

A Deep-Learning Surrogate Model Approach for Optimization of Morphing Airfoils

Hamid R. Karbasian*

*The Fields Institute for Research in Mathematical Sciences
The University of Toronto, Toronto, Ontario M5T 3J1, Canada*

Wim M. van Rees†

*Department of Mechanical Engineering
Massachusetts Institute of Technology, Cambridge, MA, 02139*

Analyzing and optimizing the aerodynamic performance of a morphing airfoil concept typically requires the numerical solution of many complex, computationally expensive fluid-structure interaction (FSI) problems. This approach becomes intractable against current developments in intelligent, programmable materials and additive manufacturing techniques, which drastically increase the design space and open novel opportunities for passively and actively morphing wings. To fully exploit these capabilities, a new paradigm for analyzing and optimizing aeroelastic structures in high-dimensional parameter spaces is required. This work presents an efficient numerical design approach for elastically morphing structures in aerodynamic flows. Our approach centers on using deep neural network surrogate models to predict the aerodynamic loading as a function of a given shape. The models are trained through a set of flow simulations around rigid stationary bodies randomly sampled from a parametrized design space of the shapes. Once trained, the surrogate model can be used to evaluate the aerodynamic performance of any structural design without the need for further costly flow or FSI simulations. Consequently, this approach can analyze and optimize airfoils within a higher-dimensional structure and structure-actuator problems than currently possible. Though the approach is general, we focus here on establishing a proof-of-concept of this idea for a 2D multi-hinged airfoil at a steady-state condition. The specific contributions are validating the surrogate model, estimating the cost benefits of this approach, and providing first insights into the approach’s capabilities. A practical optimization of a 2D morphing airfoil in steady flows demonstrates that training and using the surrogate model reduces the number of required flow solutions by several orders of magnitude compared with a fully coupled FSI approach.

I. Introduction

Morphing wings offer the possibility to design lifting aerodynamic surfaces that adapt the wing shape to operating conditions for maximum steady performance and stability [1–3]. Developments in additive manufacturing enable the fabrication of such structures out of multiple materials and with fine-grained control over stiffness and thickness distributions. On top of passive shape control, programmable materials and smart shape-changing structures further offer the potential to embed actuation within the structure, so that the shape can be actively adapted to specific flow conditions and applications [2, 4, 5]. Such actuation strategies could potentially transform the design of morphing structures to achieve near-optimal performance across operating conditions, and provide avenues for continuous shape control in unsteady flows.

To exploit these capabilities, design methods are required that are able to traverse the high-dimensional structural and control spaces that govern the underlying fluid-structure interaction (FSI) problem. Currently, typical numerical approaches to FSI optimization rely on repeated evaluation of a coupled FSI numerical solver, evaluating the aerodynamic performance at each iteration. This approach has been successfully used for parametric structural designs [6] as well as topological optimization strategies [7, 8]. The cost of the optimization can be reduced by using gradient-based information and adjoint solvers [9–11], as well as surrogate models [12, 13]. Nevertheless, in all these cases the number of expensive FSI simulations required is still fundamentally related to the number of design parameters chosen

*Postdoctoral Fellow, h.karbasian@utoronto.ca

†Assistant Professor, wvanrees@mit.edu




Fig. 1 A functional diagram as described by Fung [14] for the fundamental case of a torsional-spring-supported rigid wing segment that models bending deformation of a cantilever wing (top). The functional diagram for steady-steady deformation (bottom) consists of the aerodynamic operator that takes the initial angle of attack α to compute aerodynamic loading (lift L and moment M). The structural operator takes the aerodynamic moment and updates the twisting angle Θ , which will change the angle of attack to $\alpha + \theta$, leading to a new moment. When equilibrium is obtained, the steady-state lift L is recorded.

in the structural model. As argued above, with modern material developments and manufacturing techniques, this design space can be arbitrarily large and could exceed the number of degrees-of-freedom governing the aerodynamic problem. Moreover, in the typical approach any optimal results are strongly tied to the structural parameterization, so that revisiting that parametrization often requires a repeat of all numerical simulations and discarding the earlier results.

We propose a surrogate model based approach for FSI problems that decouples the structural optimization problem from expensive FSI simulations. The surrogate model is trained to predict the aerodynamic loading over the airfoil as a function of the deformed airfoil shape. The training occurs by running a set of simulations for aerodynamic flows around rigid airfoils with parametrically varied shapes. Once trained, the surrogate model can be used to evaluate arbitrary structural designs and optimize over high-dimensional structural design and control spaces, without reverting to further expensive flow simulations. This approach directly relates to the concept of ‘functional diagrams’ sketched in Fung [14] (Chpt 11), and reproduced in Fig. 1 for a simple single-degree-of-freedom static aeroelastic problem. Referring to this functional diagram, we propose to train a surrogate model to predict the aerodynamic loading (here the lift L and twisting moment M) from the aerodynamic conditions and kinematic deformations (here angle of attack α and twisting angle Θ). For flexible airfoils that contain more degrees of freedom than sketched in Fig. 1, the same idea can be followed when Θ is replaced with a parametrized deformation field, and L and M with the aerodynamic loading contributions required to update the structural solver associated with the deformation field.

In this work we examine a proof of concept of this idea applied to a steady-state FSI problem. Our problem description is based on a 2D hinged airfoil with N segments, connected by torsion springs with controllable stiffness parameters, as well as a global torsion spring. The aerodynamic configuration is fully described with the incident angle of attack; the deformation parameters are the hinge angles and the global twisting angle; and the structural design parameters are the hinge spring constants, hinge spring rest angles, and global torsion spring constant. We use a deep neural network to train and represent the surrogate model. For the aerodynamic simulations we use XFOIL [15], which in itself is sufficiently cheap to do large-scale optimizations without a surrogate model, so that we can validate our approach. With this setup, we can assess the accuracy of the surrogate model approach compared to full flow evaluations, and quantify the number of flow evaluations required to train the surrogate model compared to the number required to perform coupled FSI optimizations.

The rest of this work presents the problem setup (II), where we mathematically formulate a description of a deformed multi-segment airfoil. Section III discusses details of the deep learning neural network architecture and how we build the surrogate model. Section IV presents the validation cases considered for finding equilibrium deformed configurations using the surrogate model (denoted the forward problem). Section V presents validation cases for finding optimal shape and structural configurations (denoted the inverse problem), where we wrap a gradient-free optimization algorithm around our approach to the forward problem. In Section VI, we demonstrate the proposed approach to a large-scale multi-objective optimization of a flexible airfoil and analyze the results, as well as the effectiveness of the surrogate

model approach to such problems. Finally, in Section VII, we summarize our thoughts and point out possible extensions of this work towards more complex problems.

II. Problem setup

In this section we discuss the aerodynamic setting and structural design parametrization.

A. Aerodynamic setting

We consider the flow past a 2D airfoil whose centerline is represented with N segments interconnected with hinges. The shape of the centerline is uniquely defined by specifying the hinge angles γ_i $1 \leq i \leq N - 1$, where γ_i represents the angle between segments i and $i + 1$, and additionally the segment lengths L_j , with $1 \leq j \leq N$. For all airfoils considered here, we set $L_j = L/N$, with L a fixed chord length. The thickness distribution of each airfoil is given by a symmetric NACA0012 profile, where we assure C^1 continuity of the airfoil outline by replacing the centerline at each hinge by a small circular arc that smoothly connects the two adjacent segments. To evaluate aerodynamic performance, the surfaces are discretized within XFOIL using $N_p = 256$ panels. We run XFOIL with chord-based Reynolds number $Re = 10^6$ and Mach number $Ma = 0.1$, and store the simulated distribution of surface pressure. This loading is integrated to compute the net aerodynamic moment coefficient acting on each hinge, and we store these moments as $C_{m,i}$ for $1 \leq i \leq N - 1$. We further store the total lift coefficient C_L , total drag coefficient C_D and moment coefficient around the quarter-chord position $C_{m,\text{ref}}$. These force and moment coefficients are used in the FSI coupling, described in the next section.

B. Structural design and FSI coupling

The elastic response of the airfoil is governed by torsion springs attached to each of the $N - 1$ interior hinges, as well as a global torsion spring attached to the quarter-chord location to mimic an overall twist of the airfoil section. In this work, we use $N \leq 4$ and so the global torsion spring is always situated on the first segment. For the hinge torsion spring on hinge i the linear elastic response follows from $M_i = k_i(\gamma_i - \gamma_{i,0})$ where M_i is the aerodynamic moment around this hinge, and the structural parameters are the linear spring constant k_i and the spring rest angle $\gamma_{i,0}$. Expressed non-dimensionally, this relationship is $C_{m,i} = \tilde{k}_i(\gamma_i - \gamma_{i,0})$ where $\tilde{k}_i = 2k_i/(\rho c^2 U_\infty^2)$ with ρ the density, c the chord length, and U_∞ the free-stream velocity. The elastic response of the global torsion spring located at the quarter-chord location is given similarly by $C_M = \tilde{K}\Theta$, with C_M the total aerodynamic moment around the quarter-chord location, \tilde{K} the non-dimensional twisting spring constant, and Θ the current twisting angle.

For a given set of structural parameters $\{\tilde{K}, \tilde{k}_1, \gamma_{1,0}, \dots, \tilde{k}_{N-1}, \gamma_{N-1,0}\}$ and an incident angle of attack α_{inc} (further clarified below) we use an iterative process to find the equilibrium configuration of the airfoil. The update from iteration n to $n + 1$ follows the gradient descent with momentum as follows

$$\begin{aligned} \delta\gamma_i^{(n+1)} &= \beta_d \delta\gamma_i^{(n)} + [C_{m,i}^{(n)} - \tilde{k}_i(\gamma_i^{(n)} - \gamma_{i,0})], \\ \gamma_i^{(n+1)} &= \gamma_i^{(n)} + \beta \delta\gamma_i^{(n+1)}, \end{aligned} \quad (1)$$

where β and β_d are the relaxation and decay factors. For all FSI problems here, we set the decay factor to $\beta_d = 0.5$, and relaxation factor to $\beta = 0.1$. Using equation (1) and a similar equation for the twisting angle Θ we can update the hinge and twisting angles, and thus update the shape for a new flow solver evaluation that leads to updated hinge and twisting moment coefficients for the next iteration. This process is repeated until we find values for $\{\Theta, \gamma_1, \dots, \gamma_{N-1}\}$ that represent an equilibrium state between the aerodynamic loading and the response of the torsion springs.

C. Definition of angle of attack

To clarify the definition of the angle of attack throughout this process, we sketch in Figure 2 configurations of undeformed and deformed airfoils with $N = 2$ segments. We define the incident angle of attack α_{inc} as the angle between the incoming flow and the first segment of the undeformed airfoil. The value of α_{inc} is kept constant throughout the FSI iterations for any given problem. The effective angle of attack α_{eff} , on the other hand, is defined as the angle between the incoming flow and the chordline. As the airfoil deforms and twists, α_{eff} necessarily changes. The precise value of α_{eff} can be computed for both the undeformed and the deformed configurations as follows. For the undeformed




Fig. 2 Comparison of an undeformed and deformed airfoil with $N = 2$ segments, and the notation of the different angles described in the text. The flow in both panels is considered from left-to-right with angle α_{inc} with respect to the horizontal.

configuration, we have $\alpha_{\text{eff},0} = \alpha_{\text{inc}} + \Gamma_0$, where (for equal-length segments)

$$\tan(\Gamma_0) = \frac{\sum_{i=1}^{N-1} \sin(\gamma_{i,0})}{1 + \sum_{i=1}^{N-1} \cos(\gamma_{i,0})}$$

is the angle between the chordline and the first segment in the undeformed configuration. The numerator and denominator in the right-hand side can be understood respectively as the y and x positions of the trailing edge relative to the leading edge. In the deformed configuration, we have $\alpha_{\text{eff}} = \alpha_{\text{inc}} + \Theta + \Gamma = \alpha_{\text{eff},0} + \Theta + (\Gamma - \Gamma_0)$, where

$$\tan(\Gamma) = \frac{\sum_{i=1}^{N-1} \sin(\gamma_i)}{1 + \sum_{i=1}^{N-1} \cos(\gamma_i)}$$

is the angle between the chordline and the first segment in the deformed configuration. It follows that during the FSI iterations, Θ and Γ change as the airfoil twists and deforms, which modifies the effective angle of attack until equilibrium is achieved.

III. Surrogate model

The development of the surrogate model in this study is based on a deep neural network. Deep learning (DL) aims to capture non-linear manifolds in a dataset and uses those manifolds as essential features in the surrogate model. Here we leverage DL to find patterns in the dataset that relate geometric parameters to aerodynamic loads for each segment of the 2D airfoil. We use a Multi-Layer Perceptron (MLP) architecture to be able to import a set of input parameters into the DL model and receive multivariant output parameters. The input parameters that form the design space \mathcal{D} in our case are the effective angle of attack, α_{eff} , and the $N - 1$ hinge angles, γ_i . The model returns the global force and moment coefficients C_L , C_D , C_M , as well as the $N - 1$ moment coefficients around each hinge, $C_{m,i}$, which are in the physical space. Moreover, we added the inverse of drag coefficient, C_D^{-1} , to reduce the sensitivity of the output of the DL model in case the ratio between lift and drag forces (C_L/C_D) is required as design parameter.

For an airfoil with N segments, the input layer h_{in} thus has N neurons: one for α_{eff} and $N - 1$ for the hinge angles γ_i (see Fig. 3 for an example with $N = 3$). The output layer h_{out} has $N + 3$ neurons, associated with the $N - 1$ hinge torques $C_{m,i}$ and the four overall aerodynamic coefficients associated with lift, drag, inverse drag, and moment. Here we use a neural network architecture that connects the input and output layers with two hidden layers, h_1 and h_2 , of 20 neurons each, and a ‘‘tanh’’ activation function σ . We can then write

$$\begin{aligned} h_1 &= \sigma(\mathcal{W}_1 h_{\text{in}} + b_1), \\ h_2 &= \sigma(\mathcal{W}_2 h_1 + b_2), \\ h_{\text{out}} &= \sigma(\mathcal{W}_3 h_2 + b_3), \end{aligned} \quad (2)$$

where \mathcal{W}_i and b_i (with $1 \leq i \leq 3$) are the unknown weight matrix and bias, respectively. For each number of segments N , we find \mathcal{W}_i and b_i by optimizing the network against a training data set. To do so, we use Adam with a learning rate of 10^{-3} , a batch size of 64, and an epoch of 10^3 . Further, we enrich our training data set based on a ‘‘jittering’’ method to

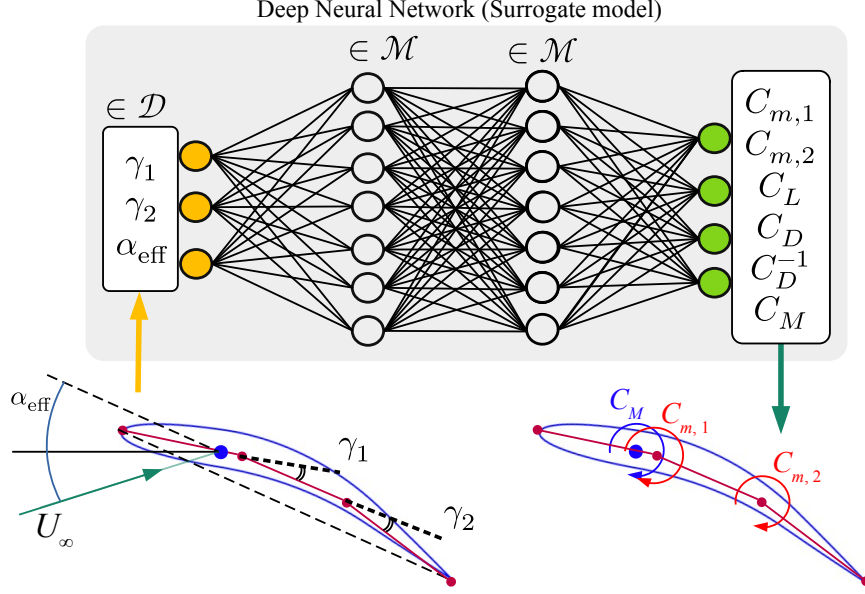


Fig. 3 Schematic of the deep neural network underlying the surrogate model, here for an airfoil with $N = 3$ segments.

synthesize new samples [16]. In practice, this approach leads to have a stable and less sensitive DL model that can learn manifolds \mathcal{M} in the design space \mathcal{D} . Moreover, having such less sensitive DL model will fit better with FSI solver, leading to a faster convergence.

The root mean-square error (RMSE) is used to measure the model's accuracy. The RMSE is minimized with respect to a dataset with N_{samp} samples. We can define this error as follows

$$\varepsilon_{\text{RMSE}} = \sqrt{\frac{1}{N_{\text{samp}}} \sum_{n=1}^{N_{\text{samp}}} (Y_{\text{pred}}^{(n)} - Y_{\text{samp}}^{(n)})^2}, \quad (3)$$

where $Y_{\text{pred}}^{(n)}$ is the prediction from the output layer of the DL model on the n th data point, and $Y_{\text{samp}}^{(n)}$ is the associated entry in the dataset. In our case, the dataset is formed by a random sampling of the design parameters in the space \mathcal{D} , running XFOIL simulations for each of the samples, and post-processing the XFOIL simulation to find the force and moment coefficients. In some cases, XFOIL fails to converge, for instance when the inflow or hinge angles are very large. When this happens, we discard the associated points from the dataset. Throughout our validation and optimization studies, we are not interested in these regions so discarding these data points is inconsequential for the results presented here.

Ultimately, a main concern in the approach is the amount of data samples required to train the network effectively. This is especially relevant as we intend to replace XFOIL with high-fidelity data from CFD solvers, which will be much more expensive to generate. To examine this in more detail, we consider an airfoil with $N = 2$ segments and one hinge angle as an example. We generate two datasets, one with 200 samples and one with 400 samples, where each sample is drawn from a multi-dimensional uniform distribution within the design space. We train two DL models with identical architecture using 80% of each dataset. The other 20% of the dataset is set aside as testing data. Figure 4 visualizes the design space built for three aerodynamic load coefficients as a function of the incident and hinge angles. The black mesh interpolates the random samples (black dots) selected as training data within the design space \mathcal{D} , and the red dots show the predicted values using the trained DL models evaluated at the testing points. The root-mean-square error between the exact and predicted values of test data $\varepsilon_{\text{RMSE}}$ for the 200- and 400-sample DL models are 3.58×10^{-4} and 1.60×10^{-4} , respectively. Although a reduction in the $\varepsilon_{\text{RMSE}}$ in the 400-sample DL model is observed, the aerodynamic coefficients behave sufficiently smooth that the 200-sample model already provides sufficient accuracy needed for the analysis. After further analysis of airfoils with different number of hinges, we proceeded with using $100 \times N$ random samples for an airfoil with N segments, where $N \leq 4$ in this work.

We note here that, as the focus in this work is on demonstration and validation of the proposed approach, and since XFOIL samples are relatively inexpensive to generate, we did not improve further upon the specific training strategy

and network architecture described. When we extend the approach to more expensive, high-fidelity CFD simulations, the efficient sampling and use of training data becomes a crucial part of the problem, and we discuss this further in section VII.




Fig. 4 Model evaluation for two different datasets for the DL model with 200 samples (a–c), and for the DL model with 400 samples (d–f). The figures show the lift (a,d), drag (b,e), and twisting moment (c,f) coefficients, plotted as a function of the instantaneous hinge angle γ_1 and the incident angle α_{inc} . In each plot the black mesh interpolates the random samples (black dots) selected as training data within the design space, and the red dots show the predicted values using the trained DL models evaluated at the testing points.

IV. Validation of forward problem

We define the forward problem as finding the equilibrium configuration (defined through the twisting and hinge angles) of an airfoil at a fixed incident angle α_{inc} and for a given set of spring constants. The equilibrium configuration is found through the approach detailed in section II.B, and we compare the result obtained between using XFOIL vs the DL surrogate model to evaluate the aerodynamic load. Below we detail the results of this comparison for three different cases of increasing complexity.

A. Case 1: single hinge, zero angle of attack and hinge rest angle

In the first case, we consider an airfoil with two segments and thus a single hinge. We set the the incident angle $\alpha_{inc} = 0^\circ$ and the rest angle of the hinge $\gamma_{1,0} = 0^\circ$. In this case, because the flow and structure are symmetric, the FSI problem should converge to a straight shape with $\Theta = \gamma_1 = 0^\circ$. To test our FSI procedure, we initialize the iterations by setting $\Theta^{(0)} = 5^\circ$ and $\gamma_1^{(0)} = 5^\circ$, respectively, and use torsional-spring stiffnesses of $\tilde{K} = \tilde{k}_1 = 1$. Figure 5a shows the initial shape (in grey) with specified initial twisting and hinge angles, and the equilibrium shapes predicted by DL (red) and XFOIL (blue), both corresponding to a symmetric airfoil. In Figure 5b, we show the error evolution of the FSI iterations when using the DL model. Here, and in the following two cases, the error is defined as the difference between the predicted angles using the DL model and the converged value obtained from the FSI iterations using the XFOIL model. We observe that the DL model converges within about 40 iterations to angles that match well with the XFOIL values.

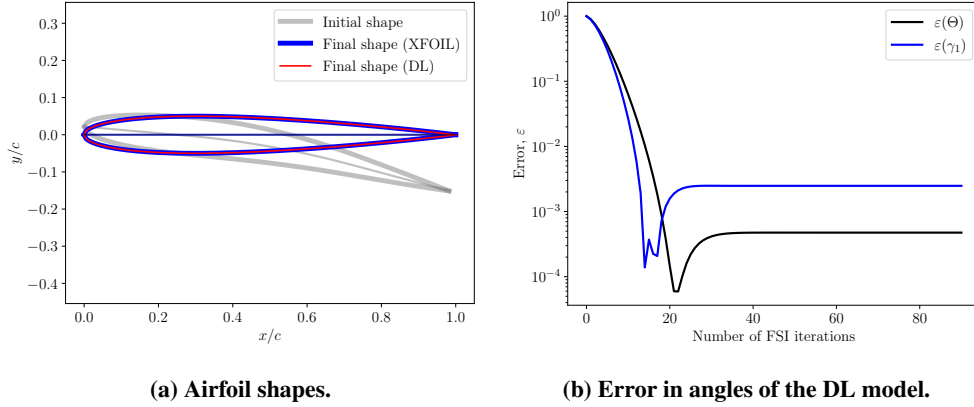


Fig. 5 Comparison of FSI solver results using XFOIL and the DL model for validation Case 1.

B. Case 2: single hinge, non-zero angle of attack and rest angle

We increase the complexity over the previous case by changing the incident angle of attack to $\alpha_{\text{inc}} = 10^\circ$ and the rest angle of the hinge to $\gamma_{1,0} = -5^\circ$, without changing the other parameters. We initialize the FSI iterations with $\Theta^{(0)} = \gamma_1^{(0)} = 0^\circ$. The torsional-spring stiffnesses are set to $\tilde{K} = \tilde{k}_1 = 0.25$. Figure 6a shows the initial shape of the airfoil and the final shapes obtained by XFOIL and DL models. The flexible airfoil deforms due to the aerodynamic moment around the hinges, and subsequently, it changes the twisting angle in such a way that the moments are in equilibrium. The final shape obtained by the DL model is in good agreement with the one obtained by XFOIL. In Figure 6b, we show the error evolution of the FSI iterations when using the DL model, showing a convergence up to roughly 1% in the hinge angle γ_1 .




Fig. 6 Comparison of FSI solver results using XFOIL and the DL model for validation Case 2.

C. Case 3: multiple hinges, non-zero angle of attack and rest angles

To extend our validation to flexible airfoils with $N = 3$ and $N = 4$ segments, we first build and train DL models for each of these segmented airfoils in the same way as for the $N = 2$ segments airfoil discussed so-far. We then validate the FSI equilibrium iterations for a configuration with $\alpha_{\text{inc}} = 5^\circ$, and rest angles of $\gamma_{i,0} = -10^\circ$ for all hinges. Moreover, the torsional-spring stiffness for the global twist and for all hinges are set to $\tilde{K} = 0.025$ and $\tilde{k}_i = 0.5$. We initialize the hinge angles to the rest angles and run the algorithm described in section II.B until equilibrium, both with XFOIL and the DL models to evaluate aerodynamic loading for a given shape. Figures 7a and 7c show that the DL models with both three and four segments are able to find sufficiently similar equilibrium shapes of the flexible airfoil compared to those obtained by XFOIL. Moreover, Figures 7b and 7d display the error for both DL models are approximately $O(10^{-3})$, indicating the DL models perform very well on these FSI test cases.

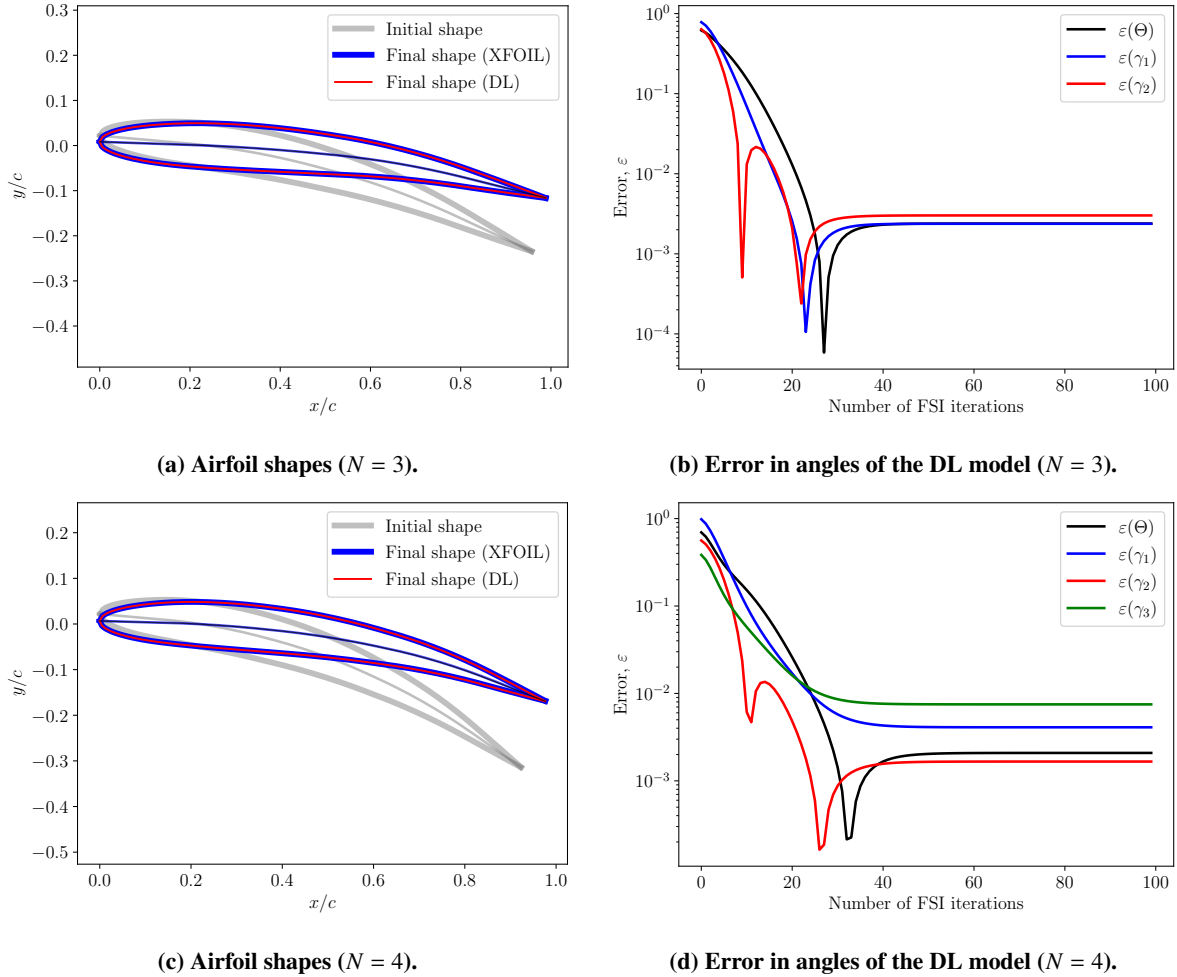


Fig. 7 Comparison of FSI solver results using XFOIL and the DL model for validation Case 3 using $N = 3$ segments (a,b) and $N = 4$ segments (c,d).

Summarizing all three cases, we conclude that the forward problem of finding an equilibrium configuration for a given structural design using the approach outlined in section II.B is reliable and robust. Further, the developed deep learning models provide an adequate surrogate approach to XFOIL for retrieving the aerodynamic loading for a given shape and incident angle.

V. Validation of the inverse problem

One of the main goals of developing the surrogate model described here is to accelerate the FSI design process. In this section we couple the forward problem described in section II.B with a gradient-free optimization framework to find optimal structural parameters of a morphing airfoil. Specifically, we perform validation studies on the optimization of both airfoil shapes and airfoil structures. For shape optimization we essentially set the spring stiffnesses to infinity, thereby ensuring that no deformation takes place; we then only optimize over the angles that immediately define the shape. For structural optimization, we optimize over the rest angles and the (finite) spring stiffnesses. In this case, for each set of structural parameters we need to find the equilibrium configuration using the procedure of section II.B before being able to evaluate the aerodynamic performance of the airfoil. Since the objective function is unchanged, we expect the structural optimization to converge to a configuration whose equilibrium deformed shape is the same as the result from the direct shape optimization.

Throughout this work, we use the Mesh Adaptive Direct Search (MADS) algorithm as optimization algorithm

because of its suitability for non-smooth and complex design spaces [17, 18]. MADS is a new version of the Generalized Pattern Search (GPS) class of algorithms [19] with features of the Coope and Price frame-based methods [20]. A brief description of MADS is provided in the Appendix, further details can be found in literature [17, 21].

A. Shape optimization

For the shape optimization we consider a rigid airfoil with $N = 2$ segments. The optimization parameters are the instantaneous incident angle α_{inc} and the instantaneous hinge angle γ_1 , with both parameters initialized to 0° . The bounds for the incident and hinge angles are set to $\alpha_{inc} = [-10^\circ, 10^\circ]$ and $\gamma_1 = [-20^\circ, 10^\circ]$. The objective function \mathcal{J} is defined as a non-dimensional function that minimizes drag at a certain target lift coefficient:

$$\mathcal{J} = \left(1.0 - \frac{C_L}{C_{L,target}}\right)^2 + \lambda \left(\frac{C_D}{C_{L,target}}\right), \quad C_{L,target} = 0.8, \quad (4)$$

where $C_{L,target}$ is the target lift coefficient. The scalar constant λ is tuned to appropriately weigh the two terms in the objective function; based on initial evaluations we choose here $\lambda = 20$. For the MADS optimizer, we scale $\Delta_m^{(0)}$ and $\Delta_p^{(0)}$ by a factor of 5, and set the minimum tolerance for $\Delta_p = 10^{-3}$ with maximum design iteration of 100.

For illustration, since both the XFOIL and DL function evaluations are computationally cheap, we compute the objective function within our entire parametric domain using both approaches. The contours of the objective function in both cases (Fig. 8) match well, and the optimizer follows a similar trajectory in both cases.

The optimized parameter values are shown in the first two rows of Table 1, with ‘RIG-XF’ denoting the results found using XFOIL as flow solver, and ‘RIG-DL’ the results found using the Deep Learning surrogate model instead of XFOIL. Since this is a rigid airfoil, we have $\Theta = 0$ and $\gamma_1 = \gamma_{1,0}$. We observe that the surrogate-model based optimization finds an optimum very similar to that of the XFOIL-based optimization, both in terms of the final angles as well as the predicted lift, drag, and objective function value. The three right-most columns indicate that the number of optimization iterations and number of function evaluations are similar between the two models. These values depend on the optimization trajectory, which itself is partially driven by a stochastic process so that a precise comparison is not informative. We do note that the cost per evaluation of the DL is about 1.5–3.5 times faster than for XFOIL on a standard laptop, so even though XFOIL is already a low-order model the DL surrogate model still offers a moderate improvement in computational cost. Once the approach is scaled to a Navier-Stokes solver instead of XFOIL, this improvement will increase by several orders of magnitude.

Case	\tilde{K}	\tilde{k}_1	α_{inc}	$\alpha_{inc} + \Theta$	$\gamma_{1,0}$	γ_1	C_L	$C_D \cdot 10^2$	\mathcal{J}	iters	fevals	time
RIG-XF	∞	∞	-3.1°	-3.1°	-11.9°	-11.9°	0.78	0.54	0.136	36	138	0.57
RIG-DL	∞	∞	-2.7°	-2.7°	-11.1°	-11.1°	0.78	0.57	0.144	52	220	0.61
FSI-1-XF	0.025	1.5	1.1°	-3.2°	-13.9°	-12.0°	0.78	0.54	0.136	26	1849	6.4
FSI-1-DL	0.025	1.5	1.3°	-2.7°	-12.9°	-11.1°	0.78	0.57	0.144	32	1882	2.31
FSI-2-XF	0.025	0.15	3.1°	-3.4°	-33.5°	-13.2°	0.77	0.53	0.135	49	7250	25.8
FSI-2-DL	0.025	0.15	1.3°	-2.7°	-29.0°	-11.4°	0.78	0.57	0.144	60	4234	4.3

Table 1 Optimization results for different cases discussed in section V. The last three columns indicate the number of iterations in the optimization loop (iters), the number of function evaluations (fevals, either XFOIL for ‘XF’ or the DL surrogate model for ‘DL’), and the total time (in minutes) taken during optimization running on a standard laptop.

B. Structural optimization

In this section, we investigate the structural optimization of a deformed airfoil using DL surrogate model. Specifically, we optimize the rest angle of the hinge $\gamma_{1,0}$ and the incident angle of attack α_{inc} , while fixing the global and hinge spring stiffnesses \tilde{K} and \tilde{k} . We set $\tilde{K} = 0.025$ and examine the optimization for two different values of \tilde{k} : the case denoted with ‘FSI-1’ has $\tilde{k} = 1.5$ and the case denoted with ‘FSI-2’ has $\tilde{k} = 0.15$.




Fig. 8 Contour map of the cost function \mathcal{J} as function of the incident angle α_{inc} and the instantaneous hinge angle γ_1 for a rigid (undeformed) airfoil, evaluated using XFOIL (left) and the DL model (right). The black lines in each contour plot show the optimization trajectory.

The evolution of the lift and drag coefficients as well as the objective function \mathcal{J} during the course of the optimization are shown in Fig. 9. The optimization algorithm requires more design iterations to convergence for the more flexible airfoils (panel (b)) compared to the more rigid airfoils (panel (a)), presumably due to a higher sensitivity between the input parameters and the objective function. Nevertheless, in all cases the optimizer finds structural parameters that yield lift coefficients close to the target lift coefficient of 0.8. The optimal parameters for α_{inc} and $\gamma_{1,0}$ across the FSI-1 and FSI-2 scenarios are reported in the last four rows of Table 1. We observe that the incident angles in both cases FSI-1 and FSI-2 are optimized to values so that, after accounting for the twisting angle Θ , the sum $\alpha_{\text{inc}} + \Theta$ closely matches the rigid airfoil incident angle found in the RIG-XF and RIG-DL cases. Similarly, the rest hinge angles $\gamma_{1,0}$ are chosen so that, after accounting for the deformation, the deformed equilibrium airfoils have hinge angles γ_1 that closely match those of the rigid airfoil. This is expected, because this means that both FSI cases deform into the rigid shape that was found to be optimal in the RIG step. To see this, Figure 10 compares the final shapes of the flexible airfoils with those obtained using XFOIL, and further plots the initial undeformed shapes from FSI-1-DL and FSI-2-DL in parts (a) and (b), respectively. We see that the initial, undeformed shapes have significant camber, which is more pronounced Figure 10b due to the lower torsional stiffness value in FSI-2. When loaded, the final deformed shapes obtained by the DL and XFOIL models are indeed practically indistinguishable and match well with the optimal rigid shape (not shown in the Figure). Consequently, the final lift and drag coefficients and the objective function \mathcal{J} are practically identical across all cases in Table 1.

The small discrepancy between the optimal design parameters between the XFOIL and DL based approaches can be analyzed further. The optimal values for C_L and C_D reported in Table 1 are evaluated using the XFOIL model for the XF cases, and using the DL model for the DL cases. However, when taking the optimal design parameters of the DL cases, and using XFOIL to evaluate the lift and drag coefficients associated with those design points, we find that the DL model errors largely account for the discrepancies. For instance, for the optimal FSI-1-DL parameters the DL model finds $C_L = 0.78$ and $C_D = 0.57 \times 10^{-2}$, whereas for those same parameters XFOIL finds $C_L = 0.75$ and $C_D = 0.53 \times 10^{-2}$. The DL objective function is $\mathcal{J} = 0.144$ whereas using XFOIL to evaluate lift and drag would give an objective function of $\mathcal{J} = 0.137$ for the same parameters. This value of the objective function is very close to the objective function value associated with the optimum design parameters for the XF-based cases in Table 1, implying that the small differences in parameter values (α_{inc} and $\gamma_{1,0}$) between the DL-based and XFOIL-based inverse design approaches are practically inconsequential for the final airfoil performance. We thus conclude that, for the approach considered here, the DL-based optimization approach leads to structural design parameters that are practically similar to those found through the XFOIL-based optimization approach, but that the DL-based prediction of the associated lift and drag coefficients is off by about 5%.




Fig. 9 Evolution of aerodynamic loads and objective function as a function of the design iteration, for the FSI-1 (a) and FSI-2 (b) inverse structural validation problems.

VI. Application to flexible airfoil optimization

Here we wish to demonstrate the utility of the DL-based surrogate model developed and validated above for a complex aerodynamic optimization problem, without having to perform additional expensive flow simulations. The optimization problem is chosen to identify the structural design of a $N = 4$ segment airfoil that shows, across four different incident angles, the smallest deviation of the lift coefficient from a target value, while also minimizing overall drag. Mathematically, we pose this problem using a scalarized multi-objective function:

$$\underset{\tilde{k}_i, \gamma_{i,0} \in \mathcal{D}}{\text{minimize}} \quad \mathcal{J} = \sum_{k=1}^4 \left(1.0 - \frac{C_L(\alpha_{\text{inc}}^{(k)})}{C_{L,\text{target}}} \right)^2 + \lambda \frac{C_D(\alpha_{\text{inc}}^{(k)})}{C_{L,\text{target}}}, \quad (5)$$

where the target lift coefficient is $C_{L,\text{target}} = 0.7$, the weight $\lambda = 10$, and the four angles over which we measure the lift and drag are $(\alpha_{\text{inc}}^{(1)}, \alpha_{\text{inc}}^{(2)}, \alpha_{\text{inc}}^{(3)}, \alpha_{\text{inc}}^{(4)}) = (0^\circ, 3^\circ, 6^\circ, 9^\circ)$. For simplicity, we remove the global twisting deformation from the parameters ($\tilde{K} \rightarrow \infty$) so that $\Theta = 0$ throughout. The optimization then considers the optimization of the three rest hinge angles $\gamma_{i,0}$, which we bound to the region $-40^\circ \leq \gamma_{i,0} \leq 10^\circ$, and the three torsional spring stiffnesses \tilde{k}_i bound to the region $0.1 \leq \tilde{k}_i \leq 2$. The maximum number of design iterations for the MADS optimizer is 500, and the minimum Δ_p is set to 10^{-4} .

Figure 11 shows the results for the multi-objective FSI optimization. In Figure 11a, the objective function reduces significantly at the initial steps of the optimization as it finds the route to the global optima, and later on, it changes slightly to tune weighting between different objective functions. In this case, the optimizer terminates after 74 design iterations, with a total number of DL evaluations of fevals = 42, 000, lasting about 38 mins on a standard laptop. In Figure 11b, the drag coefficients across the different incident angles increase early on as the terms involving the lift coefficient carry a relative higher weight in the objective function. After about 30 iterations the drag coefficients stabilize and only minor alterations to the aerodynamic coefficients and objective function are observed. Ultimately the objective function converges to $\mathcal{J} = 0.647$, associated with optimal rest hinge angles of $(\gamma_{1,0}, \gamma_{2,0}, \gamma_{3,0}) = (-2.40^\circ, -2.44^\circ, -11.57^\circ)$ and optimal non-dimensional hinge stiffnesses of $(\tilde{k}_1, \tilde{k}_2, \tilde{k}_3) = (0.10, 1.04, 1.84)$.

For reference, we repeated the optimization for a rigid airfoil with the same objective function as Eq.(5). In this case we set all torsional-spring stiffnesses $\tilde{k}_i \rightarrow \infty$ to have a rigid airfoil, and optimize only over the three instantaneous hinge angles γ_i , using the same bounds as for the rest hinge angles above. Since this optimization does not need any FSI iterations for each objective function evaluation, it only uses fevals = 280 DL evaluations for 69 optimization design iterations, taking slightly under three minutes to complete. The final objective function value is $\mathcal{J} = 1.036$, indicating that exploiting airfoil flexibility can significantly improve performance on this objective function.

The difference between the optimal flexible and optimal rigid airfoil is clearly demonstrated when plotting the lift and drag coefficients as a function of incident angle (Figure 12). For reference, we also include the results from an undeformed, rigid NACA 0012 airfoil. We observe that the flexible airfoil has a much reduced sensitivity to incident




Fig. 10 The initial (undeformed) and final (deformed) shapes of the airfoil for the cases FSI-1-DL and FSI-2-DL, and corresponding XFOIL-based shapes, all at their optimal design parameters.

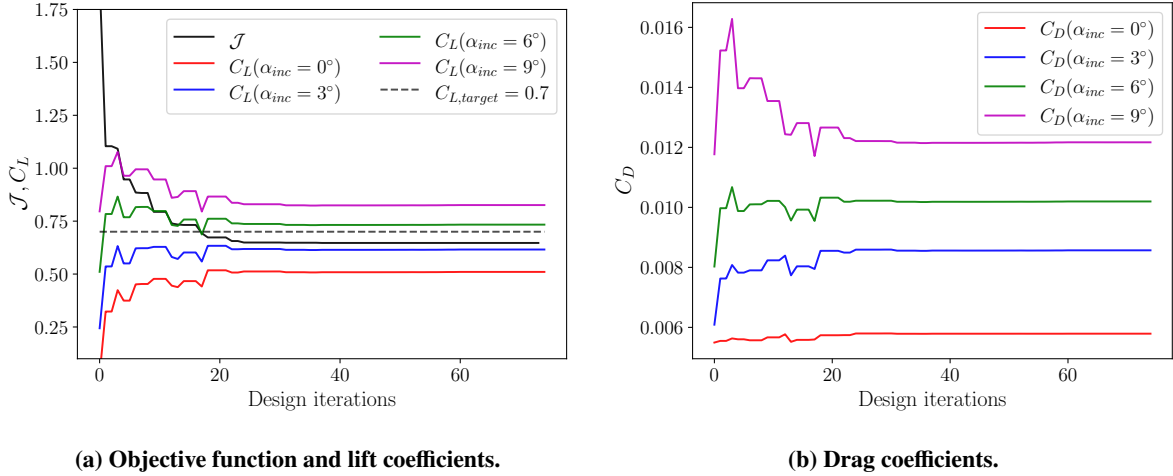


Fig. 11 The optimization progress for the multi-objective FSI optimization of section VI.

angle compared to both rigid airfoils, and a smaller drag throughout compared with the optimal rigid airfoil. The zero-lift angle for NACA 0012 is $\alpha_{inc}^{(0)} = 0^\circ$, while these angles for the optimized rigid and flexible airfoils are approximately $\alpha_{inc}^{(0)} = -3^\circ$ and $\alpha_{inc}^{(0)} = -8^\circ$, respectively.

Figure 13 shows the shape of the optimized flexible airfoil before deformation (left panel, in gray) and after deformation (in red) at incident angles of 0° , 6° , and 9° . Furthermore, the shape of the optimized rigid airfoil (in blue) is added for comparison. As shown, the optimized flexible and rigid airfoils have a similar shape at $\alpha_{inc} = 6^\circ$, where the lift curves coincide around $C_{L,target} = 0.7$. However, in two other incident angles, we see considerable differences in the deformed shape of the optimized flexible and optimized rigid airfoils that improve this cost function for the flexible airfoil.

VII. Conclusion

In this work, we demonstrate a surrogate-model approach for analyzing and optimizing complex FSI problems for morphing airfoils under steady-state flow conditions. After training the surrogate model to predict aerodynamic loading as a function of a given deformed shape and incident angle, the model can be used to compute equilibrium shapes and




Fig. 12 The lift and drag coefficient curves for the NACA 0012, optimized rigid, and optimized flexible airfoils.




Fig. 13 Shape of the optimized flexible airfoil before and after deformation at different incident angles. Note that the frame of reference is chosen so that in each panel the inflow is aligned with the x -axis.

associated aerodynamic performance for any given structural design parameters. Consequently, large-scale optimizations within high-dimensional design spaces become feasible without additional overhead of expensive coupled flow-structure simulations. Our final demonstration showed that using a surrogate model trained with only 400 flow simulations around rigid, specified shapes, one can effectively tackle complex, multi-objective structural optimization problems that would otherwise require orders of magnitude more flow simulations. The results implied that the optimized flexible airfoil significantly improved in the form of lower drag, less sensitivity of the lift curve to the incident angles, and delays in stall compared to the optimized rigid airfoil.

Several immediate extensions of this approach can be envisioned. First, the extensions to three-dimensional problems and continuously flexible airfoils are of interest. In both cases, the current segmented airfoil approach would be replaced by more complex parametrizations of the shape. Future work in this direction can directly build upon airfoil shape optimization techniques, e.g. [22]. On the other side, the surrogate model predictions would now involve continuous spatially varying aerodynamic load distributions. To reduce the dimensionality of that space, a projection onto a suitable basis would be desired, with a basis choice that can effectively retain the most important dimensions of these distributions.

Second, the extension of this technique to transient problems is of significant interest to incorporate the stability of flexible airfoils in steady flows, as well as their response to unsteady flow conditions into the optimization process. This is significantly harder to the point that generic parametric approach are unlikely to be realistic [23]. Instead, approaches for different types of instabilities and classes of structures would need to be devised, for example along the lines of the approach presented in [24].

Finally, handling more complex problems translates directly into increasing the dimensionality of the input and

output parameters of the surrogate model, which calls for more data and more complex network architectures. The latter would require further work into the DL network architecture, and what the most beneficial alterations are when the dimensionality of the input and output parameters governing this problem increase. However, given that ultimately the surrogate model aims to replace expensive CFD simulations, generating sufficient data points will be the primary bottleneck for extending the proposed approach towards more complex problems. Though this trend is ultimately inescapable, one can likely significantly improve upon our presented approach for generating and sampling training data, by choosing training points through more careful sampling techniques [25], exploiting multi-fidelity flow simulations [26], and using active learning approaches [27, 28]. Lastly, one can consider numerical approaches that reduce the cost of CFD simulations themselves, such as higher-order boundary discretizations [29] or multiresolution mesh adaptation techniques [30], so that more training data can be generated within a given computational budget.

Appendix

MADS is a new version of the Generalized Pattern Search (GPS) class of algorithms [19] with features of the Coope and Price frame-based methods [20]. At each design iteration, MADS generates trial points with different design parameters, then tries to find the best performance among the evaluated trial points. All these trial points lie on a mesh constructed by a finite set of n_D directions, defined by the mesh size parameter, $\Delta_m^{(j)}$. Also, we define \mathbf{D} as a matrix of directions that is a positive spanning set, and a non-negative integer combination of the directions [21, 31]. The mesh points $\mathbf{M}^{(j)}$ are defined by

$$\mathbf{M}^{(j)} = \bigcup_{\mathcal{S}^{(j)} \in \mathcal{D}} \left\{ \mathcal{S}^{(j)} + \Delta_m^{(j)} \mathbf{D}z : z \in \mathbb{N}^{n_D} \right\}, \quad (6)$$

where $\mathcal{S}^{(j)}$ is the current optimal point at j^{th} design iteration. The trial points $\mathbf{P}^{(j)}$ in the design space \mathcal{D} are also defined by

$$\mathbf{P}^{(j)} = \left\{ \mathcal{S}^{(j)} + \Delta_m^{(j)} d : d \in \mathbf{D}^{(j)} \right\} \subset \mathbf{M}^{(j)}. \quad (7)$$

Figure 14 shows the MADS two strategic steps in the optimization process. These two steps are called the search and

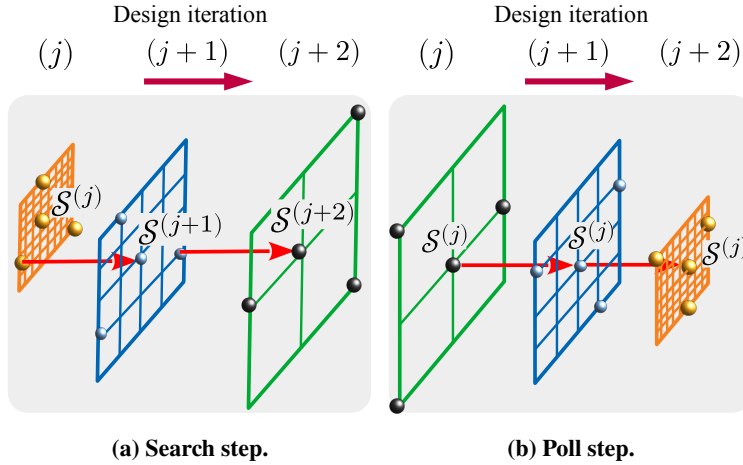


Fig. 14 Schematic of the search and poll steps in the MADS algorithm.

poll steps. In the search step, after function evaluation at trial points $\mathbf{P}^{(j)}$, if there is a new optimum among the case considered in $\mathbf{P}^{(j)}$, then the MADS increases the mesh and poll size parameters (Figure 14a). However, if no optimum is found among the trial points in the poll step, then the MADS reduces the mesh and poll size parameters simultaneously to reduce the searching zone (Figure 14b). In this study, we set the maximum frame parameter at j^{th} design iteration to $\Delta_m^{(j)} = \Delta_p^{(j)} = 1$. Then, the possible mesh size parameter for the next design iteration, $j + 1$, is defined by

$$\Delta_m^{(j+1)} = \begin{cases} \frac{\Delta_m^{(j)}}{4} & \text{for unsuccessful step at } j^{\text{th}}, \\ 4\Delta_m^{(j)} & \text{for successful step at } j^{\text{th}}, \\ \Delta_m^{(j)} & \text{for successful step at } j^{\text{th}} \text{ with } \Delta_m^{(j)} = 1, \end{cases} \quad (8)$$

and the poll size parameter is

$$\Delta_p^{(j)} = \begin{cases} n_s \sqrt{\Delta_m^{(j)}} \geq \Delta_m^{(j)} & \text{if minimal mode,} \\ \sqrt{\Delta_m^{(j)}} \geq \Delta_m^{(j)} & \text{if maximal mode.} \end{cases} \quad (9)$$

where n_s is the number of design parameters in design space \mathcal{D} . The maximal and minimal modes represent the $2n_s$ and $n_s + 1$ number of trial points in \mathbf{P} selected at each optimization step.

Acknowledgments

WMvR would like to thank Christian Viteri and Brady Sullivan for their exploratory research in the context of an Undergraduate Research Opportunities Program (UROP) at MIT. Further, WMvR would like to thank Dr Gregg Abate from the Airforce Office of Scientific Research (AFOSR) and Drs Philip Beran, Nathan Wukie, and Alexander Pankonien from the Airforce Research Laboratory (AFRL) for kindly sharing their thoughts on an early iteration of this topic.

References

- [1] Thill, C., Etches, J., Bond, I., Potter, K., and Weaver, P., “Morphing skins,” *The Aeronautical Journal*, Vol. 112, No. 1129, 2008, pp. 117–139. <https://doi.org/10.1017/s000192400002062>.
- [2] Goh, G. D., Agarwala, S., Goh, G. L., Dikshit, V., Sing, S. L., and Yeong, W. Y., “Additive manufacturing in unmanned aerial vehicles (UAVs): Challenges and potential,” *Aerospace Science and Technology*, Vol. 63, 2017, pp. 140–151. <https://doi.org/10.1016/j.ast.2016.12.019>.
- [3] Li, D., Zhao, S., Ronch, A. D., Xiang, J., Drofelnik, J., Li, Y., Zhang, L., Wu, Y., Kintscher, M., Monner, H. P., Rudenko, A., Guo, S., Yin, W., Kirn, J., Storm, S., and Breuker, R. D., “A review of modelling and analysis of morphing wings,” *Progress in Aerospace Sciences*, Vol. 100, 2018, pp. 46–62. <https://doi.org/10.1016/j.paerosci.2018.06.002>.
- [4] van Rees, W. M., Vouga, E., and Mahadevan, L., “Growth patterns for shape-shifting elastic bilayers,” *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 114, No. 44, 2017, pp. 11597–11602. <https://doi.org/10.1073/pnas.1709025114>.
- [5] Boley, J. W., van Rees, W. M., Lissandrello, C., Horenstein, M. N., Truby, R. L., Kotikian, A., Lewis, J. A., and Mahadevan, L., “Shape-shifting structured lattices via multimaterial 4D printing,” *Proceedings of the National Academy of Sciences*, Vol. 116, No. 42, 2019, pp. 20856–20862. <https://doi.org/10.1073/pnas.1908806116>.
- [6] Lian, Y., Shyy, W., Viieru, D., and Zhang, B., “Membrane wing aerodynamics for micro air vehicles,” *Progress in Aerospace Sciences*, Vol. 39, No. 6-7, 2003, pp. 425–465. [https://doi.org/10.1016/s0376-0421\(03\)00076-9](https://doi.org/10.1016/s0376-0421(03)00076-9).
- [7] Maute, K., and Allen, M., “Conceptual design of aeroelastic structures by topology optimization,” *Structural and Multidisciplinary Optimization*, Vol. 27, No. 1-2, 2004, pp. 27–42. <https://doi.org/10.1007/s00158-003-0362-z>.
- [8] Stanford, B., and Ifju, P., “Aeroelastic topology optimization of membrane structures for micro air vehicles,” *Structural and Multidisciplinary Optimization*, Vol. 38, No. 3, 2008, pp. 301–316. <https://doi.org/10.1007/s00158-008-0292-x>.
- [9] Lund, E., Moller, H., and Jakobsen, L., “Shape design optimization of steady fluid-structure interaction problems with large displacements,” *19th AIAA Applied Aerodynamics Conference*, American Institute of Aeronautics and Astronautics, 2001. <https://doi.org/10.2514/6.2001-1624>.
- [10] Jenkins, N., and Maute, K., “An immersed boundary approach for shape and topology optimization of stationary fluid-structure interaction problems,” *Structural and Multidisciplinary Optimization*, Vol. 54, No. 5, 2016, pp. 1191–1208. <https://doi.org/10.1007/s00158-016-1467-5>.
- [11] Sanchez, R., Albring, T., Palacios, R., Gauger, N. R., Economon, T. D., and Alonso, J. J., “Coupled adjoint-based sensitivities in large-displacement fluid-structure interaction using algorithmic differentiation,” *International Journal for Numerical Methods in Engineering*, Vol. 113, No. 7, 2017, pp. 1081–1107. <https://doi.org/10.1002/nme.5700>.
- [12] Han, Z.-H., Zhang, K.-S., Song, W.-P., and Qiao, Z.-D., “Optimization of Active Flow Control over an Airfoil Using a Surrogate-Management Framework,” *Journal of Aircraft*, Vol. 47, No. 2, 2010, pp. 603–612. <https://doi.org/10.2514/1.45899>.

- [13] Berci, M., Gaskell, P. H., Hewson, R. W., and Toropov, V. V., “Multifidelity metamodel building as a route to aeroelastic optimization of flexible wings,” *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, Vol. 225, No. 9, 2011, pp. 2115–2137. <https://doi.org/10.1177/0954406211403549>.
- [14] Fung, Y. C., *An introduction to the theory of aeroelasticity*, Dover Publications Inc, 1993.
- [15] Drela, M., “XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils,” *Lecture Notes in Engineering*, Springer Berlin Heidelberg, 1989, pp. 1–12. https://doi.org/10.1007/978-3-642-84010-4_1.
- [16] Iwana, B. K., and Uchida, S., “An empirical survey of data augmentation for time series classification with neural networks,” *Plos one*, Vol. 16, No. 7, 2021, p. e0254841.
- [17] Karbasian, H. R., and Vermeire, B. C., “Gradient-free aerodynamic shape optimization using large eddy simulation,” *Computers & Fluids*, Vol. 232, 2022, p. 105185.
- [18] Aubry, A., Karbasian, H. R., and Vermeire, B. C., “High-fidelity gradient-free optimization of low-pressure turbine cascades,” *Computers & Fluids*, Vol. 248, 2022, p. 105668.
- [19] Booker, A. J., Jr., J. E. D., Frank, P. D., Serafini, D. B., Torczon, V., and Trosset, M. W., “A rigorous framework for optimization of expensive functions by surrogates,” *Structural Optimization*, Vol. 17, 1999, pp. 1–13.
- [20] Cooper, I. D., and Price, C. J., “On the convergence of grid-based methods for unconstrained optimization,” *SIAM Journal on Optimization*, Vol. 11, 2001, pp. 859–869.
- [21] Audet, C., Savard, G., and Zghal, W., “A mesh adaptive direct search algorithm for multiobjective optimization,” *European Journal of Operational Research*, Vol. 204, 2009, pp. 545–556.
- [22] Samareh, J. A., “Survey of Shape Parameterization Techniques for High-Fidelity Multidisciplinary Shape Optimization,” *AIAA Journal*, Vol. 39, No. 5, 2001, pp. 877–884. <https://doi.org/10.2514/2.1391>.
- [23] Benner, P., Gugercin, S., and Willcox, K., “A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems,” *SIAM Review*, Vol. 57, No. 4, 2015, pp. 483–531. <https://doi.org/10.1137/130932715>.
- [24] Thelen, A. S., Leifsson, L. T., and Beran, P. S., “Aeroelastic Flutter Prediction using Multi-fidelity Modeling of the Aerodynamic Influence Coefficients,” *AIAA Scitech 2019 Forum*, American Institute of Aeronautics and Astronautics, 2019. <https://doi.org/10.2514/6.2019-0609>.
- [25] Forrester, A., Sobester, A., and Keane, A., *Engineering design via surrogate modelling: a practical guide*, John Wiley & Sons, 2008.
- [26] Fuchi, K., Wolf, E. M., Makhija, D., Wukie, N. A., Schrock, C. R., and Beran, P. S., “Enhancement of Low Fidelity Fluid Simulations using Machine Learning,” *AIAA Scitech 2020 Forum*, American Institute of Aeronautics and Astronautics, 2020. <https://doi.org/10.2514/6.2020-1409>.
- [27] Alizadeh, R., Allen, J. K., and Mistree, F., “Managing computational complexity using surrogate models: a critical review,” *Research in Engineering Design*, Vol. 31, No. 3, 2020, pp. 275–298. <https://doi.org/10.1007/s00163-020-00336-7>.
- [28] Blanchard, A., and Sapsis, T., “Output-Weighted Optimal Sampling for Bayesian Experimental Design and Uncertainty Quantification,” *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 9, No. 2, 2021, pp. 564–592. <https://doi.org/10.1137/20m1347486>.
- [29] Gabbard, J., Gillis, T., Chatelain, P., and van Rees, W. M., “An immersed interface method for the 2D vorticity-velocity Navier-Stokes equations with multiple bodies,” *Journal of Computational Physics*, 2022, p. 111339.
- [30] Gillis, T., and van Rees, W. M., “MURPHY—A Scalable Multiresolution Framework for Scientific Computing on 3D Block-Structured Collocated Grids,” *SIAM Journal on Scientific Computing*, Vol. 44, No. 5, 2022, pp. C367–C398. <https://doi.org/10.1137/21M141676X>.
- [31] Audet, C., and JR., J. E. D., “Mesh adaptive direct search algorithms for constrained optimization,” *SIAM Journal on Optimization*, Vol. 17, No. 1, 2006, pp. 188–217.