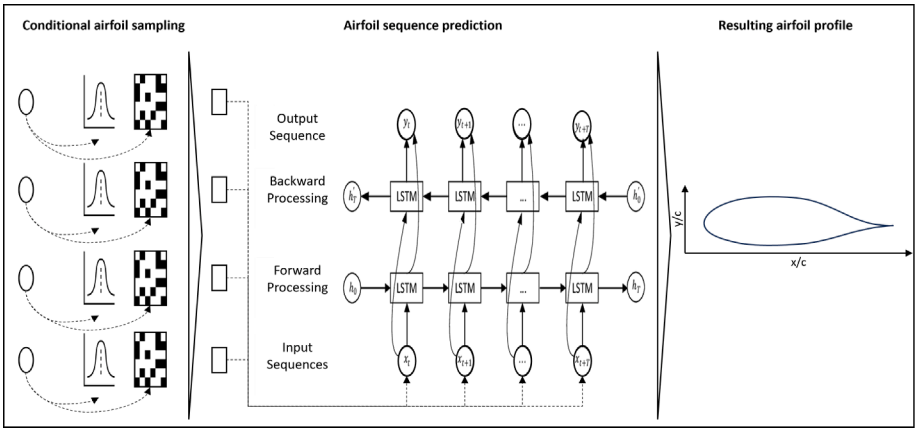

Fig. 2. Depiction of the inference process applicable when generating an airfoil.

Table 3
Search space for BLSTM hyperparemeter random search.

Parameter	Lower boundary	Upper boundary
Number of hidden layers	1	5
Units per layer	1	1024
Batch size	8	4096
RNN dropout probability	0	0.5
Learning rate	0.0001	0.1
Dropout probability	0	0.5

early stopping is employed as a regularisation technique to further mitigate the risk of overfitting. Here early stopping used to terminate training if no validation improvement is realised in the past 15 epochs.

The BLSTM model that yielded the lowest validation set RMSE of 8.52×10^{-7} (test set RMSE of 9.74×10^{-7}) had the following architecture: (1) an input layer with 23 neurons, (2) two hidden layers, each consisting of (a) an LSTM layer with 90 dense units, and (b) a BLSTM layer with 180 dense units, and (3) a fully connected output layer with 160 dense units. This means that on average the difference between the predicted airfoil profile x and y coordinates and the airfoil profiles in the unseen test set is 9.74×10^{-7}, indicating a high degree of accuracy in unseen airfoil reconstruction and demonstrating the generalisation ability of the model.

4.2. Conditioned airfoil generation

The conditional airfoil generation capacity is showcased in this section by means of a case study. In this case study, the proposed framework is used to generate 20 airfoils sampled in pairs of five, with each pair conditioned on a different airfoil cluster, i.e. clusters $5, 10, 15,$ 20. The resulting 20 airfoils sampled from the four mentioned clusters can be seen in Fig. 3.

From this figure, it can be seen that airfoils within the same cluster share similar airfoil profiles whereas there are notable differences between the profile characteristics when comparing airfoils from different clusters. Also clear from this figure is that there are no saw-tooth-like distributions on the upper or lower surface of the respective airfoils generated. It is also noted that no impractical, or non-airfoil, shapes result from the mapping between the lower-order parsimonious variable space to the higher-order x and y coordinate space. Finally, as visually depicted in Fig. 3, the airfoils within the same cluster share similar global profile characteristics but have slight local variations in nose radius, trailing edge reflex, camber, and thickness distributions, etc. This case study therefore highlights the fact that the BGMM, with the chosen model parameters, is able to effectively separate the 23 parsimonious variable characteristics of the 325 clusters and hence enabling accurate conditional airfoil generation when combined with the BLSTM.

4.3. Clark-Y airfoil reconstruction

As a validation of the proposed framework's ability to generalise on unseen airfoils, it is applied to the reconstruction of an out-of-sample airfoil profile. This reconstructed airfoil is then analysed to determine if it has the same airfoil profile characteristics and similar 2D airfoil performance when compared with the baseline design. The chosen airfoil for reconstruction analysis is the Clark-Y airfoil. The airfoil is reconstructed by means of calculating its 23 parsimonious variables numerically, using the x and y coordinates of the baseline airfoil, and then using these calculated parsimonious variables in the BLSTM to predict the profile of the reconstructed airfoil. The reconstruction results can be seen in Figure 4*a*. It is clear that there are minor deviations of the reconstructed airfoil from the baseline airfoil. When repanelling both the reconstructed airfoil and the baseline airfoil to have the same x coordinate base, the Mean Squared Error (MSE) is 9.93×10^{-7}. More specifically, this MSE represents the average Euclidean distance between the y coordinate position of the reconstructed and the baseline airfoils. This indicates that there is sufficient accuracy in profile reconstruction from the 23 parsimonious variables. For the reconstruction, the coefficient of determination is 0.999 indicating that there is a high degree ofexplained variance and hence correlation between the reconstructed and baseline airfoils.

To analyse the effect that the profile difference between the reconstructed airfoil and the baseline airfoil has on 2D aerodynamic performance, XFOIL simulations are employed. Here, the results of two XFOIL simulations are evaluated, i.e. a low Reynolds number case and a high Reynolds number case. The low Reynolds number case is set at 500 000, and the high Reynolds number case is set at 5,000,000, both with a Mach number setting of 0.15. The results of this XFOIL analysis is depicted in Fig. 4b. Here it can be seen that for both simulations, the reconstructed and baseline airfoils have similar lift and drag performance with an R^2 of 0.999 for the low Reynolds number case and a R^2 of 0.998 for the high Reynolds number case respectively.

This high average R^2 (0.999) for the two cases investigated, illustrates the generalisation capacity of the proposed framework in the reconstruction of an out-of-sample airfoil. Also seen from this analysis is that there is further scope for improving model reconstruction perfor-

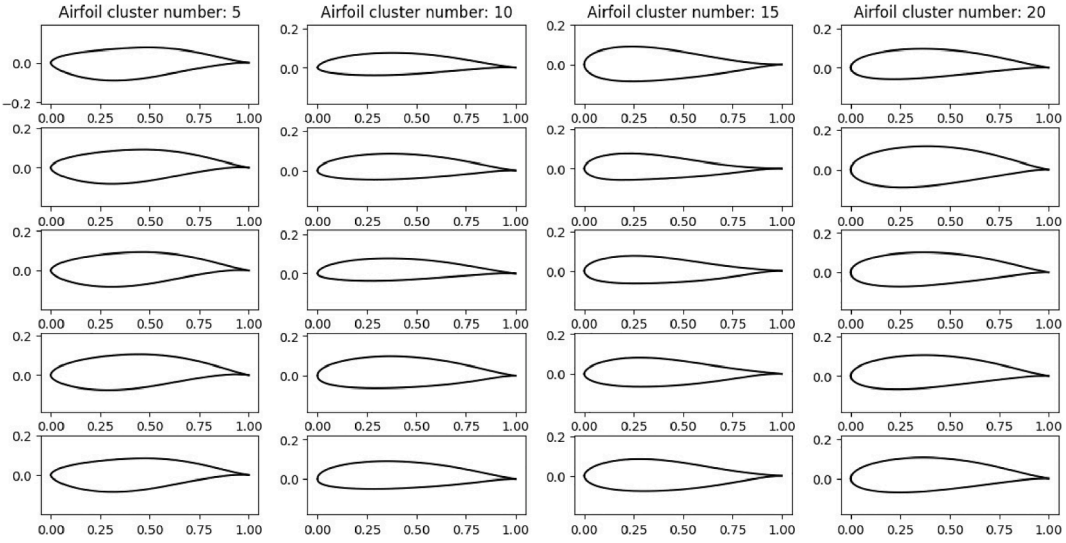

Fig. 3. Randomly sampled airfoils conditioned on clusters $5, 10, 15,$ and 20.

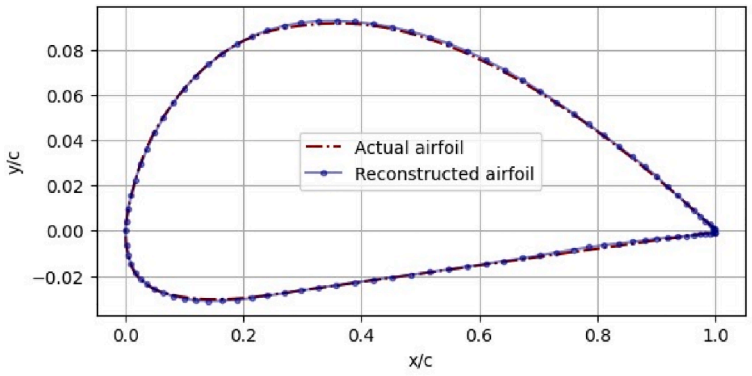

(a) Profiles of the baseline and reconstructed Clark-Y airfoil

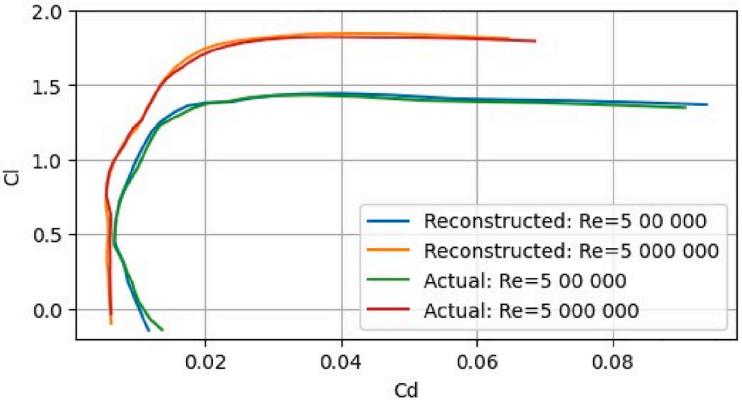

(b) XFOIL simulated drag polar of the reconstructed airfoil andthe baseline airfoil for two different operating conditions

Fig. 4. Reconstruction error analysis in terms of profile and 2D aerodynamic performance comparison.

mance. This can be achieved by further optimising the model parameters and underlying architectures as well as increasing the volume, variety, and veracity of the airfoil shape dataset.

4.4. Targeted airfoil generation efficiency comparison

As mentioned before, one of the key advantages of the proposed framework is the ability to efficiently generate airfoil samples with explicitly defined shape and structural characteristics. To quantify the efficiency improvements in targeted airfoil generation, we perform a case study, comparing the proposed framework with the SPM. In this study, the proposed framework and the SPM are used to generate 100 airfoils respectively, each with the shape characteristics within the upper and lower bounds as defined in Table 4.

For the proposed framework, these 100 airfoils are generated by randomly perturbing the input latent space of the BGMM for a randomly sampled airfoil cluster. The shape characteristics of the resulting airfoil are then evaluated to ascertain if they fall within the desired upper and lower bounds. If the shape characteristics of interest fall within the bounds, the BLSTM is used to generate the airfoil x and y coordinates, if not, the sample is disregarded and a new sample is generated by restarting the process. This process is repeated until a data set results with 100 airfoils – each possessing the shape characteristics as defined in Table 4.

Table 4
Shape characteristics and their respective upper and lower bounds used to guide constrained airfoil generation when comparing the computational complexity of SPM and the proposed framework.

Constraint	Lower bound	Upper bound
Max. thickness magnitude	0.120	0.155
Max. thickness location	0.350	0.600
Max. top magnitude	0.080	0.150
Max. top magnitude location	0.300	0.650
Max. camber location	0.350	0.650
Airfoil area	0.100	0.135

For the SPM, the generation of these 100 airfoil samples involves firstly perturbing the six parametric coefficients, within their respective upper and lower bounds, to generate an airfoil profile in terms of x and y coordinates. The coordinates of this sample are then used in an additional calculation procedure where the airfoils' shape characteristics are numerically determined. The shape characteristics are then evaluated to determine if they fall within the desiredupper and lower boundaries. If this is the case, the sample is added to the data set, if not, the process is restarted. Once again, this process is repeated until the data set contains 100 airfoil profiles, each with the desired shape characteristics.

On average, when executed using the same hardware and with comparable implementations, it takes the proposed framework approximately 0.69 s to generate a single airfoil with explicitly constrained shape characteristics, compared to the 87.03 s it takes the SPM for the same task. This indicates that, compared to the SPM, the proposed framework achieves approximately a 126 times reduction in the time required to generate an airfoil with specific shape characteristics. There are mainly two reasons for this reduction. The first is due to the design of the proposed framework. The proposed framework decouples the generation of an airfoil such that the profile is first generated in

terms of the 23 descriptive features before the x and y coordinate generation is invoked. This holds the advantage of being able to reject an airfoil if the 23 descriptive features are not within the desirable ranges before the more expensive x and y coordinate generation process is executed. This is in contrast with the SPM, and the majority of constructive and destructive airfoil generation methods, that first generate airfoils in terms of x and y coordinates before their shape and structural characteristics are numerically extracted by addi-tional downstream processes. The second reason for the improved efficiency of the proposed framework is due to the method's proclivity for generating practical airfoils as opposed to the SPM, which has a high propensity of generating impractical airfoils when randomly perturbing the six input parameters.

It is interesting to note that, not only is the proposed framework more efficient in generating airfoils with explicitly defined characteristics but also are the resulting profiles more diverse than those generated by the SPM. To highlight this property refer to Fig. 5, depicting six randomly sampled airfoil profiles from the 100 generated airfoils from the SPM and proposed framework data sets. From this figure, it can be seen that there is more variability, and hence diversity, in the airfoil profiles generated with the proposed framework compared to the airfoil profiles generated with the SPM.

To quantify this difference in variability, the mean, standard deviation (SD), and coefficient of variation (CV) (Wu & Huang, 2023) are calculated for various shape characteristics for each of the 100 airfoils generated with the SPM and the proposed framework respectively. Here, the CV is defined as:

$$CV = \frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(x_i - \frac{1}{n}\sum_{i=1}^{n}x_i\right)^2}}{\frac{1}{n}\sum_{i=1}^{n}x_i} \times 100\% \tag{9}$$

Where x_i is the shape characteristic of the i^{th} airfoil sample and n is the totalnumber of airfoils. The results of this analysis can be seen in Table 5.

From this table, it can be seen that – on average – the airfoils generated with the proposed framework have a higher CV than those generated with the SPM. The airfoils generated with the proposed framework have more variability in the position and magnitude of the total airfoil thickness, the position and location of the maximum top surface magnitude, the maximum camber magnitude, and the airfoil nose radius; whereas the SPM-generated airfoils have more variability in the airfoil area and the location of the maximum camber. This increased variability is a result of the proposed framework being able to condition on various airfoil clusters that already possess desirable shape characteristics (within the defined bounds) on average, and then bringing minor profile alterations by perturbing the underlying BGMM latent code for that specific cluster. This characteristic of the proposed framework to produce an airfoil population with more diversity, whilst adhering to explicitly defined constraints, is especially favourable when constrained airfoil optimisation cases are considered. This is because the more diversity injected in the optimisation routine, the higher the probability of finding an airfoil with favorable performance characteristics – as will be seen later in Section 5.3.

Table 5
Comparison of the variance in the shape characteristics of the 100 airfoils generated with the SPM and the proposed framework (PF).

Feature	Mean		SD		CV	
	PF	SPM	PF	SPM	PF	SPM
Max. thickness	0.1484	0.1486	0.0053	0.0042	3.55	2.84
Max. thickness location	0.4069	0.4154	0.0412	0.0363	10.12	8.74
Max. top magnitude	0.0984	0.0899	0.0104	0.0052	10.54	5.83
Max. top magnitude location	0.4543	0.4361	0.0405	0.0309	8.92	7.08
Max. camber	0.0252	0.0154	0.0108	0.0056	42.65	36.16
Max. camber location	0.5492	0.5126	0.0839	0.0905	15.28	17.65
Nose radius	0.0194	0.0202	0.0051	0.0052	26.52	25.56
Airfoil area	0.1040	0.1052	0.003	0.0032	2.87	3.09

4.5. Padulo completeness discussion

The proposed framework is evaluated in terms of its adherence to Padulo's five criteria, i.e. completeness, orthogonality, flawlessness, parsimonious, and practicality.

The first dimension of evaluation is that of completeness, which requires an airfoil parametrisation method to be able to describe any airfoil profile to a specific degree of accuracy. In the proposed method, the completeness of the method is mainly a function of the underlying training data set, which is a function of the airfoil generation method used and the complexity (and variety) of the descriptive features used in the airfoil parameterisation framework. For the airfoil generation method used, as the data set is mined from airfoils generated with the SPM, the proposed framework's completeness is directly linked to the completeness of the SPM. This is because the deep learning architecture chosen is not able to produce airfoil samples outside of the probability distribution of shape and structural features of the underlying training set (while the airfoil generator is able to generate unseen airfoils, these are linked to the features of the underlying training set.).

This means that the designer can choose the level of completeness required by strategically expanding or narrowing the diversity of the underlying training set. This completeness expansion can be done by either: (1) making use of multiple constructive or destructive airfoil generation models when mining airfoil profile data, (2) using a larger number of descriptive features in the airfoil parameterisation framework, or (3) a combination of multiple airfoil generation techniques and a more complex airfoil parameterisation framework. Similarly, the

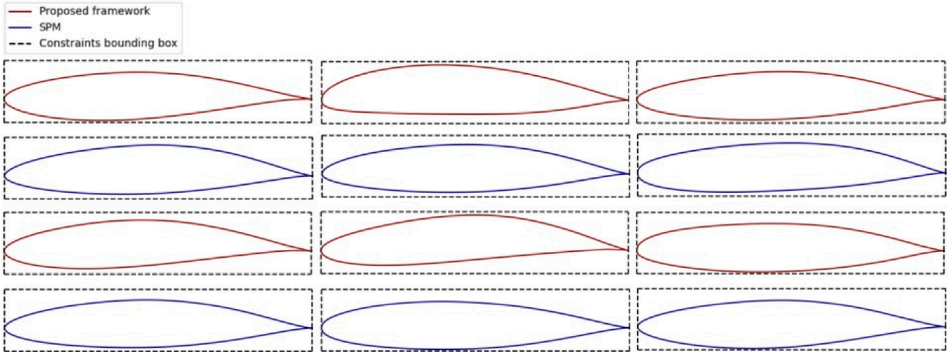

Fig. 5. Six randomly sampled airfoils resulting from the targeted airfoil generation case study with the SPM and proposed framework, showcasing the differences in variability between the airfoils generated with the respective methods.

narrowing of the proposed methods' airfoil coverage can be achieved by making use of a single method for airfoil generation and imposing constraints on the inputs to the chosen generation model – thereby limiting the airfoil profile diversity of the proposed method, or by using simpler and fever variables in the airfoil parameterisation framework.

The next dimension of analysis is the orthogonality of the proposed framework, that is, the ability to represent a unique airfoil for a unique set of inputs.

The proposed framework has orthogonality implicit in its design. This is because the BLSTM will always yield a unique set of outputs for a unique set of inputs seeing that the underlying matrix multiplication result is unique for each unique input.

The same can not be said when discussing the proposed framework's adherence to the parsimony criteria, i.e. the ability to introduce significant change to the main geometric features of the generated airfoil with limited parameters. This is because a total of 23 variables, 17 shape and 6 structural, are used to parametrise any given airfoil with the proposed framework. It is important to note that there is an inverse relationship between completeness and parsimony such that, the more complete a given airfoil generation method, the more parameters are required to describe that specific airfoil's profile. Stated differently, maximising the coverage (completeness) of an airfoil generation method comes at the cost of increasing the number of variables required to describe an airfoil profile, and hence reduces the framework's adherence to Padulo's parsimony criteria. Given that the main requirement of the proposed method is to enable explicit airfoil generation for a diverse airfoil population, the method fails to completely adhere to Padulo's parsimony criteria.

The next dimension of analysis is the flawlessness of the proposed framework. The flawlessness of the proposed framework is a function of the underlying data set and the deep learning architecture chosen. Seeing that the BLSTM is trained to map the 23 descriptive features to the corresponding airfoil profile in terms of its x and y coordinates, the model is only capable of generating impractical airfoil shapes if:

- there are representative samples of such airfoil profiles in the training data set,
- the 23 descriptive features of an impractical airfoil is supplied as input to the BLSTM, or
- the 23 descriptive features of an airfoil sample outside the distribution of the training data set is supplied to the input of the BLSTM.

Great care was taken with the construction of the airfoil corpus (relating the 23 descriptive features to the corresponding x and y coordinates) to ensure that there are no impractical airfoils, thus mitigating the risk that could implicitly arise from low-quality airfoil data. With this risk mitigated, there still is the possibility of generating impractical airfoil profiles with the BLSTM, if the descriptive features used as input to the BLSTM fall outside of the training distribution, or are representative of impractical airfoil shapes. The use of the BGMM is the key mechanism used to mitigate this risk. This is because, for a given cluster choice when randomly sampling the latent code from the normal distribution, the BGMM will always yield 23 descriptive features that closely resemble the underlying distribution of the training data. This means that under the condition that the BGMM is correctly constructed using sufficient training data – the BLSTM will always be supplied with: (1) the 23 descriptive features of a practical and realistic airfoil sample, and (2) the 23 descriptive features of an airfoil that has the same underlying probability distribution and co-variance as in the training data. This means the proposed framework can, under these conditions, be considered as flawless and adheres to this criteria of Padulo.

The final dimension of analysis is the proposed framework's adherence to Padulo's intuitiveness criteria. That is the property of the parameters used in generating the airfoil profile having explicit physical and intuitive meaning such that it simplifies generation bound setting and design constraining. As seen in Section 2, the 17 shape and 6 structural parameters all have direct intrinsic and intuitive meaning thereby facilitating targeted airfoil generation. For this reason, the proposed framework can be considered to adhere to Padulo's intuitiveness criteria.

5. Optimisation framework

The development of the optimisation framework, it's application to multiobjective constrained airfoil optimisation, and the analysis of the optimisationresults are discussed in the preceding sections.

5.1. Problem formulation

With the newly developed airfoil generation framework, multiobjective airfoil optimisation with constraints on airfoil shape and structural characteristics is possible. To illustrate this flexibility, an airfoil optimisation case study is considered where the main goal is to develop an airfoil with thick sectional properties for use as the root airfoil in sailplane cross-country applications. This optimisation should take place whilst adhering to explicitly defined shape and structural constraints defined upfront as input to the optimisation pipeline. The upper and lower boundaries of the shape and structural constraints can be seen in Table 6. These constraints should be adhered to whilst the lift-to-drag performance over a range of operating conditions is maximised. This objective can mathematically be expressed as:

$$\beta = \sum_{i}^{15} \sum_{j}^{5e6} \left[\frac{C_{l_i}}{C_{d_i}}\right]_j \text{ with } \begin{cases} i = \{-15, -14.9, \ldots, 15\} \\ j = \{2e6, 3e6, \ldots, 5e6\} \end{cases} \tag{10}$$

with C_l and C_d the 2D lift and drag coefficients respectively for the angle of attack i and Reynolds number j. As mentioned, the second objective evaluated during the airfoil optimisation is the maximisation of the airfoil area. The final fitness function used to guide the optimisation algorithm can therefore beexpressed as:

$$\text{Fitness} = (\beta \text{ x } \delta) - (\gamma \text{ x } \sigma) \tag{11}$$

where β is the objective function as defined in Equation (10), δ is the total airfoil area, γ is the total penalty value for any deviation from the desirable parsimonious characteristics, and σ is a scaling parameter used the scale the importance of the shape deviation penalty in the final fitness function.

5.2. Optimisation process

The proposed framework is embedded into an optimisation pipeline as shown in Fig. 6. The optimisation pipeline requires the following four elements: (1) a flow solver, (2) a method to produce realistic airfoil parsimonious variables, (3) a method for converting airfoil parsimonious variables to raw x and y coordinates, and (4) an algorithm to efficiently exploit the parsimonious design space in order to optimise the underlying objective function.

As a flow solver, the panel code XFOIL is used to evaluate the aerodynamic performance. In each new XFOIL run the boundary layer is re-initialised and the N_{crit} value is set to 9. Each candidate airfoil generated is smoothed and re-panelled using 300 panels. Finally, 500

Table 6
Upper and lower bounds on airfoil shape and structural features.

Feature	Lower bound	Upper bound
Max. thickness	0.145	0.200
Max. thickness location	0.350	0.600
Max. camber	0.015	0.055
Max. top magnitude	0.100	0.150
Max. top magnitude location	0.300	0.650
Airfoil area	0.120	0.175

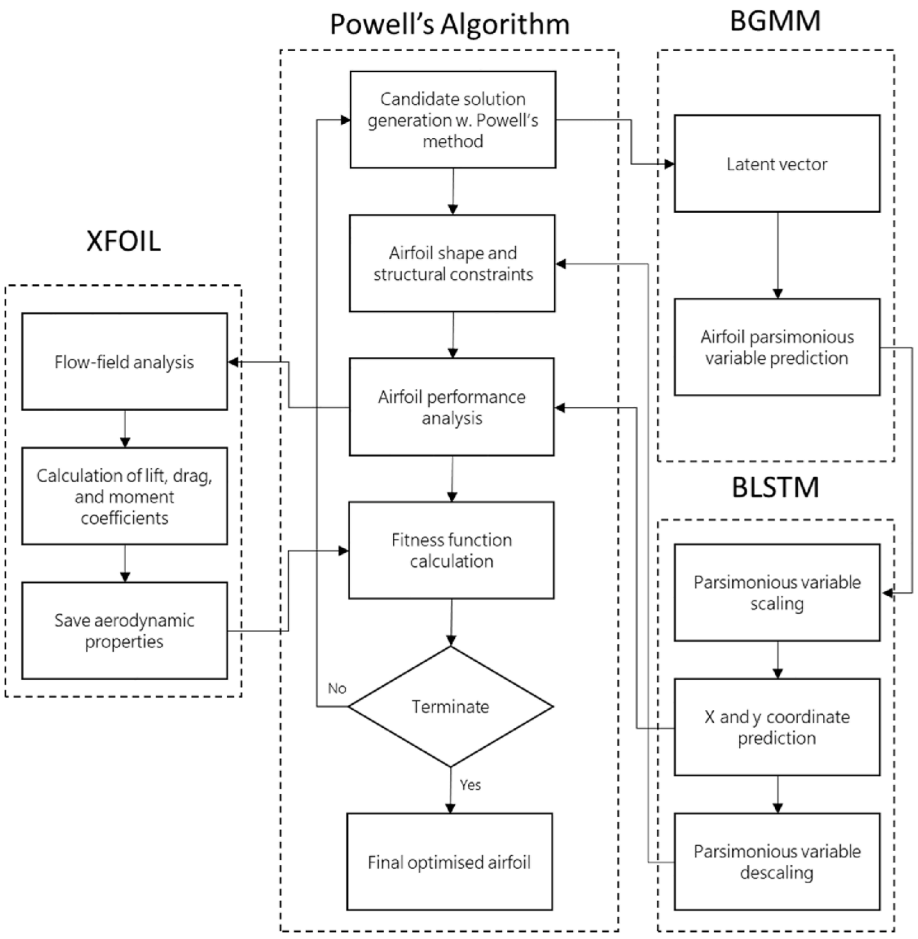

Fig. 6. Airfoil optimisation with constraints set on the shape and structural characteristics.

iterations for each XFOIL simulation are used to allow the global Newton-Raphson solver to converge.

For optimisation, the Powell method (Brent, 1973) is used to adjust the BGMM latent vector in the direction most likely to optimise the objective function – whilst adhering to the defined constraints. As seen in Equation (10), in each optimisation loop, the aerodynamic performance is evaluated over the Reynolds number ranges of 2e6 to 5e6 (in increments of 1e6) with angle of attack ranging from -15° to 15° (in increments of 0.1).

Finally, for realistic airfoil generation in terms of 23 parsimoniously defined airfoil characteristics and sequence mapping from the parsimonious domain to the x and y coordinate domain, the proposed framework as discussed in Section 4.1 is used.

5.3. Comparison with traditional optimisation techniques

To evaluate the efficacy of the proposed framework, its optimisation results are compared with those of a traditional optimisation routine. This traditional optimisation routine follows the same process flow as depicted in Fig. 6, except here, the SPM is used as the airfoil generation technique. This means that in the traditional optimisation pipeline, there is an additional calculation step where the airfoil shape characteristics are numerically calculated from the x and y coordinate representation from the generated airfoil. From here all steps regarding

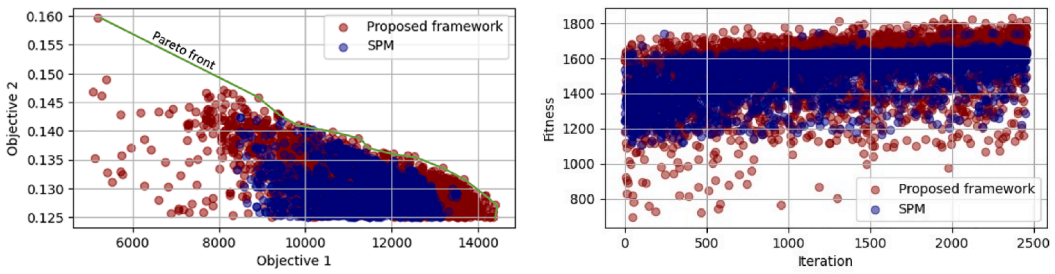

(a) Resulting airfoil performance distribution and (b) Solution fitness evolution for SPM and proPareto front for the two objectives of interest posed framework optimisation procedures

Fig. 7. Comparison between the solution distributions for the two optimisation objectives as well as the fitness evolution for the SPM and proposed framework constrained optimisation procedures. (a) Resulting airfoil performance distribution and (b) Solution fitness evolution for SPM and proPareto front for the two objectives of interest posed framework optimisation procedures.

constraining and optimisation objectives remain the same.

In both the traditional and the proposed framework optimisation pipelines, a total of 62,500 candidate airfoils are generated during the optimisation process. Of these candidate solutions, roughly 2400 fall within the upper and lower defined constraints. In Fig. 7a the Pareto front, formed by these respectivesolutions, can be seen. Here, objective 1 refers to Equation (10) whereas objective 2 refers to the total airfoil area. From this figure, it is interesting to note the differences in the variance in the performance of the airfoil designs generated by the SPM and proposed framework optimisation pipelines. The proposed framework produces designs with areas closer to the upper and lower boundaries as defined in Table 6. This indicates that the proposed framework is capable of a more thorough search of the design space as it navigates through the various airfoil clusters used to condition the airfoil generation. This is opposed to the more narrow exploration done by the SPM optimisation pipeline. This is likely because, although wide bounds are set on the SPM parameters meaning diverse designs will result, the method is capable of producing impractical airfoil profiles. This has the result of the optimisation algorithm more prudently exploring the design space near solutions already found to adhere to the airfoil shape constraints.

It is important to note that the two objectives evaluated here have contradictory effects, where high area leads to reduced overall lift-to-drag performance. This is because an increase in area is proportional to an increase in an airfoil's drag coefficient. However, when considering root airfoils, a high area is required to ensure the structural integrity of the wing. Therefore a balance between performance and area is required and the optimal design is defined as the airfoil,or set of airfoils, that has the highest area and lift-to-drag performance that is, the airfoil that evaluates to the maximum value of Equation (11).

Fig. 7bdepicts the fitness evolution of the airfoils optimised with the SPM and the proposed framework, respectively. It can be seen that, for the first 250 iterations, the airfoils produced with the SPM and the proposed framework have comparable fitness values; however, there is more variability present in the fitness of the airfoils produced with the proposed framework. From iteration 250 onwards, the fitness of the airfoils produced with the proposed framework is consistently higher than the fitness of those airfoils produced with the SPM. This positive difference between the proposed framework fitness and the SPM fitness gradually starts to increase, reaching a maximum at the point of simulation termination. This indicates that the average increase per generation is higher for the proposed framework than for the SPM. This comparison highlights the efficiency improvements in constrained optimisation achieved with the proposed framework compared to the SPM. In this case study, the proposed framework was not only capable of finding more favorable designs but also did so in fewer iterations on average, resulting in reduced computational requirements during optimisation.

The airfoil profiles that maximises Equation (11), while adhering to the upfront defined shape constraints, can be seen in Fig. 8. The SPM optimised airfoil has a maximum thickness of 0.1981 at $x = 0.390$, a maximum camber of 0.0504 at $x = 0.401$, a maximum top thickness of 0.1495 at $x = 0.392$, and a total area of 0.1291. The proposed framework optimised airfoil has a maximum thickness of 0.1872 at $x = 0.380$, a maximum camber of 0.0527 at $x = 0.442$, a maximum top thickness of 0.1490 at $x = 0.4150$, and a total area of 0.1233. From this analysis it can therefore be concluded that both the optimised airfoils adhere to the upfront defined constraints set as input to the optimisation pipeline.

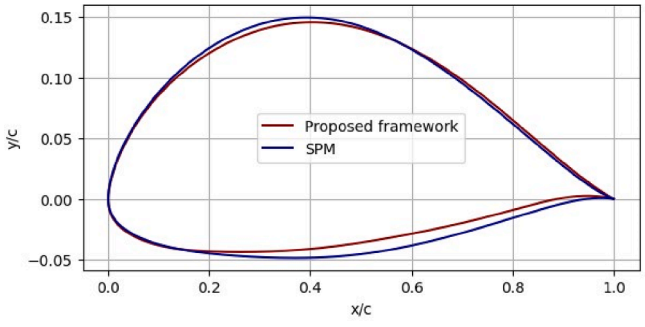

Fig. 8. Optimised airfoil profiles resulting from the constrained SPM and proposed framework optimisation pipeline.

5.4. Further analyses

As seen in Fig. 8, there are distinct differences between the profiles of the SPM optimised airfoil and the proposed framework optimised airfoil. Specifically, the proposed framework optimised airfoil has a higher degree of camber, slightly lower surface area, and interestingly, a higher degree of curvature present at the airfoil trailing edge.

To analyse the differences between the 2D lift-to-drag performance of the two optimised airfoils, XFOIL is employed. The optimised airfoils are analysed at a Reynolds number of 2e6 and 5e6 respectively, with angle of attack ranging from -15° to 15° (similar to the conditions for which they were optimised). The results of this analysis can be seen in Fig. 9. Here, it is observed that the proposed framework optimised airfoil: (1) significantly outperforms the SPM optimised airfoil between the angle of attack ranges of -15° and 6°, with an average increase of 19.75% in lift-to-drag performance over this region, and (2) has lower lift-to-drag performance between the angle of attack ranges of 7° and 15°, with an average decrease of 2.52% in lift-to-drag performance over this region. Overall, the average relative performance increase of the proposed framework optimised airfoil over the SPM optimised airfoil is 9.04% for the two Reynolds numbers and angle of attack range investigated. In this analysis, the proposed framework was therefore able to achieve improved airfoil optimisation results compared to the SPM. when controlling for all other parameters.

One of the reasons for this was alluded to in Section5.3, where it was shown how the proposed framework is able to explore the input design space more efficiently. To investigate this property of the proposed framework, an analysis is conducted into the exploration versus exploitation decisions made.

Fig. 10 shows the fitness evolution of the solutions that adhered to all shape constraints, as evaluated during the optimisation process. Here, the optimisation results are segmented by airfoil cluster to highlight the exploration versus exploitation process at play. As can be seen in this figure, up to 1 250 iterations a strong exploration process is at play,

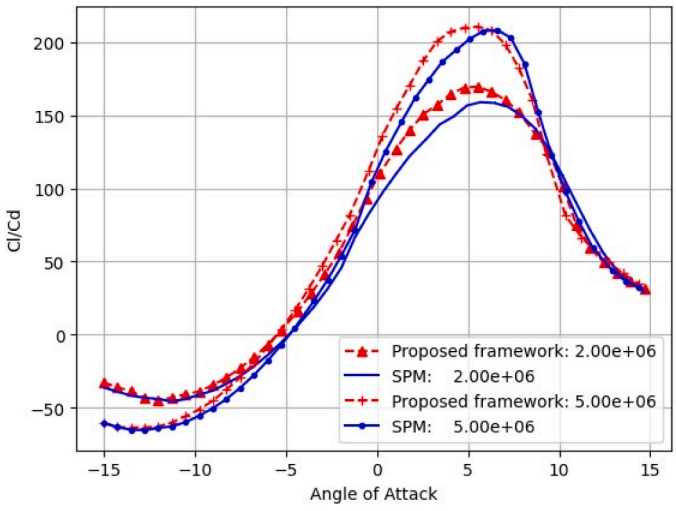

Fig. 9. XFOIL performance comparison for the SPM optimised and proposed framework optimised airfoils at two different operating conditions.

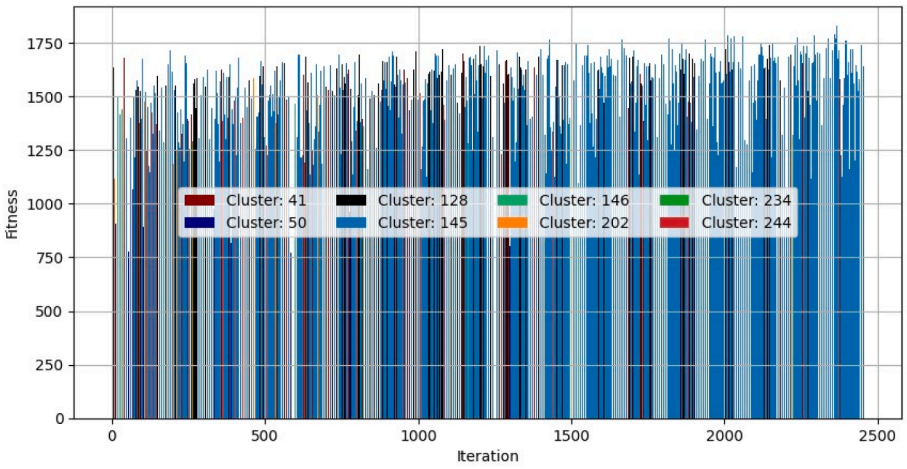

Fig. 10. Airfoil fitness (Eq. (10) segmented by airfoil cluster number, measured per iteration during optimisation pipeline execution – highlighting the evolution of the exploration versus exploitation decisions made by the proposed framework.

where the algorithm is conditioning on a wider variety of airfoil clusters that adhere to defined shape constraints. In this process, the algorithm explores eight of the 325 clusters and optimises the BGMM latent dimensions in order to improve airfoil performance.

From iteration 1250 up to termination it is clear that a strong exploitation process is present where latent codes of the clusters that adheres to the shape constraints and has proven to yield good performance are modified. The main airfoil cluster here is cluster 145, and it is observed that as more optimisation iterations are executed, the stronger the presence of this airfoil cluster up till optimisation termination. This ability to initially perform wide design explo-ration before refining the designs of the best performing clusters is one of the main factors that enables the proposed framework to achieve better results than the traditional method employed in this case study.

6. Discussion

As seen in Section 5.4 the proposed framework has demonstrated promising results in multi-objective airfoil optimisation, with multiple constraints imposed on the airfoil shape characteristics. One of the reasons for this is that the proposed framework has a well behaved input design space. This means that similar airfoil designs are adjacent in a higher dimensional space and therefore provide a strong signal to follow during optimisation for non-gradient based optimisation algorithms. That is, during optimisation, this well behaved input design space means that, once the optimisation algorithm has found a cluster and latent code combination that yields good results, minor perturbations to the latent code/cluster combinations will result in gradual changes in airfoil perfor-mance. This leads to small performance increases or decreases over successive optimisation iterations and hence provides a strong optimisation signal.

Another reason for improved optimisation performance when using the proposed framework is the explicit control over airfoil generation. The proposed framework provides local control through exploitation of BGMM latent codes and global control via strategic cluster selection. This means that once the optimisation algorithm finds airfoil clusters with suitable performance, that adheres to the all imposed constraints, then local variations can gradually be introduced in an attempt to maximise the objective function.

Finally, an important property of the proposed framework that adds significantly to its optimisation performance is the fact that it has a very low probability of generating impractical airfoil shapes. This means the optimisation algorithm is injected with more diversity in successive iterations and can spend time improving the airfoil design rather than learning what designs are practical or not. This further adds to the optimisation efficacy when employing the proposed framework.

7. Conclusion

An airfoil parameterisation framework has been developed that enables conditional airfoil generation with desired shape and structural characteristics. This is achieved by leveraging BLSTMs and BGMMs to map the lower-order parsimonious representation of a selected airfoil cluster, to its higher order x and y coordinate representation.

A set of 17 shape and six structural parameters are identified that are used as the set of 23 parsimonious variables to explicitly express an airfoil profile in a lower-order, interpretable manner. An airfoil corpus for training the BLSTM and BGMM is acquired by means of setting up a data mining pipeline, extracting the 23 parsimonious variables numerically for the x and y coordinates generated with the SPM, although any one (or multiple) constructive and/or destructive airfoil generation methods can be used in this endeavour.

It is shown that the proposed framework is capable of conditional airfoil generation by means of specifying the airfoil cluster of interest and perturbation of the BGMM latent codes. Furthermore, it is demonstrated how the proposed framework enables robust and explicit control over various airfoil shape characteristics such as total airfoil thickness, maximum camber magnitude, etc.

The proposed framework is then used in multi-objective airfoil optimisation under various shape constraints. Comparison between the proposed framework and traditional methods reveals the former's ability to explore a wider design space efficiently, leading to greater variability in airfoil designs that comply with defined constraints. The efficiency in design space exploration resulted in an overall improvement of 9.04% in the lift-to-drag performance of the airfoil optimised with the proposed framework.

Overall, this work demonstrates the effectiveness of the proposed framework in constructing, analysing, and optimising airfoil designs. By efficiently exploring the design space, maintaining compliance with shape constraints, and achieving enhanced aerodynamic performance, the framework presents promising potential for advancing airfoil design in various engineering applications.

The proposed framework, however, has some limitations that should be the subject of future study to enhance the efficacy and utility of the proposed framework. The first is the coverage of the framework. Although the selected 23 parsimonious variables ensure coverage of a significantly large and diverse set of airfoil profiles, it lacks the capacity to sufficiently represent airfoils with complex tail geometry. This is because, in the current set of 23 parsimonious variables, there are no parameters for describing the trailing edge wedge angle or the trailing

edge thickness. The current set of descriptive parameters also has no way of describing the leading edge when the top and bottom surfaces have differing leading edge radii. Furthermore, adding additional descriptive features to the framework such as the curvature present at the location of the maximum camber magnitude may further extend the complexity of airfoils that can be represented by the proposed framework, as well as enhance the accuracy of the mapping from the lower-order domain to the higher-order domain.

Finally, in order to increase the generation efficiency, as well as reduce the sequence mapping inference requirements, it is recommended that alternative architectures be investigated. For example, replacing the BGMM architecture with a Generative Adversarial Architecture may hold the advantage of reducing the latent dimension size whilst improving the modelling of the underlying conditional data distribution. Also, using alternative architectures for the sequence mapping problem, such as Transformers, may lead to more accurate airfoil reconstruction as well as reduced computational requirements – in both training and inference.

CRediT authorship contribution statement

Vincent le Roux: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Marelie H. Davel:** Conceptualization, Methodology, Software, Validation, Formal analysis, Writing – original draft, Writing – review & editing, Supervision. **Johan Bosman:** Conceptualization, Methodology, Resources, Writing – original draft, Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

Achour, G., Sung, W. J., Pinon-Fischer, O. J., & Mavris, D. N. (2020). Development of a conditional generative adversarial network for airfoil shape optimization. In *In AIAA Scitech 2020 Forum*. https://doi.org/10.2514/6.2020-2261

Brent, R. P. (1973). Algorithms for minimization without derivatives. *PrenticeHall.*

Chen, W., & Fuge, M. (2018). Béziergan: Automatic generation of smooth curves from interpretable low-dimensional parameters. URL: https://arxiv.org/abs/1808.08871. https://doi.org/10.48550/ARXIV.1808.08871.

Christie, R., Robinson, M., Tejero, F., & MacManus, D. G. (2019). The use of hybrid intuitive class shape transformation curves in aerodynamic design. *Aerospace Science and Technology, 95*, 105473. URL: https://www.sciencedirect.com/science/article/pii/S1270963819305450. https://doi.org/10.1016/j.ast.2019.105473.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial networks. URL: https://arxiv.org/abs/1406.2661. https://doi.org/10.48550/ARXIV.1406.2661.

Hamdi, Y., Boubaker, H., Rabhi, B., Qahtani, A. M., Alharithi, F. S., Almutiry, O., Dhahri, H., & Alimi, A. M. (2022). Deep learned blstm for online handwriting modeling simulating the beta-elliptic approach. *Engineering Science and Technology, an International Journal, 35*, 101215. URL: https://www.sciencedirect.com/science/article/pii/S2215098622001240. https://doi.org/10.1016/j.jestch.2022.101215.

Hicks, R. M., & Henne, P. A. (1978). Wing design by numerical optimization. *Journal of Aircraft, 15*, 407–412. arXiv: https://doi.org/10.2514/3.58379.

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation, 9*, 1735–1780. URL: https://doi.org/10.1162/neco.1997.9.8.1735.

Kingma, D. P., & Welling, M. (2019). An introduction to variational autoencoders. *Foundations and Trends in Machine Learning, 12*, 307–392. https://doi.org/10.1561/2200000056

Kou, J., Botero-Bolívar, L., Ballano, R., Marino, O., de Santana, L., Valero, E., & Ferrer, E. (2023). Aeroacoustic airfoil shape optimization enhanced by autoencoders. *Expert Systems with Applications, 217*, 119513. URL: https://www.sciencedirect.com/science/article/pii/S0957417423000143. https://doi.org/10.1016/j.eswa.2023.119513.

Kulfan, B. M. (2008). Universal parametric geometry representation method. *Journal of Aircraft, 45*, 142–158. https://doi.org/10.2514/1.29958.

Li, D., Yao, S., Liu, Y.-H., Wang, S., & Sun, X.-H. (2016). In *Efficient design space exploration via statistical sampling and adaboost learning*. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2897937.2898012.

Lu, J. (2021). A survey on bayesian inference for gaussian mixture model. arXiv: 2108.11753.

Lu, X., Huang, J., Song, L., & Li, J. (2018). An improved geometric parameter airfoil parameterization method. *Aerospace Science and Technology, 78*, 241–247. URL: https://www.sciencedirect.com/science/article/pii/S1270963817304686. https://doi.org/10.1016/j.ast.2018.04.025.

Masters, D. A., Taylor, N. J., Rendall, T. C. S., Allen, C. B., & Poole, D. J. (2017). Geometric comparison of aerofoil shape parameterization methods. *AIAA Journal, 55*, 1575–1589. https://doi.org/10.2514/1.J054943

McLachlan, G. J., Krishnan, T., & Ng, S. K. (2004). The em algorithm. URL: http://hdl.handle.net/10419/22198.

Mukesh, R., Lingadurai, K., & Selvakumar, U. (2014). Airfoil shape optimization using non-traditional optimization technique and its validation. *Journal of King Saud University - Engineering Sciences, 26*, 191–197. https://doi.org/10.1016/j.jksues.2013.04.003. Thermal and Micro structure Properties.

Nemati, M., & Jahangirian, A. (2020). Robust aerodynamic morphing shape optimization for high-lift missions. *Aerospace Science and Technology, 103*, 105897. URL: https://www.sciencedirect.com/science/article/pii/S1270963820305794. https://doi.org/10.1016/j.ast.2020.105897.

Padulo, M., Maginot, J., Guenov, M., & Holden, C. (2009). Airfoil design under uncertainty with robust geometric parameterization. In *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. URL: https://arc.aiaa.org/doi/abs/10.2514/6.2009-2270. https://doi.org/10.2514/6.2009-2270. arXiv: https://arc.aiaa.org/doi/pdf/10.2514/6.2009-2270.

Piegl, L., & Tiller, W. (1995). The NURBS book. https://doi.org/10.1007/978-3-642-97385-7.

Samareh, J. A. (2001). Survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimization. *AIAA Journal, 39*, 877–884. https://doi.org/10.2514/2.1391

Schmidt, R. M. (2019). Recurrent neural networks (rnns): A gentle introduction and overview. arXiv:1912.05911.

Sekar, V., Zhang, M., Shu, C., & Khoo, B. C. (2019). Inverse design of airfoil using a deep convolutional neural network. *AIAA Journal, 57*, 993–1003. https://doi.org/10.2514/1.J057894

Sheikh, H. M., Lee, S., Wang, J., & Marcus, P. S. (2023). Airfoil optimization using Design-by-Morphing. *Journal of Computational Design and Engineering, 10*, 1443–1459. URL: https: //doi.org/10.1093/jcde/qwad059. 10.1093/jcde/qwad059 arXiv:https://academic.oup.com/jcde/article-pdf/10/4/1443/50882006/qwad059.pdf.

Sherstinsky, A. (2020). Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena, 404*, Article 132306. https://doi.org/10.1016/j.physd.2019.132306

Sobieczky, H. (1999). Parametric airfoils and wings. Recent Development of Aerodynamic Design Methodologies, (pp. 71–87).

Toal, D. J. J., Bressloff, N. W., Keane, A. J., & Holden, C. M. E. (2010). Geometric filtration using proper orthogonal decomposition for aerodynamic design optimization. *AIAA Journal, 48*, 916–928. URL: 10.2514/1.41420 doi 10.2514/1.41420 arXiv:https://doi.org/10.2514/1.41420.

Wada, K., Suzuki, K., & Yonekura, K. (2024). Physics-guided training of gan to improve accuracy in airfoil design synthesis. *Computer Methods in Applied Mechanics and Engineering, 421*, 116746. URL: https://www.sciencedirect.com/science/article/pii/S0045782524000021. https://doi.org/10.1016/j.cma.2024.116746.

Wang, J., Li, R., He, C., Chen, H., Cheng, R., Zhai, C., & Zhang, M. (2022). An inverse design method for supercritical airfoil based on conditional generative models. *Chinese Journal of Aeronautics, 35*, 62–74. URL: https://www.sciencedirect.com/science/article/pii/S1000936121000662. https://doi.org/10.1016/j.cja.2021.03.006.

Wu, X., & Huang, X. (2023). Screening of urban environmental vulnerability indicators based on coefficient of variation and anti-image correlation matrix method. Ecological Indicators, 150, 110196. URL: https://www.sciencedirect.com/science/article/pii/S1470160X23003382. https://doi.org/10.1016/j.ecolind.2023.110196.

Ziemkiewicz, D. (2017). Simple analytic equation for airfoil shape description. *AIAA J. DOI, 10*(2514/1), J055986.