Aerodynamic Sensitivities Over Separable Shape Tensors

Olga A. Doronina *
National Renewable Energy Laboratory, Golden, CO, 80401

Bumseok Lee†
National Renewable Energy Laboratory, Golden, CO, USA, 80401

Zachary J. Grey‡
National Institute of Standards and Technology, Boulder, CO, 80305

Andrew Glaws§
National Renewable Energy Laboratory, Golden, CO, 80401

We present a comprehensive aerodynamic sensitivity analysis of airfoil parametrization informed by separable shape tensors. This parametrization approach uniquely benefits the design process by isolating various well-studied shape characteristics, such as airfoil thickness, and providing a well-regulated low-dimensional parameter domain for aerodynamic designs. Exploring the aerodynamic sensitivities of this novel parametrization can provide valuable insights for more robust designs and future manufacturing efforts. We construct a data-driven parameter space of airfoils using principal geodesic analysis of separable shape tensors informed by a curated database containing almost twenty thousand suitable engineering airfoils. Analyzing the shape reconstruction error and the maximum mean discrepancy between joint distributions of aerodynamic quantities, we study the dimensionality of the learned parameter space. This simple numerical experiment demonstrates a dramatic dimension reduction that retains design effectiveness and promotes regularity of the shape representations. Finally, we generate new airfoils and use the `HAM2D` RANS solver to predict lift, drag, and moment coefficients. We compute multiple sensitivity metrics to quantify and assert the consistency of parameter influence on the aerodynamic quantities. We also explore low-dimensional polynomial ridge approximations to motivate physical intuitions and offer explanations of the approximated sensitivities.

Nomenclature

*olga.doronina@nrel.gov, Researcher, Computational Science Center
†Researcher, National Wind Technology Center
‡Mathematician, Mathematical Analysis and Modeling Group
§Researcher, Computational Science Center

$\mathbf{x}_i$, $\boldsymbol{X}$	=	i^{th} shape landmark and set of all shape landmarks
n	=	number of shape landmarks
q	=	ambient shape dimension
N	=	number of airfoil shapes in a given data set
$\mathcal{G}(\cdot,\cdot)$	=	Grassmann manifold/Grassmannian
π	=	projection of airfoil shape to Grassmannian
$[\tilde{\boldsymbol{X}}]$	=	equivalence class of shapes defined by the representative element $\tilde{\boldsymbol{X}}$
$\boldsymbol{B}$	=	shape center of mass
$\boldsymbol{P}$	=	symmetric, positive-definite matrix defining polar standardization of airfoil shapes
t_i, $\mathbf{t}$	=	i^{th} normal coordinate and set of all normal coordinates
m	=	dimension of full shape parametrization
r	=	dimension of reduced shape parametrization
$\zeta_{(\cdot)}$	=	upper (u) or lower (ℓ) airfoil surface defined by CST
ψ_{N2}^{N1}	=	class function with class paramters $N1$, $N2$
$\phi_{(\cdot)}$	=	shape function for the upper (u) or lower (ℓ) airfoil surface
$a_{i,(\cdot)}$	=	CST coefficients for the upper (u) or lower (ℓ) airfoil surface
p	=	polynomial degree for CST shape functions
f, $f_{\mathbf{t}}$	=	given function and its output for some input $\mathbf{t}$
S_i, S_{T_i}	=	first-order and total Sobol' index for the i^{th} input parameter
$\boldsymbol{C}$	=	active subspace matrix
ρ	=	probability density function over input space
λ_i, $\boldsymbol{\Lambda}$	=	i^{th} eigenvalue and set of all eigenvalues of $\boldsymbol{C}$
$\boldsymbol{W}$	=	matrix defining the reduced-dimensional subspace
α_i	=	activity score for the i^{th} input parameter
$\tilde{f}$	=	reduced dimensional polynomial approximation of f
C_l	=	lift coefficient
C_d	=	drag coefficient
C_m	=	pitching moment coefficient
α	=	angle of attack
x/c, y/c	=	chord-normalized horizontal and vertical coordinates

I. Introduction

The design and analysis of airfoils are critical to the overall construction of larger and more complex three-dimensional aerodynamic structures, such as wind turbine blades [1, 2], wings and propellers for various aircraft [3–5], and turbine fan blades [6, 7]. These structures are often expected to meet various aerodynamic, structural, acoustic, and economic criteria, making their design a complex, multidisciplinary optimization problem [8–10]. A wide variety of airfoils have been designed that exhibit different properties, either to find a desired balance among these criteria or to perform a specific task while operating within a larger system, e.g., structural airfoils near the root of a wind turbine blade versus high-lift airfoils near the tip. The development of new airfoil designs remains an active area of research, especially with recent progress in machine learning methods [11, 12]

Among the first steps in the airfoil design process is the selection of the design space and the method to represent the airfoil shapes. Common approaches to airfoil parametrization can be categorized into deformative and constructive methods [13]. Deformative methods, such as Hicks-Henne bump functions [14], Bézier curves [15], and radial basis function domain elements [16], apply deformations to a baseline airfoil shape. In contrast, constructive methods, such as B-splines [17], parametric sections (PARSEC) [18], and, more popular recently, the class-shape transformation (CST) [19, 20], represent the airfoil entirely through a set of parameters, enabling the creation of shapes from scratch. However, all these representations lack orthogonality in the design space, an important property that ensures that the parameters defining a shape are independent. Interdependent parameters can result in poor mapping between design variables and shape geometry and cause ill-conditioned optimization problems.

Derivation of a set of orthogonal modes to represent airfoil shapes has been an active area of research in recent years. Toal et al. [21] first proposed using the proper orthogonal decomposition (POD) of an ensemble of airfoils represented by discrete coordinates. The authors applied POD to a small ensemble of airfoils and demonstrated that it can reduce the number of parameters and simultaneously filter out poorly performing geometries. Poole et al. [22] applied a POD to a more extensive library of airfoil shapes, including various airfoil families, and derived more generic design variables. Masters et al. [13] compared deformative methods and constructive methods, including the singular value decomposition (SVD) technique, by evaluating their ability to reconstruct shapes from a diverse dataset of over 2,000 airfoils and found that the SVD method provides the best reconstruction and allows to reduce the number of parameters. Recently, Li et al. [23] suggested using more intuitive and practical modes of camber and thickness instead of modes of shape coordinates and showed that it simplifies the design space bounds. A comprehensive review of modal parameterizations is provided by Li et al. [24]. Another SVD parametrization method, which requires no training dataset of airfoils, was proposed by Kedward et al. [25]. First, Kedward et al. [26] proposed the bounding of parametric derivative by imposing linear constraints on the discrete difference matrix for mesh point control or B-splines to achieve smoothness constraints. Then, Kedward et al. [25] further developed this idea by replacing linear constraints on discrete difference operator with discrete difference operator modes. The authors demonstrated that their method can be easily

applied to nonstandard shapes.

All these approaches demonstrate different characteristics in terms of their ability to represent various airfoil features, their overall expressiveness of new shapes, the number of parameters needed, and the regularity of the parameter space [13]. This last concern is becoming increasingly relevant as machine learning and other data-driven tools are being leveraged for airfoil design [27–30]. In these cases, a well-regularized design space of shapes can accelerate the training of the surrogate models and improve their ability to generalize to new designs.

Recent work has introduced a novel airfoil parametrization technique based on separable shape tensors that learns a matrix manifold representation of the airfoil shapes from a given dataset [31]. This approach shares some similarities with the other data-driven methods discussed above. However, in the above approaches, POD or SVD modes of shape landmarks are calculated extrinsically over the ambient space rather than defining the parameterization intrinsically over the space of relevant airfoil shapes. This can impact the regularity of the design space resulting in large regions of the parameter space mapping to invalid shapes and forcing designers to significantly reduce their parameter ranges or apply posthoc fixes to ensure valid shapes. Further, this intrinsic representation separates affine variations in the airfoil shapes (e.g., those corresponding to translation, scaling, stretching, or rotation of the shape) from higher-order variations defined over the Grassmann manifold (or Grassmannian), which capture the smaller scale, non-linear changes in the airfoil shapes. This separability confers several advantages to the definition of the airfoil design space. First, it enhances the interpretability of the airfoil representation by isolating fundamental aerodynamic variations in shape, such as the airfoil thickness, from other more complex non-linear variations. The affine variations are usually constrained or carefully chosen for 3D blade definition—e.g., the thickness of airfoils near the root has to satisfy structural requirements—while the higher-order variations on their own can considerably improve aerodynamics. The second major advantage of this separability is its regularizing effect on individual shapes across the design space. Separating the two classes of shape variations makes the resulting parameter spaces relatively controlled and more stable to perturbations, even for drastically different shapes. This is intuitively achieved through an improved notion of distance between shapes [32].

In practice, the separable shape tensor approach has been shown to effectively support the development of machine learning models and the optimization of airfoil characteristics. Specifically, the separable shape tensor representation has been used to perform drag minimization of the RAE2822 transonic airfoil in viscous flow cases [33]. Follow-on research integrated this parametrization with surrogate-based optimization to demonstrate dimensionality reduction capabilities and the ability to optimize airfoils with fixed affine deformations (i.e., fixed airfoil thickness) using the base test case of the RAE2822 airfoil. Invertible neural network models have been trained to using separable shape tensors to enable efficient design workflows for design of airfoil cross-sections with larger blade optimization workflow [34]. Despite this emerging research that has examined the use of these Grassmann-based shape parametrizations for predictive modeling and design, no work has performed a formal sensitivity analysis to help understand how these parametrizations affect variation in critical aerodynamic quantities of interest. Such analysis can more broadly enable the future use

of the separable shape tensor framework for aerodynamic design and optimization by identifying critical parameters driving changes in the outputs to accelerate convergence, informing the choice of optimization algorithm by uncovering critical properties of the objective function, and quantifying trade-offs in different properties to guide constraints or multi-objective solutions.

This study aims to fill the gap between the separable shape tensors parametrization and its effects on aerodynamic quantities by performing a variety of sensitivity analyses and motivating physically intuitive explanations of the results. In doing so, we believe this work can help to guide and motivate the use of these tools for more targeted aerodynamic design and optimization problems. Sensitivity analysis is a systematic approach to quantify the impact of given input parameters on specific target quantities of interest for some physical process [35]. This can include local sensitivities that highlight the behavior of the physical processes in the neighborhood of a point in the parameter space or global sensitivities that capture the relative importance of each parameter—or the interactions between multiple parameters—over a full distribution of values. Insights from sensitivity analysis can inform component-wise dimension reduction by discovering insensitive parameters that can be fixed at a nominal value [36–38]. Alternatively, the results from sensitivity analysis can inform other practical considerations, such as where design trade-offs can be made between multiple objectives [39, 40] or how tight manufacturing tolerances need to be set with respect to different aspects of an aerodynamic structure [41, 42].

Quantifying and understanding parametric sensitivities in aerodynamic applications is critical to robust design processes and safety analyses of a wide array of complex engineering applications [43]. Previous works have performed sensitivity analysis to understand the implications of various airfoil designs on aerodynamics [44–48]. Additionally, Liu and Lee [46] examined the sensitivities of different airfoil shape parameters with regard to aeroacoustic sound levels. Similarly, [49] used adjoint-based sensitivity analysis to explore the impacts of different active control techniques to reduce drag forces on an airfoil. Other related works have used sensitivity analysis tools to understand the importance of different turbulence model parameters to the pressure distribution over a wing [50] and the impact of atmospheric conditions on fatigue loads in wind turbines [51, 52]. However, sensitivity analysis results and their interpretations depend on the parametrization and the design space used to characterize the system under consideration. To our knowledge, no previous study has explored the aerodynamic sensitivities of the Grassmann-based separable shape tensors representation for airfoils, except an initial presentation outlined in Grey and Constantine [53], which concentrated on shape sensitivities rather than parameter sensitivities.

The remainder of the paper is structured as follows. Section II provides a brief background of the key techniques leveraged in this work, including separable shape tensors representation, CFD modeling of airfoils, and sensitivity analysis methods. Section III details the methodology used to perform this sensitivity analysis. Section IV presents and discusses the results of the sensitivity analysis. Section V concludes the paper with a summary of the key takeaways and future research directions.

II. Background

A. Separable Shape Tensors for Airfoil Design

In this work, we explore the aerodynamic implications of a separable shape tensors framework [31, 54]. This novel approach learns a data-driven space of shapes as elements of the Grassmann manifold (or Grassmannian), $\mathcal{G}(n, q)$. In the case of shapes defined by planar curves (such as airfoils), we have $q = 2$, and the Grassmann manifold is a space of all 2-dimensional subspaces of $\mathbb{R}^n$, where n is the number of landmarks used to discretely represent the shape. The Grassmannian representations of the airfoils are constructed through a landmark standardization process, such that linear features (e.g., anisotropic scale variations, rotations, and reflections) are ignored. Differences between shapes on Grassmann manifold are isolated to higher order deformations, and perturbations that can be parameterized through a basis representation learned via principal geodesic analysis (PGA) [55, 56]. Thus, this separable shape tensors approach provides a framework for studying sensitivities to more nuanced shape variations independent of larger, linear shape deformations that typically arise in related sensitivity analysis [53, 57]. The remainder of this section provides a brief mathematical overview of the separable shape tensors parametrization. More details on the development and supporting algorithms for implementing this parametrization can be found in [31]. Figure 1 provides a schematic overview of the separable shape tensors approach for the airfoil design.

Following the developments in [31], we define an airfoil shape by a discrete, ordered sequence of landmarks $(\mathbf{x}_i) \subset \mathbb{R}^2$, over $i = 1, \ldots, n$, which form a full-rank $n \times 2$ matrix $\boldsymbol{X} = \begin{pmatrix} \mathbf{x}_1 & \ldots & \mathbf{x}_n \end{pmatrix}^\top \in \mathbb{R}^{n\times 2}$—i.e., points and lines are not considered valid shapes. With the separable shape tensors approach, this discrete shape $\boldsymbol{X}$ is mapped to an element of the Grassmannian through the canonical projection, $\pi : \mathbb{R}^{n\times 2} \to \mathcal{G}(n, 2)$, onto an equivalence class of landmark-affine standardizations which is invariant to general linear transformations. In other words, this projection maps a given airfoil shape $\boldsymbol{X}$ to the underlying equivalence class of shapes $[\widetilde{\boldsymbol{X}}]$ such that the standardized representative $\widetilde{\boldsymbol{X}}$ has orthonormal columns. Numerically, standardization is achieved using the polar decomposition $\boldsymbol{X} - \boldsymbol{B}(\boldsymbol{X}) = \widetilde{\boldsymbol{X}}\boldsymbol{P}$ where $\boldsymbol{B}(\boldsymbol{X}) = \frac{1}{n}\mathbf{1}_{n,n}\boldsymbol{X}$, with $\mathbf{1}_{n,n}$ an n-by-n matrix of ones, computes the center-of-mass to remove shape translations—i.e., we take the polar decomposition of $(\boldsymbol{I} - \frac{1}{n}\mathbf{1}_{n,n})\boldsymbol{X}$. The polar decomposition of the centered shape, $\boldsymbol{X} - \boldsymbol{B}(\boldsymbol{X}) \in \mathbb{R}^{n\times 2}$, is easiest to interpret using the singular value decomposition (SVD). Specifically, given the thin SVD $(\boldsymbol{X} - \boldsymbol{B}(\boldsymbol{X}))^\top = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\top$ such that $\boldsymbol{U} \in \mathbb{R}^{2\times 2}$ is orthogonal, $\boldsymbol{\Sigma} \in \mathbb{R}^{2\times 2}$ is diagonal, and $\boldsymbol{V} \in \mathbb{R}^{n\times 2}$ has orthonormal columns, we construct the 2-by-2 matrix $\boldsymbol{P} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^\top$ paired with a representative Grassmannian element $\widetilde{\boldsymbol{X}} = \boldsymbol{V}\boldsymbol{U}^\top$. The result is equivalent to the polar standardization, [31] i.e., $\boldsymbol{X} - \boldsymbol{B}(\boldsymbol{X}) = \boldsymbol{V}\boldsymbol{\Sigma}\boldsymbol{U}^\top = \boldsymbol{V}\boldsymbol{U}^\top\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^\top = \widetilde{\boldsymbol{X}}\boldsymbol{P}$. Thus, our original shape is separated into the symmetric positive definite matrix $\boldsymbol{P} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^\top$ representing scale variations (up to rotations and reflections) as a right action on the representative Grassmannian element, $\widetilde{\boldsymbol{X}} = \boldsymbol{V}\boldsymbol{U}^\top$. To maintain a separation of scale variations and nonlinear undulations, we independently study learned deformations of shape over nonlinear perturbations $[\widetilde{\boldsymbol{X}}](\mathbf{t})$— equivalence classes modulo the right action of general linear deformations—parametrized by a

vector $\mathbf{t} \in \mathbb{R}^m$ and linear scale deformations $\boldsymbol{P}(\boldsymbol{\ell})$ parametrized by a vector $\boldsymbol{\ell} \in \mathbb{R}^3$. Thus, either separate parameter set can be fixed or varied to deform (centered) shapes, i.e., $\boldsymbol{X}(\mathbf{t}, \boldsymbol{\ell}) = \widetilde{\boldsymbol{X}}(\mathbf{t})\boldsymbol{P}(\boldsymbol{\ell})$ with representative $\widetilde{\boldsymbol{X}}(\mathbf{t})$ taken from $[\widetilde{\boldsymbol{X}}](\mathbf{t}) \in \mathcal{G}(n, 2)$. We elaborate on a data-driven procedure for determining $\boldsymbol{t}$ and $\boldsymbol{\ell}$ in later sections.

This approach is data-driven and relies on a provided dataset of physical shapes $\{\boldsymbol{X}_i\}$ to learn a manifold-valued domain. Moreover, separable shape tensors necessitate a fixed reparametrization of shapes via smoothing or interpolation of landmarks to produce a consistent n total landmarks per shape. The reparametrization of data is often hand-picked to promote refinements around important leading-edge and trailing-edge features of airfoils. Separable shape tensors representations offer the advantage of separately averaging out the effects of $\boldsymbol{\ell}$, representing a dominating influence of variations in scale, while utilizing a learned subspace of $\mathbf{t}$'s to explore the remaining nonlinear variations in the shape representation as $\boldsymbol{X}(\mathbf{t}) = \widetilde{\boldsymbol{X}}(\mathbf{t})\overline{\boldsymbol{P}}$—e.g., we can study shapes with $\overline{\boldsymbol{P}}$ as some notion of average scale. To obtain this parametrization, we use principal geodesic analysis (PGA) [55, 56], which extends the classical principal component analysis (PCA) [58] to Riemannian manifolds, such as the Grassmannian. PGA is a data-driven approach that determines principal components as important directions in a tangent space defined at an intrinsic mean (e.g., the Karcher mean) over the manifold-valued data. The resulting space of so-called normal coordinates (PGA space) is akin to a latent parameter space, denoted $\mathbf{t} = (t_1, t_2, \ldots, t_m)^\top \in \mathbb{R}^m$, of a nonlinear AI-based generative model of airfoil shapes. Note that the maximum dimension of this space is bounded above as $m \leq 2(n-2)$ [59]. Variations in these latent parameters account for fundamental differences in shapes that cannot be achieved through simple linear transformations (e.g., stretching, scaling, shearing, rotation). Thus, characterizing sensitivities of aerodynamic performance with respect to these parameters can provide important, nuanced insights into how small changes in airfoil shape can impact performance. However, we are not precluded from the possibility of generating non-physical self-intersecting shapes. Consequently, we incorporate a simple hard constraint to check for piecewise linear intersections over the generated shape landmarks and reject samples which result in intersections—these are often a small proportion of randomly generated sets.

B. The Class-Shape Transformation

Another airfoil parametrization referenced in this work is the class-shape transformation (CST) [19, 20]. The CST parametrization is a popular method for airfoil representation that encodes the upper and lower surfaces of the airfoil shape by the coefficients of a truncated Bernstein polynomial series. Given a normalized chordal coordinate $c \in [0, 1]$, the CST representation of the airfoil upper (u) and lower (ℓ) surfaces is given by ζ_u and ζ_ℓ respectively such that

$$\zeta_{(\cdot)}(c) = \psi_{N2}^{N1}(c)\varphi_{(\cdot)}(c) + c\zeta_{T_{(\cdot)}}, \tag{1}$$

where $\psi_{N2}^{N1}(c) = c^{N1}(1-c)^{N2}$ is the class function with N_1 and N_2 defining the airfoil class. This work uses $N_1 = 0.5$ and $N_2 = 1$, which correspond to blunt-nosed airfoils. The shape function is given by the Bernstein polynomial

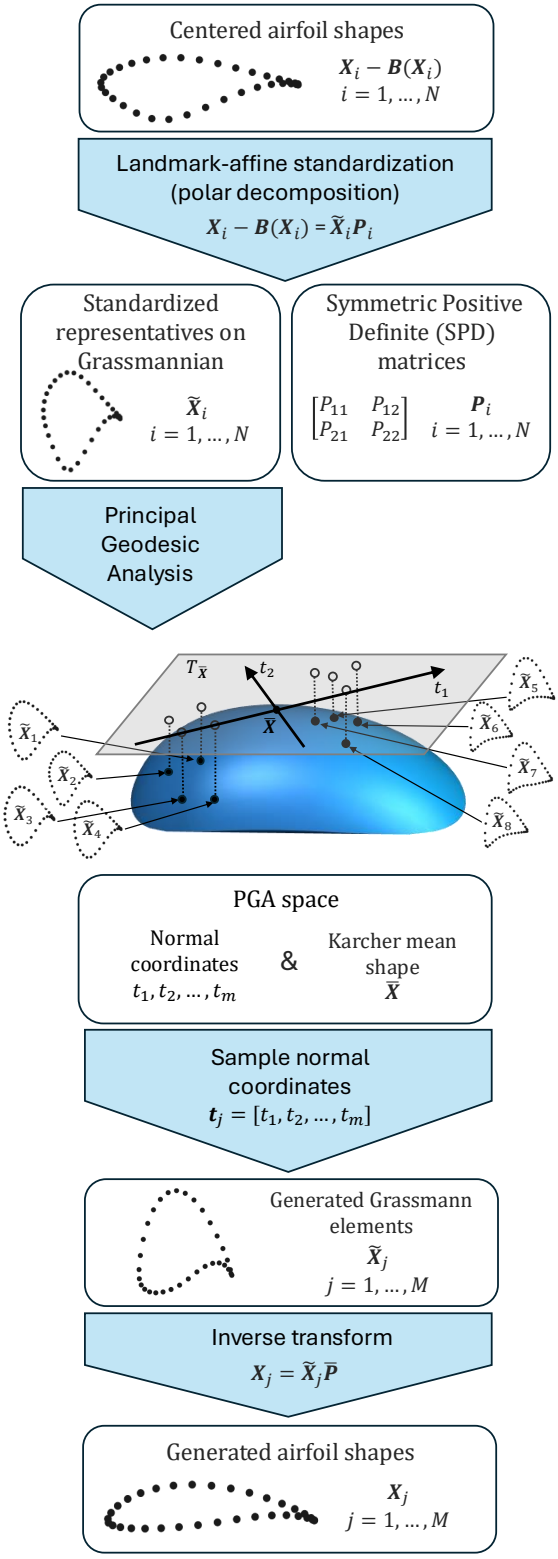

Fig. 1 Diagram representation of separable shape tensors for airfoil design.

expansion of the respective surfaces,

$$\varphi_{(\cdot)}(c) = \sum_{i=0}^{p} a_{i,(\cdot)} \binom{p}{i} c^i (1-c)^{p-i}, \tag{2}$$

with free parameters, $a_{i,u}$ for the upper and $a_{i,\ell}$ for the lower surface, that can be adjusted to define different airfoil shapes. Finally, the upper/lower trailing edge thickness of the airfoil is represented by $\zeta_{T_{(\cdot)}}$; again, with one per surface.

Given this definition, a p^{th}-order CST parametrization requires $m = 2(p+2)$ parameters to define both the upper and lower surfaces of an airfoil. This includes $2(p+1)$ free parameters in Bernstein polynomials and 2 parameters that account for the trailing edge thickness. Ceze et al. [60] studied the performance of this method to reconstruct airfoil shapes and determined that roughly $m = 20$ total terms were necessary to accurately define airfoil shapes. These insights have supported the use of CST parametrization in various airfoil design studies [61, 62]. It is natural to perhaps employ this parametrization as a global representation of shape. However, in [31], it is noted that the global CST parameter sweeps over $a_{i,(\cdot)}$'s spanning a diverse set of airfoils can create undesirable shapes. Regardless, locally, this representation offers a very useful engineering parametrization. In this work, we use 8^{th}-order CST parametrization to pre-process otherwise noisy data and get a consistent number of landmarks for all airfoil shapes.

C. CFD Simulations of 2-D Airfoil

To obtain aerodynamic responses for our generated airfoil shapes, we leverage CFD simulations using an in-house, finite-volume, Reynolds-averaged Navier-Stokes (RANS) flow solver, `HAM2D` [63]. `HAM2D` is distinguished by its ability to identify line structures on unstructured meshes and apply line-based schemes for spatial reconstruction and implicit inversion, similar to a structured mesh-based flow solver. In this work, we employ a fifth-order Weighted Essentially Non-Oscillatory (WENO) scheme [64] for spatial reconstruction, with Roe's flux difference scheme [65] for inviscid flux, and second-order central differencing for viscous flux. A preconditioned Generalized Minimal Residuals (GMRES) method [66] is applied for implicit integration. The Spalart-Allmaras one-equation turbulence model [67] is used for the turbulence closure, and the Medida-Baeder two-equations transition model [68] is applied to account for the effects of laminar-turbulent transition. The two-dimensional airfoil grid is generated with 400 points on the airfoil surface, the dimensionless wall-normal spacing of 1, and an outer boundary located at 300 chord lengths away from the wall. The CFD simulations are performed at a freestream Mach number of 0.1, Reynolds number of 9×10^6, and two angles of attack, 4° and 12°. For laminar-turbulent simulations, the freestream turbulence intensity is set to 0.1%. To manage CFD cases involving thousands of airfoils, all stages of CFD simulations, including mesh generation, flow solver setup, CFD runs, and post-processing of results, are automated using Python and bash scripts. Details of the mesh generation, grid sensitivity study, and validations of the simulation results for airfoils at Reynolds numbers ranging from 3 million to 15 million are available in Jung, et al. [69].

D. Sensitivity Metrics

Sensitivity analysis seeks to quantify the impact that uncertainty or changes in certain parameters (i.e., model inputs) have on other parameters (i.e., model outputs). That is, we may consider a given relationship $f_{\mathbf{t}} = f(\mathbf{t})$, where $\mathbf{t} \in \mathbb{R}^m$ is a collection of model inputs whose uncertainty is represented by the probability density function ρ, and $f_{\mathbf{t}} \in \mathbb{R}$ is the model output. Different sensitivity analysis methods can quantify these impacts in a local sense (i.e., in the neighborhood of a particular point, $\mathbf{t}_0$) [70] or in a global sense (i.e., over all valid values of $\mathbf{t}$) [35]. In this work, we consider several global sensitivity metrics to provide a robust characterization of the relationship between shape parameters and the aerodynamic quantities of interest.

The first method we consider is a variance-based metric of sensitivity known as Sobol' indices [71]. First-order Sobol' indices capture the ratio of the total response variance attributed to a particular parameter and are defined as

$$S_i = \frac{\mathrm{Var}\left[\mathbb{E}\left[f_{\mathbf{t}} \mid t_i\right]\right]}{\mathrm{Var}\left[f_{\mathbf{t}}\right]}, \qquad i = 1, \ldots, m. \tag{3}$$

These indices capture the individual impact of each parameter independent of their interactions with other parameters, while higher-order Sobol' indices capture sensitivities to interactions between parameters. The collection of all the first- and higher-order sensitivities for a particular parameter accounts for its total contribution to the variation in the output. This is referred to as the total Sobol' indices and is expressed as

$$S_{T_i} = 1 - \frac{\mathrm{Var}\left[\mathbb{E}\left[f_{\mathbf{t}} \mid \mathbf{t}_{\sim i}\right]\right]}{\mathrm{Var}\left[f_{\mathbf{t}}\right]}, \qquad i = 1, \ldots, m, \tag{4}$$

where $\mathbf{t}_{\sim i}$ is the collection of all elements of $\mathbf{t}$ except for the i^{th}. If $S_{T_i} = 0$ then the t_i has no impact on $f_{\mathbf{t}}$.

Another class of sensitivity metrics leverages subspace-based dimension reduction that seeks to represent most of the variation in the output of interest in terms of a reduced set of linear combinations of the original input parameters. First, we consider the activity score metric [72], which is based on the gradient-based dimension reduction framework called active subspaces [73]. Active subspaces are defined by the eigenspace of the mean outer product of gradients,

$$\boldsymbol{C} = \int \nabla f(\mathbf{t}) \nabla f(\mathbf{t})^{\top} \rho(\mathbf{t}) d\mathbf{t} = \boldsymbol{W} \boldsymbol{\Lambda} \boldsymbol{W}^{\top}, \tag{5}$$

where $\boldsymbol{\Lambda} = \mathrm{diag}\left(\lambda_1, \lambda_2, \ldots, \lambda_m\right)$ is a diagonal matrix of eigenvalues of $\boldsymbol{C}$ sorted from largest to smallest and $\boldsymbol{W} = \left[\mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_m\right]$ are the associated eigenvectors. The eigenvalue λ_i has been shown to correspond to the mean squared directional derivative of f along the direction $\mathbf{w}_i$. The r-dimensional active subspace of f is defined as the span of r eigenvectors associated with the largest eigenvalues. In practice, computing the active subspace requires the calculation or approximation of the gradient ∇f, which may be difficult to obtain. For this work, we approximate the gradients at a collection of Monte Carlo samples by constructing a local linear model around each sample using its k

nearest neighbors. Given an active subspace for $f(\mathbf{t})$, the activity score for the i^{th} parameter t_i is defined as

$$\alpha_i(r) = \sum_{j=1}^{r} \lambda_j w_{i,j}^2. \tag{6}$$

Next, we consider the sensitivity metric defined by the ridge directions obtained from a dimension reduction method known as polynomial ridge approximation [74]. Polynomial ridge approximation learns a low-dimensional subspace representation by efficiently solving a nonlinear least-squares problem,

$$\min_{\substack{[\boldsymbol{W}]\in\mathcal{G}(m,r)\\ \widetilde{f}\in\mathcal{P}_d}} \left\| f(\mathbf{t}) - \widetilde{f}(\boldsymbol{W}^\top \mathbf{t}) \right\|_2^2, \tag{7}$$

where $\widetilde{f}$ is a d-degree polynomial approximation of f defined over a low-dimensional subspace of the input space—coincidentally also an element of the Grassmannian but unrelated to notions of shape. The joint optimization is simplified and solved using a variable projection technique. By setting $r = 1$, we consider a one-dimensional reduction of f and can view the weights in $\boldsymbol{W}$, now just a column vector, as a rough measure for the relative amount of variation in f that is driven by the corresponding parameter [75]. The degree to which this sensitivity can be trusted is quantified by the error in the low-dimensional approximation. This can be qualitatively evaluated by examining the shadow plots that show the relationship between f and the learned one-dimensional parameter [76].

III. Methodology

This study workflow involves two steps: (i) building a data-driven parameterization of airfoils based on principal geodesic analysis (PGA) of separable shape tensors and (ii) performing sensitivity analyses of aerodynamic quantities over this parameter space. To build the data-driven parameter space, we curate a database of almost twenty thousand trusted airfoil shapes from a variety of sources. We then represent these airfoils by separable shape tensors and apply PGA to get the parametrization. This can be done in a matter of minutes and without high-performance computing. Once the parametrization has been learned, we sample this parameter space and use CFD simulation software to obtain corresponding aerodynamic properties. Finally, we apply the various sensitivity metrics to the collection of aerodynamic outputs to highlight various shape features that impact performance. In this section, we provide the technical details for this approach before exploring the sensitivity analysis results in the next section.

A. The Curated Database

Our goal is to develop a curated set of airfoils with various characteristics that will inform the data-driven parameter space. We begin by gathering airfoils from the extensive BigFoil database [77], which is a large database that consolidates airfoil data from the University of Illinois Urbana-Champaign (UIUC) airfoil database [78], the JavaFoil database [79],

the NACA-TR-824 database [80], and other sources. The diverse origins of these airfoils mean that, while they cover a wide range of interesting features worth studying, airfoil landmark data can often be sparse, noisy, or incomplete. To ensure a reasonable design space of airfoils, we performed a pre-processing to clean the data and identify undesirable shape characteristics unsuitable for this analysis. Specifically, we removed airfoils with too few landmarks to adequately describe the shape (i.e., 25 or fewer landmarks) and those whose landmark data appeared to be noisy upon visual inspection. We also removed airfoils with shape characteristics outside of this work interest—e.g., airfoils with a sharp leading edge or an extreme camber line as well as those exceeding a threshold of 0.025 for the trailing edge thickness. Further, we note that the BigFoil dataset contains a disproportionate number of variations of NACA airfoils. Specifically, it contains parametric sweeps of NACA airfoils with incrementally increasing thickness and camber. To avoid a bias towards these shapes in our design space, we thinned out the collection of these airfoils by selecting every fourth step in the parameter sweeps. The final step involves regularizing the airfoils by fitting the shapes with an 8^{th}-order CST parametrization and then removing airfoils with high reconstruction errors (i.e., greater than 10^{-2} root mean squared error of the given landmarks). This data pre-processing resulted in $2,343$ airfoils being removed to produce a set of $6,164$ airfoils with $1,001$ landmarks resampled from 8^{th}-order CST representation with x coordinates following cosine distribution along the chord.

To augment this data set, we also generated additional airfoils using CST parametrization. We identified 13 baseline airfoils from the NREL 5MW [81] and IEA 15MW reference wind turbines [82] and used least-squares fits of CST parametrizations with fixed order $p = 8$ to reparametrize and sample shape landmarks. We then sampled $1,000$ unique airfoils by uniformly perturbing all 18 CST coefficients (9 coefficients for upper and lower surfaces polynomials) by up to $\pm 20\%$ around the baseline for each airfoil, resulting in $13,000$ additional airfoil shapes with 1,001 shape landmarks whose x-coordinates follows a cosine distribution along the chord.

The final curated database of airfoils combines the set of 6,164 airfoils from BigFoil and 13,000 airfoils defined by random perturbations to CST coefficients of baseline airfoils. Thus, we have a total of $N = 19,164$ airfoils, each with 1,001 landmarks, to define data-driven design space using PGA of separable shape tensors. This curated database of airfoils was made available online [83]. Figure 2 provides a statistical overview of the final curated database and highlights the range of characteristics of airfoils used to inform the data-driven parameter space.

B. Construction of the PGA space

With the curated airfoil dataset, we follow the separable shape tensors framework [31] described in Section II.A, which decomposes each airfoil shape, $\boldsymbol{X}_i$ for $i = 1, \ldots, N$, into a representative Grassmannian element, $\widetilde{\boldsymbol{X}}_i$, and fixed linear transformation defined by an SPD matrix, $\boldsymbol{P}_i$. Our focus in this work was to explore the nuanced aerodynamic sensitivities of nonlinear (undulating) variations in shapes which are not conflated with linear variations. To remove variance in linear transformation and obtain a consistent right-inverse from the Grassmann manifold to the physical space,

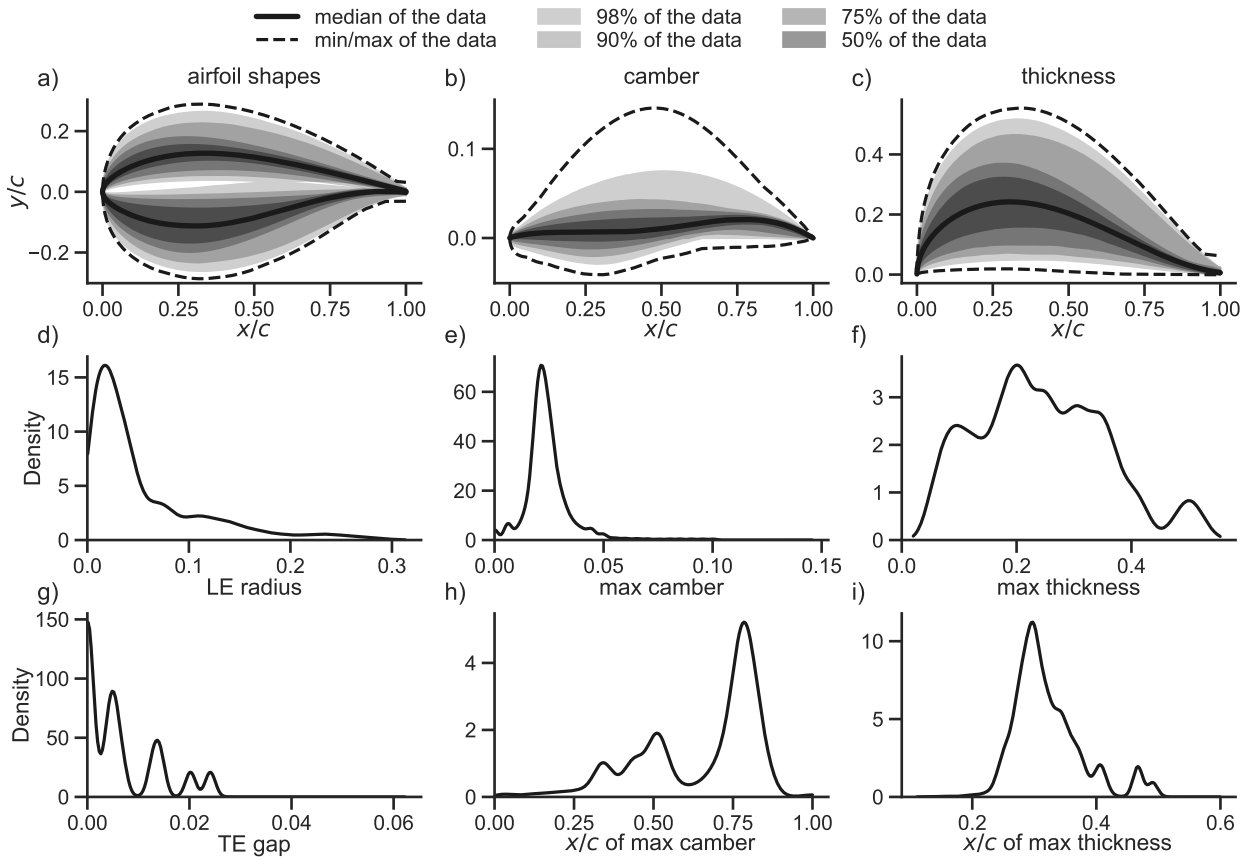

Fig. 2 Statistical characterization of the curated database, including the chordal distribution of airfoil shape, camber, and thickness (top row) and the distribution of various other shape characteristics (bottom two rows).

we compute a mean transform $\overline{\boldsymbol{P}}$ as the Karcher mean of all matrices $\{\boldsymbol{P}_i\}$ over the SPD manifold [84]—constituting an average of scale variations. In this sense, we average-out the effects of parameter variations over ambiguous $\boldsymbol{\ell} \in \mathbb{R}^3$ and only study nonlinear undulating variations in the shapes with fixed linear scales. Then we apply PGA [56] to the set of representative Grassmannian elements, $\{\widetilde{\boldsymbol{X}}_i\}$, to obtain data-driven (normal) coordinates $\boldsymbol{t}$ for shape parametrization. We obtain the set of latent variables $\{\mathbf{t}_i\} \subset \mathbb{R}^m$ with PGA to capture fundamental nonlinear shape deformations over the set of airfoils. Figure. 1 schematically summarizes this procedure, which details and supporting algorithms provided in [31]. Figure 3 visualizes the amount and direction of perturbation to the Karcher mean shape by each normal coordinate, $\mathbf{t} \in \mathbb{R}^m$. We only depict the first 18 PGA modes because the curated airfoil data has been smoothed with an 18-term least-squares CST fit—i.e., order $p = 8$ for upper and lower surface Bernstein polynomial expansions. Smoothing restricts the ambient dimensionality of the shapes to $m = 18$ as a mechanism for removing noise from the data, and the first 18 PGA modes completely describe shape variations in the curated dataset—i.e., the PGA eigenvalues drop to machine precision beyond the first 18 modes. In other words, the data-driven PGA routine perfectly captures the chosen ambient $(m = 18)$-dimensional fits up to machine precision.

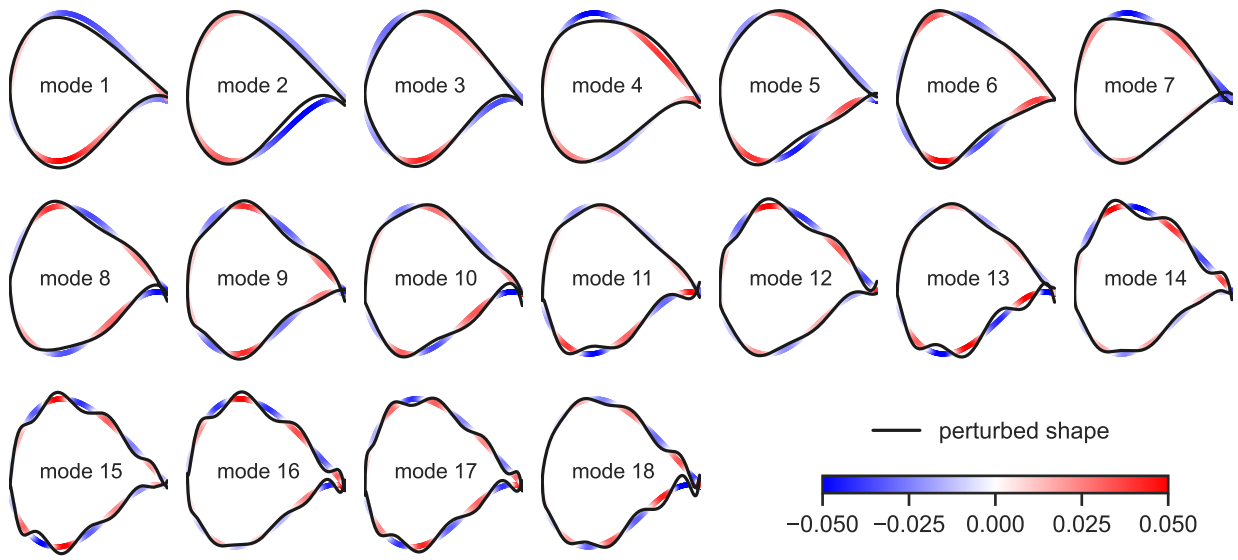

Fig. 3 Visualization of shape variations (black line) about the Karcher mean (colored line) for perturbations to each PGA parameter. The line color indicates the inner product of shape unit normal and the mode perturbation.

C. Selecting Dimensionality

With a curated database of $m = 18$ dimensional shapes, we explore how much of the full dimensionality is required to produce reasonable coverage of output aerodynamic quantities of interest. Also note that the ordered PGA modes exhibit more oscillatory characteristics with increasing mode number, as shown in Fig. 3. This creates a trade-off between the shape's physical expressivity versus regularization against non-physical airfoil features (e.g., divots or kinks). Thus, one critical question that arises during the construction of the design space is: how many PGA modes are necessary to explore a relevant space of airfoil shapes while providing sufficient regularization against non-physical shapes? That is, what is the idealized effective dimension for our PGA parametrization? To address this question, we analyze the reconstruction of airfoil shapes and their aerodynamic properties for variable parametrization dimensions.

For this analysis, we systematically selected 100 airfoils to span a wide range of the characteristics present in the curated dataset. We included the baseline airfoils from the NREL 5MW [81] and IEA 15MW reference wind turbines [82], the largest perturbations to the baseline CST parameters, airfoils that had high shape reconstruction error, and airfoils that are far from the Karcher mean shape over the dataset (i.e., distance-based outlier shapes). This selection ensures that we capture a large amount of variability in the airfoils chosen for this analysis and don't unfairly bias the subset of 100 hand-picked designs towards shapes that are inadvertently easier to represent or not representative of the wind application.

Figure 4 shows a steady decay in the selected shapes' reconstruction errors (calculated as the maximum Euclidean distance between corresponding landmarks of the initial and reconstructed shapes) without an obvious point where the change in errors suddenly levels out. As such, we must consider the trade-off between the expressiveness of the PGA

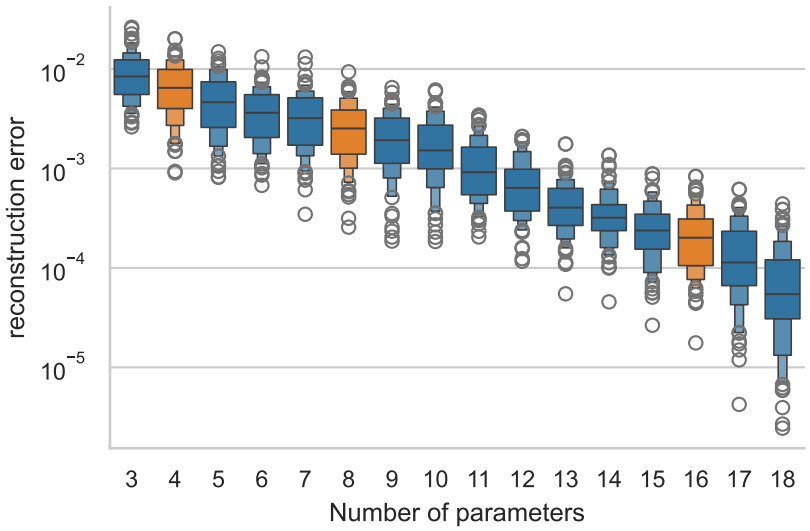

Fig. 4 Letter-value plot of shape reconstruction errors over the set of 100 airfoils as a function of PGA dimension.

representations for various dimensions and the complications that arise from working in a higher dimensional parameter space. We choose the $r = 4, 8$, and 16 dimensions as cases for further investigation.

Specific to this inquiry is the impact of dimension on the recovery of relevant aerodynamic quantities: lift coefficient, C_L, drag coefficient, C_D, and moment coefficient, C_M. CFD simulations were performed at low ($\alpha = 4°$) and high ($\alpha = 12°$) angles of attack for each of the 100 hand-picked airfoils and their reconstructions with $r = 4, 8$, and 16 parameters using the finite-volume RANS flow solver, `HAM2D` described in detail in Section II.C. Figure 5 shows the distribution of absolute errors (normalized by the mean of the true values) in the recovered aerodynamic quantities. In this figure, the higher angle of attack has a slightly larger error, the errors between the $r = 4$ and $r = 8$ cases are relatively similar, and a much larger dimension ($r = 16$) is required to noticeably reduce the error for all aerodynamic quantities. While we see that lower-dimension reconstructions produce larger errors, the focus is to identify a design space that is sufficiently expressive to cover the complete range of aerodynamic quantities of interest. Towards this end, we consider the distribution of the lift, drag, and moment coefficients produced by the airfoil reconstructions across different dimensions.

To identify this effective dimension, we examine the generated joint distributions of aerodynamic responses. Figure 6 qualitatively highlights that the first four PGA dimensions appear sufficient for recovering the statistical moments of these aerodynamic coefficients. Quantitatively, we employ maximum mean discrepancy [85] to test if the shapes generated with different dimensionality significantly modify the joint distribution of the three aerodynamic quantities of interest. That is, we compare each of the joint distributions of aerodynamic coefficients generated by the low-dimensional ($r = 4, 8, 16$) airfoil representations to the coefficients generated by the original airfoils, resulting in three different approximations of the maximum mean discrepancy. In this numerical experiment, we found no statistically significant

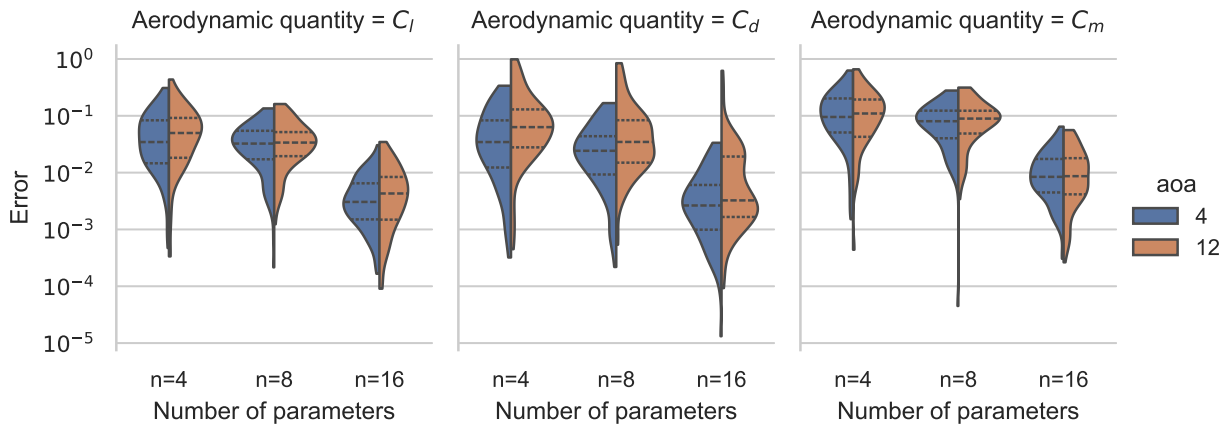

Fig. 5 The distribution of error in the recovery of aerodynamic quantities for low ($\alpha = 4°$) and high ($\alpha = 12°$) angles of attack, with respect to dimensionality of shape parametrization.

difference between the joint distributions of the aerodynamic quantities obtained from shapes constructed by any of the reduced PGA representations. We subsequently conclude that, despite the larger errors in the reconstruction of shapes and aerodynamic quantities with reduced dimension shown in Fig. 4 and Fig. 5, the joint distribution of aerodynamic quantities is still well described by shapes with as few as four dimensions. Intuitively, this suggests that even with as few as four parameters for shape representation, we can obtain sufficient coverage in the distribution of aerodynamic responses that would drive subsequent optimization and design problems. Empirically, we conclude that these different reduced-dimension designs span an equivalent range of aerodynamic quantities of interest.

An additional benefit of keeping only the first four PGA modes is that we reduce the oscillatory characteristics that are depicted in higher modes in Fig. 3. This truncation of ordered terms in the normal coordinate basis expansion offers, intuitively, a kind of explicit regularization in the shape representation to prevent undesirable high-order oscillations in generated designs. Consequently, we select the $r = 4$ case for the analysis performed in the remainder of this paper. This learned design space of four parameters informs a dramatic dimension reduction for generating new shapes over the ordered normal coordinates restricted to $\mathbf{t} = (t_1, t_2, t_3, t_4)^\top \in \mathbb{R}^r$. With this parameter space, we randomly generate designs to study changes in aerodynamic responses using CFD simulations.

D. Sampling PGA Space

The final step in our sensitivity analysis is to generate new reduced-dimension airfoil shapes by sampling the constructed PGA design space and obtaining aerodynamic quantities for the airfoils. Figure 7 visualizes the joint and marginal distributions of normal coordinates for the original curated dataset of airfoils and the new samples of the design space. Note that the original dataset results in a relatively irregular distribution of points in this space with several outliers in each dimension. To better manage the design space for sensitivity analysis, the new samples are drawn with uniform probability from an ellipsoid defined over the normal coordinates $\{t_1, t_2, t_3, t_4\}$. This sampling distribution

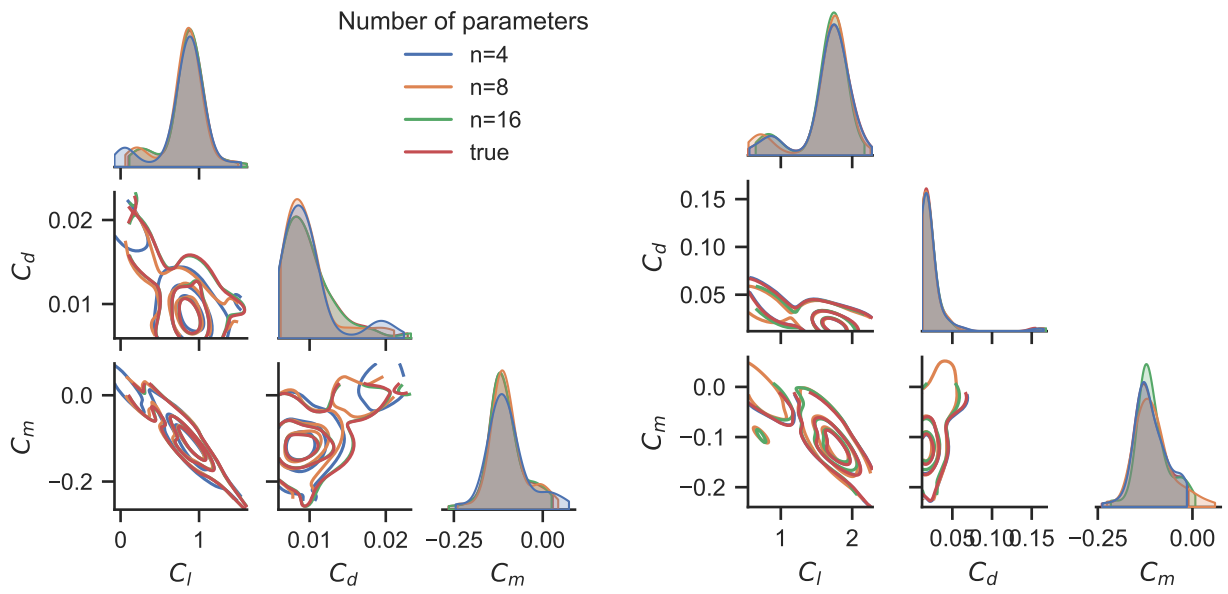

(a) $\alpha = 4°$ **(b)** $\alpha = 12°$

Fig. 6 The joint and marginal distributions for the output aerodynamic quantities C_l, C_d, and C_m for a variety of dimensionalities.

allows for a balanced exploration of the parameter space while avoid extreme designs that could occur in the corners of a bounding hypercube. The ellipsoid is centered over the middle 95% of the original dataset, with the diameter along each axis defined to cover 99% of this dataset. This ensures that the sampling space covers most of the curated airfoils without allowing outliers to impact our sensitivity analysis. It is also worth noting that global sensitivity analysis depends heavily on the probability distribution underlying the design space since it weighs the relative importance of different shapes. To avoid biasing our sensitivity analysis, we draw samples of PGA parameters with uniform probability from this ellipsoid. Sampling constraints are utilized to detect the possibility of self-intersecting shapes. Specifically, airfoils with simultaneously large negative t_1 and t_2 values resulted in self-intersecting shapes, where the intersection occurred near the tail edge of the airfoil. Such shapes are always checked against a hard constraint in the sampling procedure and removed from the sampled set to appropriately regularize the generative model.

The Grassmannian elements obtained from this sampling process were mapped back to physical airfoil shapes using the inverse SPD matrix fixed at the Karcher mean of the SPD manifold. This constitutes a notion of shapes generated with some fixed average scale. This also ensured that the aerodynamic sensitivity analysis was not impacted by these well-studied factors—e.g., thicker airfoils are known to increase drag—but could be focused on the nuanced shape deformations arising from the separable shape tensors representations. Figure 7 also shows the distributions of the airfoil shapes and the camber and thickness of the airfoils along the chord.

Finally, we ran CFD simulations for the resulting set of 8,996 airfoils using the `HAM2D` solver. From these simulations,

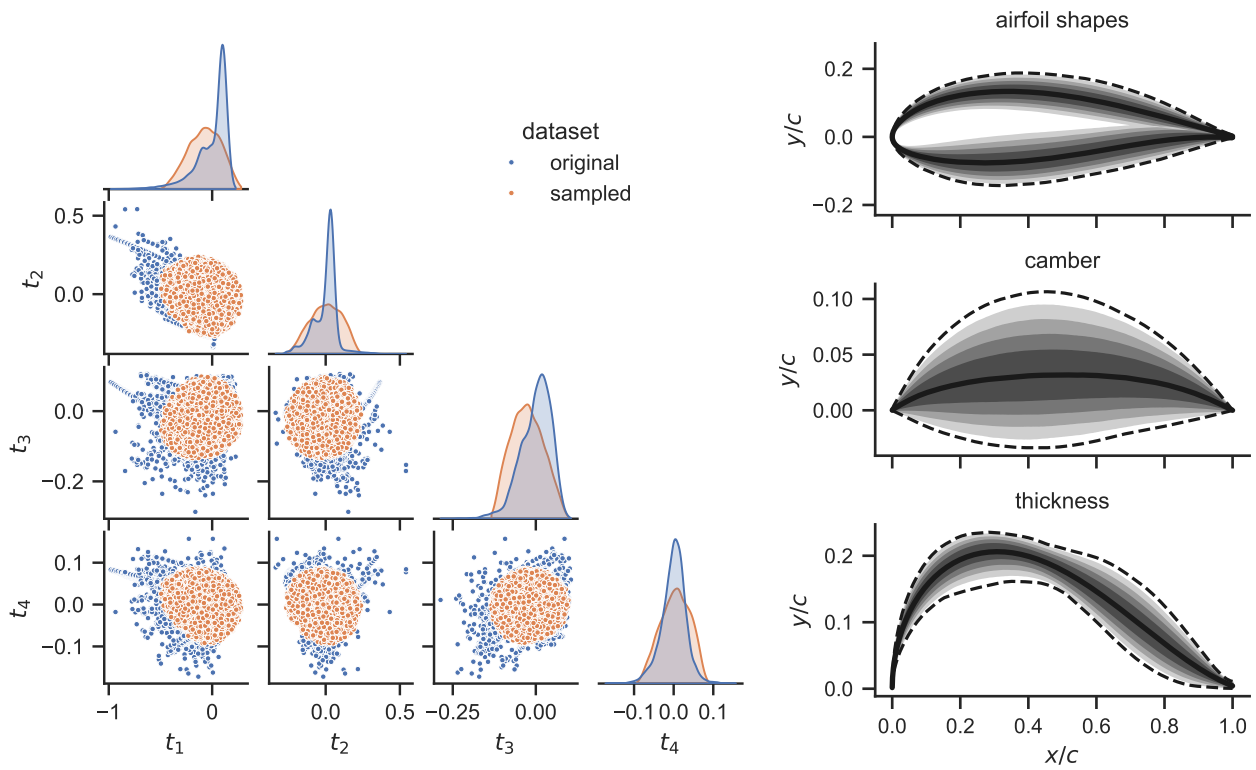

(a) Joint and marginal distributions of the normal coordinates for the original curated dataset and sampled airfoils.

(b) The distribution of sampled airfoil shapes and camber and thickness along the chord.

Fig. 7 A visualization of the sampled airfoils

we obtain values for the lift, drag, and moment coefficients for each of these airfoils for low ($\alpha = 4°$) and high ($\alpha = 12°$) angles of attack. This CFD data has been made publicly available on the Open Energy Data Initiative (OEDI) webpage [86]. In the next section, we explore the results of the sensitivity analysis for these airfoil shapes and discuss how the insights can help to improve airfoil design in the future.

IV. Results

We now explore the sensitivities of the computed aerodynamic quantities with respect to the PGA parameters used to generate the shapes based on the separable shape tensors approach. Given the novelty of this shape representation framework, it is critical to understand how these learned parameters can impact aerodynamic performance to inform their use for airfoil design. We begin by computing the sensitivities of the various aerodynamic quantities with respect to the normal coordinates using multiple approaches described in Section II.D. We look for consistency across these metrics and examine their uncertainty. We perform a deeper exploration of the subspace-based sensitivity approach that identifies combinations of parameters that capture most of the variation in the output quantities of interest. This allows us to translate sensitivities with respect to shape parameters into physical variations in the airfoil shapes.

A. Sensitivity indices

We begin by computing a variety of sensitivity metrics for each aerodynamic quantity (C_l, C_d, and C_m) for each angle of attack ($\alpha = 4°, 12°$) with respect to the four normal coordinates $\{t_1, t_2, t_3, t_4\}$. The sensitivity metrics we compute here are the total Sobol' indices from Eq. (4), the first-order Sobol' indices from Eq. (3), the activity scores from Eq. (6), and the magnitude of the subspace weights from the polynomial ridge approximation, i.e., the vector $\boldsymbol{W}$, for $r = 1$, from Eq. (7). The total and first-order Sobol' indices are computed using the Sensitivity Analysis Library in Python (SALib) [87] with a surrogate-based approach that fits a Gaussian process model to the data to estimate the indices. The activity scores are computed using a local linear approximation method to obtain gradients at each point to compute the active subspace [88]. The polynomial ridge approximation is performed using the Parameter Space Dimension Reduction (PSDR) Toolbox [89].

Figure 8 shows these sensitivity metrics computed for each aerodynamic quantity with respect to the normal coordinates. Despite each of the sensitivity metrics capturing different characteristics of the input-output relationship (e.g., statistical, gradient, or subspace-based), we see a strong agreement across the methods in conveying the relative importance of each of the PGA modes. In particular, the lift coefficient, C_l, heavily depends on the first two modes, with the latter two modes impacting this quantity very little. Conversely, the drag coefficient, C_d, is strongly influenced by these latter modes, albeit differently depending on the angle of attack—i.e., at the low angle of attack, dependence is stronger on the fourth mode, while at the high angle of attack it is stronger on the third. Finally, the moment coefficient has a more equal dependency on all of the PGA modes, slightly favoring the first two, with a small variation in sensitivity between the low and high angle of attack.

Next, we address the question of how well we can trust the sensitivity metrics presented in Figure 8. To do this, we compute the same metrics using randomly sampled subsets of the full dataset to estimate the convergence. Figure 9 shows the errors between the metrics over the subsets and the full dataset of 8,996 airfoils. We perform this convergence study for subsets of size N = 50, 100, 300, 500, 1000, 3000, 5000, and 8000 and report the average error across 200 resamplings. This analysis shows that the Sobol' metrics show weaker convergence than the ridge approximation and activity scores. However, once we have more than $N = 1,000$ subsampled shapes (or roughly 11% of the full dataset), the variation in the computed sensitivity metrics is typically less than 0.01. This amount of variation would not drastically impact interpretations of the results since the differences between high and low sensitivity indices were generally an order of magnitude larger.

B. Ridge Subspaces and Airfoil Sweeps

The results presented in the previous section provide insight into how to leverage the PGA space to improve airfoil designs. However, the interpretations of the sensitivity metrics computed with respect to the normal coordinates are somewhat disconnected from the physical variations that are driving these aerodynamic changes. To address this, we

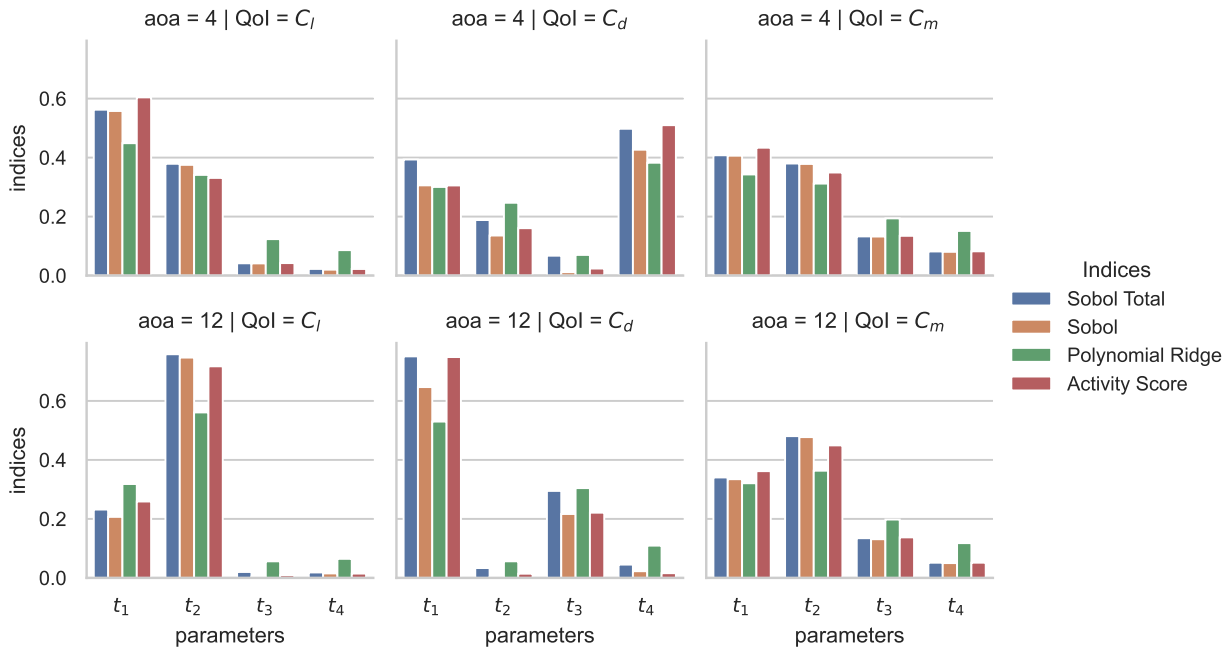

Fig. 8 Comparison of four sensitivity metrics for aerodynamic quantities C_l, C_d, and C_m, with respect to the normal coordinates $\{t_1, t_2, t_3, t_4\}$ for the low ($\alpha = 4°$) and high ($\alpha = 12°$) angles of attack.

explore the results from the polynomial ridge approximation more deeply.

Subspace-based dimension reduction defines low-dimensional projections of input parameters that capture exploitable structure in the data. It can be done in an unsupervised manner using PCA [58] or in a supervised manner using methods such as active subspaces [73] or sufficient dimension reduction [90]. For this work, we use the CFD results to inform a data-driven polynomial ridge approximation approach to identifying exploitable low-dimensional structure [74]. As discussed in Section II.D, this approach finds the linear combination of input parameters that minimizes a polynomial fit to the output quantities. This combination of parameters provides insight into how we can maximally vary each aerodynamic quantity by changing the PGA values. This can be critical for informing aerodynamic design processes where we may seek parameter variations that can reduce drag without impacting lift. Further, low-dimensional structure can be exploited to identify ranges of target objective values [91], to improve multi-objective optimization via subspace-based Pareto tracing [40], or to inform constraints for targeted sampling and efficient design space exploration [92]. Moreover, subspace-based dimension reduction can help to uncover local optima or other regions within the domain that could prove problematic for optimization in high-dimensional problems.

Applying this method to the generated data, we identify strong one-dimensional subspace approximations for the lift, drag, and moment coefficients that exhibit highly accurate quadratic approximations for each separate response. Specifically, we find that the coefficient of determination for each fit is $R^2 > 0.91$ for all quantities (see Table 1). This identification of low-dimensional structure in aerodynamic shapes is consistent with related work [53, 57]. However, the

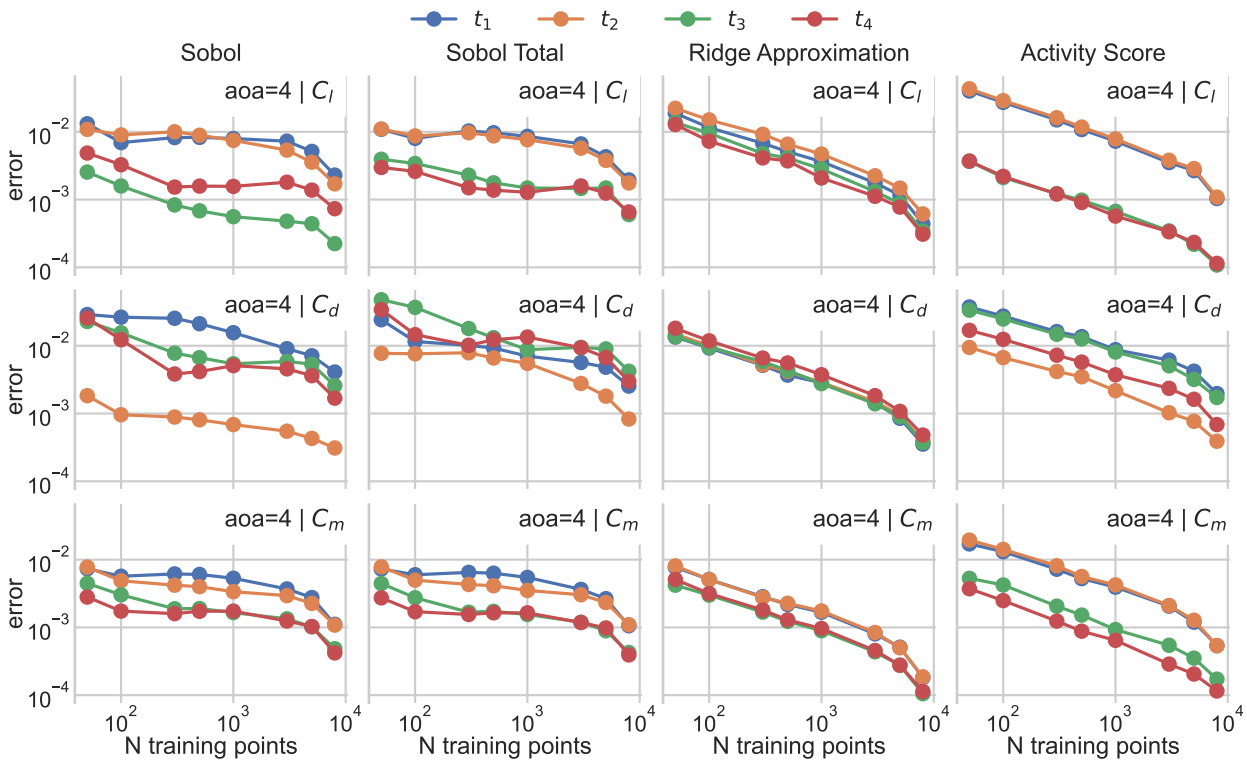

Fig. 9 Convergence of sensitivity metrics with increasing number of data points for cases with angle of attack $\alpha = 4°$.

resulting deformations are notably distinct and arguably novel due to differences in the underlying parametrizations, which are no longer conflated with variations in the generalized scale of the shape, which is known to dominate aerodynamic sensitivities.

Figure 10 contains shadow plots (top rows) visualizing the goodness-of-fit for each polynomial ridge approximation over the computed one-dimensional active subspace $\boldsymbol{W}^\top \mathbf{t}$. The entries of $\boldsymbol{W}$ are given above each plot for each aerodynamic quantity and angle of attack. These values provide insight into the relative influence and the direction of the influence of each normal coordinate t_i on the given output. Larger magnitudes imply that the given aerodynamic quantity is more sensitive to changes in the parameter, and positive values imply a positive correlation with the aerodynamic quantity.

By sampling the PGA space along these one-dimensional subspaces for each aerodynamic quantity and angle of attack, we obtain the physical deformations to the Karcher mean airfoil along this subspace, shown in the bottom rows of Figures 10a and 10b. These deformations drive the changes in the aerodynamic performance of the airfoils the most. In particular, the lift coefficient is increased by increasing the overall camber of the shape. However, the manner in which this is achieved is different between the low and high angles of attack. Specifically, the higher angle of attack seeks to increase the camber closer to the trailing edge of the airfoil, while the lower angle of attack does so closer to the middle

Table 1 Values of coefficient of determination, R^2, for polynomial ridge approximation of aerodynamic quantities using different subspace dimension r and different polynomial degree d.

		aoa = 4^o			aoa = 12^o		
		C_l	C_d	C_m	C_l	C_d	C_m
$r = 1$	$d = 1$	0.992	0.898	0.996	0.972	0.866	0.992
	$d = 2$	0.994	0.914	0.997	0.972	0.965	0.993
	$d = 3$	0.994	0.915	0.997	0.972	0.967	0.994
$r = 2$	$d = 2$	0.999	0.963	0.999	0.997	0.977	0.999
	$d = 3$	1.0	0.966	0.999	0.997	0.98	1.0
$r = 3$	$d = 3$	1.0	0.983	1.0	0.999	0.988	1.0

of the airfoil. The drag coefficient similarly increased by increasing camber, and this change is optimized by pushing the max camber point forward for the low angle of attack designs and backward for the high angle of attack designs. Finally, as noted in the previous section, the airfoil shape variations are nearly identical for the moment coefficient, regardless of the angle of attack. This analysis demonstrates how these surrogates and sensitivities can be utilized to inform more sophisticated and nuanced airfoil designs.

V. Conclusion

In this work, we provided the first in-depth sensitivity study of a novel airfoil parametrization with regard to critical aerodynamic quantities. This airfoil parametrization, obtained using the separable shape tensors framework, provides unique benefits to the design process: (i) it isolates various well-studied shape characteristics, e.g., airfoil thickness, and (ii) provides a low-dimensional and well-regularized parameter domain that is well-suited for analysis using ML and AI algorithms. The separable shape tensors approach is data-driven and relies on provided data to learn the shape parametrization. Thus, we constructed a curated dataset of existing airfoil shapes to inform the data-driven parameter space for airfoil parametrization. Both the curated dataset and corresponding parameter space were made available online [83]. In addition, we performed a dimensionality study and demonstrated how to identify the optimal trade-off between the expressivity and regularity of the new parametrization.

For our sensitivity analysis study, we used the `HAM2D` CFD solver to compute aerodynamic properties for the generated airfoil shapes at a low and high angle of attack. This data has been made available through the Open Energy Data Initiative [86]. We then computed multiple sensitivity metrics to quantify the impacts of the various shape parameters on the lift, drag, and moment coefficients. These quantities are critical for the design and optimization of airfoils and other aerodynamic shapes. Thus, this sensitivity analysis highlighted the key insights into the use of separable shape tensors for shape design. We found that these quantities were primarily driven by two PGA parameters: the leading PGA parameter and a second PGA parameter, the identity of which varied depending on the specific quantity and angle of attack. The sensitivities were relatively consistent across all of the computed metrics. We then explored the

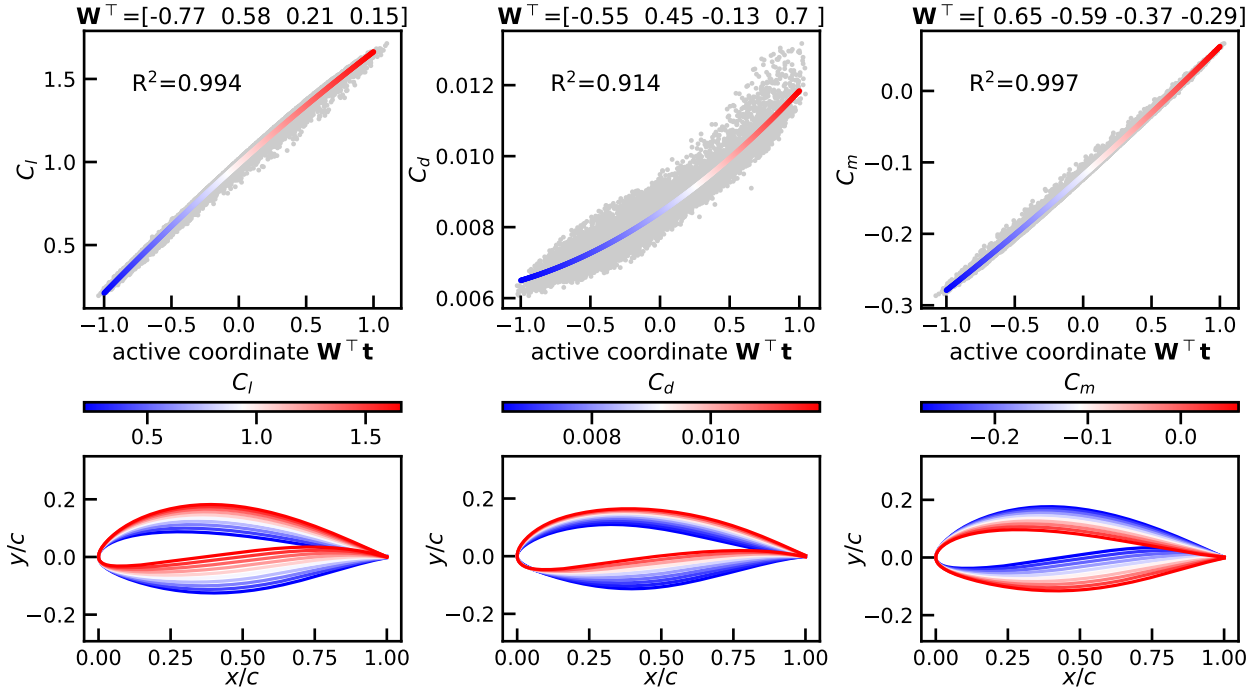

(a) $\alpha = 4°$.

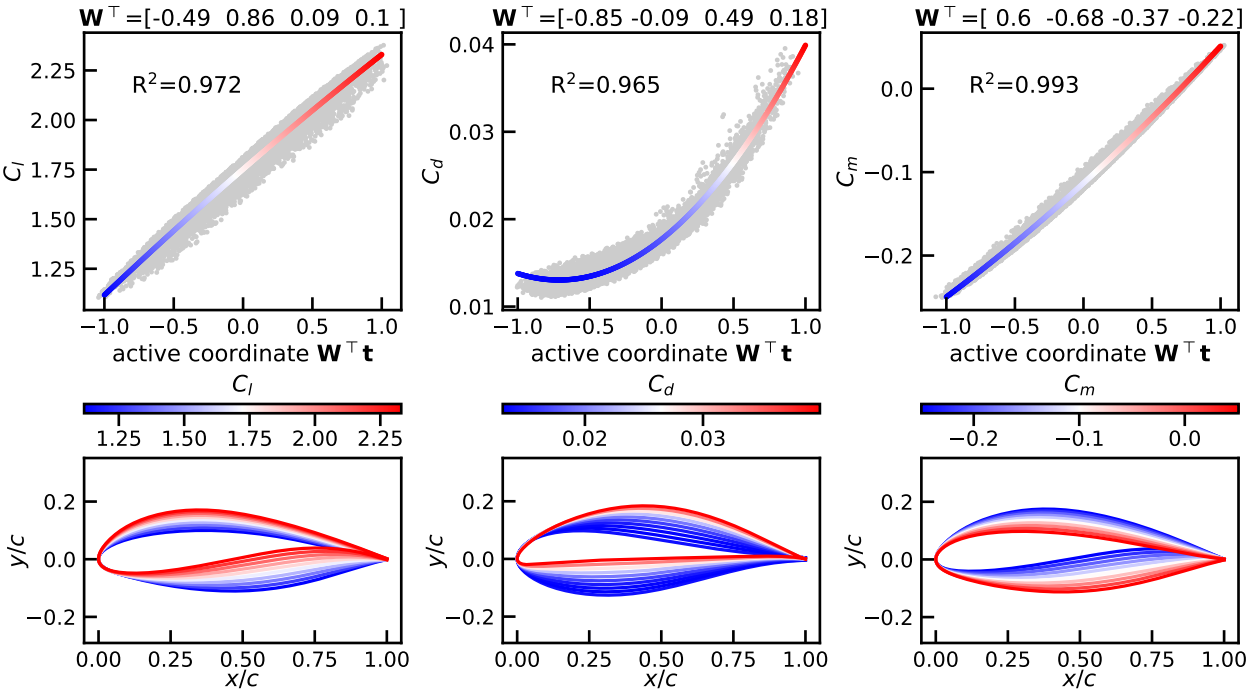

(b) $\alpha = 12°$.

Fig. 10 One-dimensional ridge approximations for aerodynamic responses (top rows) and airfoil shape sweeps along the corresponding subspace (lower rows). The color corresponds to the value of the predicted aerodynamic response.

physical interpretation of these sensitivities by investigating the airfoil deformations obtained by sweeping across the parameters identified using the subspace-based approach. We found one-dimensional subspaces in airfoil parameter space that allow us to deform the airfoils in a manner that maximally varies the aerodynamic quantities. These showed specific shape variations that drive changes in performance and where these variations should occur. Such insights can inform manufacturing and airfoil designs for various applications by identifying shape variations that have either small or large impacts on performance.

This work provides an initial investigation of how the separable shape tensors framework can help to inform airfoil design processes. One key aspect of separable shape tensors not explored in this work is their natural extension to three-dimensional wind and blade shape designs. Such blade shapes can often be represented using interpolations of a series of airfoil shapes at various spanwise locations. By controlling large-scale shape variations (e.g., thickness and rotation), separable shape tensors can represent a cohesive family of airfoils using limited PGA coefficients whose interpolation avoids undesirable features, such as divots or kinks. Further, the blade shape can be manipulated using parallel translation of the PGA parameters to apply consistent deformations to all of the component airfoils. Sections 2.5 and 3.2 of Grey et al. [31] provide a detailed discussion of this extension to three dimensions. Future sensitivity analysis work will seek to characterize design impacts for the 3D aerodynamic structures.

Acknowledgments

We are grateful to Michael Quayle for providing airfoil data through access to the BigFoil database for this work. This work was authored in part by the National Renewable Energy Laboratory (NREL), operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. This work was supported by the Laboratory Directed Research and Development (LDRD) Program at NREL and the AI4WIND project sponsored by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. A portion of this research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. This work is U.S. Government work and not protected by U.S. copyright. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

References

[1] Grasso, F., "Development of thick airfoils for wind turbines," *Journal of Aircraft*, Vol. 50, No. 3, 2013, pp. 975–981. https://doi.org/10.2514/1.C032182.

[2] Zhu, W. J., Shen, W. Z., and Sørensen, J. N., "Integrated airfoil and blade design method for large wind turbines," *Renewable Energy*, Vol. 70, 2014, pp. 172–183. https://doi.org/10.1016/j.renene.2014.02.057.

[3] Zhang, M., and Rizzi, A. W., "Assessment of Inverse and Direct Methods for Airfoil and Wing Design," *Simulation-Driven Modeling and Optimization: ASDOM, Reykjavik, August 2014*, Springer, 2016, pp. 75–109. https://doi.org/10.1007/978-3-319-27517-8_4.

[4] Zhang, Y., Fang, X., Chen, H., Fu, S., Duan, Z., and Zhang, Y., "Supercritical natural laminar flow airfoil optimization for regional aircraft wing design," *Aerospace Science and Technology*, Vol. 43, 2015, pp. 152–164. https://doi.org/10.1016/j.ast.2015.02.024.

[5] Smith, M. J., Potsdam, M., Wong, T.-C., Baeder, J. D., and Phanse, S., "Evaluation of Computational Fluid Dynamics to Determine Two-Dimensional Airfoil Characteristics for Rotorcraft Applications," *Journal of the American Helicopter Society*, Vol. 51, No. 1, 2006, pp. 70–79. https://doi.org/10.4050/1.3092879.

[6] Kim, J.-H., Choi, J.-H., Husain, A., and Kim, K.-Y., "Performance enhancement of axial fan blade through multi-objective optimization techniques," *Journal of Mechanical Science and Technology*, Vol. 24, 2010, pp. 2059–2066. https://doi.org/10.1007/s12206-010-0619-6.

[7] Lin, S., and Tsai, M., "An integrated study of the design method for small axial-flow fans, based on the airfoil theory," *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, Vol. 225, No. 4, 2011, pp. 885–895. https://doi.org/10.1243/09544062JMES2231.

[8] Sobieszczanski-Sobieski, J., and Haftka, R. T., "Multidisciplinary aerospace design optimization: survey of recent developments," *Structural optimization*, Vol. 14, 1997, pp. 1–23. https://doi.org/10.1007/BF01197554.

[9] Sadraey, M. H., *Aircraft Design: A Systems Engineering Approach*, John Wiley & Sons, 2012. https://doi.org/10.1002/9781118352700.

[10] Martins, J. R., "Aerodynamic design optimization: Challenges and perspectives," *Computers & Fluids*, Vol. 239, 2022, p. 105391. https://doi.org/10.1016/j.compfluid.2022.105391.

[11] Anand, A., Marepally, K., Muneeb Safdar, M., and Baeder, J. D., "A novel approach to inverse design of wind turbine airfoils using tandem neural networks," *Wind Energy*, 2024. https://doi.org/10.1002/we.2918.

[12] Mozafari, M., Sadeghimalekabadi, M., Fardi, A., Bruecker, C., and Masdari, M., "Aeroacoustic investigation of a ducted wind turbine employing bio-inspired airfoil profiles," *Physics of Fluids*, Vol. 36, No. 4, 2024. https://doi.org/10.1063/5.0204050.

[13] Masters, D. A., Taylor, N. J., Rendall, T., Allen, C. B., and Poole, D. J., "Geometric comparison of aerofoil shape parameterization methods," *AIAA Journal*, Vol. 55, No. 5, 2017, pp. 1575–1589. https://doi.org/10.2514/1.J054943.

[14] Hicks, R. M., and Henne, P. A., "Wing design by numerical optimization," *Journal of Aircraft*, Vol. 15, No. 7, 1978, pp. 407–412. https://doi.org/10.2514/3.58379.

[15] Venkataraman, P., "A new procedure for airfoil definition," *13th Applied Aerodynamics Conference*, 1995, p. 1875. https://doi.org/10.2514/6.1995-1875.

[16] Morris, A. M., Allen, C. B., and Rendall, T. C. S., "CFD-based optimization of aerofoils using radial basis functions for domain element parameterization and mesh deformation," *International Journal for Numerical Methods in Fluids*, Vol. 58, No. 8, 2008, pp. 827–860. https://doi.org/10.1002/fld.1769.

[17] Braibant, V., and Fleury, C., "Shape optimal design using B-splines," *Computer Methods in Applied Mechanics and Engineering*, Vol. 44, No. 3, 1984, pp. 247–267. https://doi.org/10.1016/0045-7825(84)90132-4.

[18] Sobieczky, H., "Parametric airfoils and wings," *Recent Development of Aerodynamic Design Methodologies: Inverse Design and Optimization*, Springer, 1999, pp. 71–87. https://doi.org/10.1007/978-3-322-89952-1_4.

[19] Kulfan, B., and Bussoletti, J., ""Fundamental" Parameteric Geometry Representations for Aircraft Component Shapes," *11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2006, p. 6948. https://doi.org/doi.org/10.2514/6.2006-6948.

[20] Kulfan, B. M., "Universal parametric geometry representation method," *Journal of aircraft*, Vol. 45, No. 1, 2008, pp. 142–158. https://doi.org/10.2514/1.29958.

[21] Toal, D., Bressloff, N., Keane, A., and Holden, C., "Geometric Filtration Using Proper Orthogonal Decomposition for Aerodynamic Design Optimization," *AIAA Journal*, Vol. 48, No. 5, 2010, pp. 916–928. https://doi.org/10.2514/1.41420.

[22] Poole, D., Allen, C., and Rendall, T., "Metric-Based Mathematical Derivation of Efficient Airfoil Design Variables," *AIAA Journal*, Vol. 53, 2015, pp. 1349–1361. https://doi.org/10.2514/1.J053427.

[23] Li, J., Bouhlel, M. A., and Martins, J. R. R. A., "Data-Based Approach for Fast Airfoil Analysis and Optimization," *AIAA Journal*, Vol. 57, No. 2, 2019, pp. 581–596. https://doi.org/10.2514/1.J057129.

[24] Li, J., Du, X., and Martins, J. R. R. A., "Machine Learning in Aerodynamic Shape Optimization," *Progress in Aerospace Sciences*, Vol. 134, 2022, p. 100849. https://doi.org/10.1016/j.paerosci.2022.100849.

[25] Kedward, L. J., Allen, C. B., Poole, D. J., and Rendall, T. C. S., "Generic Modal Design Variables for Efficient Aerodynamic Optimization," *AIAA Journal*, Vol. 61, No. 2, 2023, pp. 739–755. https://doi.org/10.2514/1.J061727.

[26] Kedward, L. J., Allen, C. B., and Rendall, T. C. S., "Gradient-Limiting Shape Control for Efficient Aerodynamic Optimization," *AIAA Journal*, Vol. 58, No. 9, 2020, pp. 3748–3764. https://doi.org/10.2514/1.J058977.

[27] Sekar, V., Zhang, M., Shu, C., and Khoo, B. C., "Inverse design of airfoil using a deep convolutional neural network," *AIAA Journal*, Vol. 57, No. 3, 2019, pp. 993–1003. https://doi.org/10.2514/1.J057894.

[28] Glaws, A., King, R. N., Vijayakumar, G., and Ananthan, S., "Invertible neural networks for airfoil design," *AIAA journal*, Vol. 60, No. 5, 2022, pp. 3035–3047. https://doi.org/10.2514/1.J060866.

[29] Wang, Y., Shimada, K., and Barati Farimani, A., "Airfoil GAN: encoding and synthesizing airfoils for aerodynamic shape optimization," *Journal of Computational Design and Engineering*, Vol. 10, No. 4, 2023, pp. 1350–1362. https://doi.org/10.1093/jcde/qwad046.

[30] Wei, Z., Yang, A., Li, J., Bauerheim, M., Liem, R. P., and Fua, P., "DeepGeo: Deep Geometric Mapping for Automated and Effective Parameterization in Aerodynamic Shape Optimization," *AIAA AVIATION Forum*, 2024. https://doi.org/10.2514/6.2024-3839.

[31] Grey, Z. J., Doronina, O. A., and Glaws, A., "Separable shape tensors for aerodynamic design," *Journal of Computational Design and Engineering*, Vol. 10, No. 1, 2023, pp. 468–487. https://doi.org/10.1093/jcde/qwac140.

[32] Yataka, R., Hirashima, K., and Shiraishi, M., "Grassmann manifold flows for stable shape generation," *Advances in Neural Information Processing Systems*, Vol. 36, 2023, pp. 72377–72411.

[33] Zhang, Y., Pang, B., Li, X., and Chen, G., "Aerodynamic shape optimization with Grassmannian shape parameterization method," *Energies*, Vol. 15, No. 20, 2022, p. 7722. https://doi.org/10.3390/en15207722.

[34] Jasa, J., Glaws, A., Bortolotti, P., Vijayakumar, G., and Barter, G., "Wind turbine blade design with airfoil shape control using invertible neural networks," *Journal of Physics: Conference Series*, Vol. 2265, IOP Publishing, 2022, p. 042052. https://doi.org/10.1088/1742-6596/2265/4/042052.

[35] Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., and Tarantola, S., *Global Sensitivity Analysis: The Primer*, John Wiley & Sons, 2008. https://doi.org/10.1002/9780470725184.

[36] Miro, S., Hartmann, D., and Schanz, T., "Global sensitivity analysis for subsoil parameter estimation in mechanized tunneling," *Computers and Geotechnics*, Vol. 56, 2014, pp. 80–88. https://doi.org/10.1016/j.compgeo.2013.11.003.

[37] Cheng, Q., Zhao, H., Zhang, G., Gu, P., and Cai, L., "An analytical approach for crucial geometric errors identification of multi-axis machine tool based on global sensitivity analysis," *The International Journal of Advanced Manufacturing Technology*, Vol. 75, 2014, pp. 107–121. https://doi.org/10.1007/s00170-014-6133-8.

[38] Ellinger, J., Semm, T., and Zaeh, M. F., "Dimensionality Reduction of High-Fidelity Machine Tool Models by Using Global Sensitivity Analysis," *Journal of Manufacturing Science and Engineering*, Vol. 144, No. 5, 2022, p. 051010. https://doi.org/10.1115/1.4052710.

[39] Ebrahimi-Moghadam, A., Ildarabadi, P., Aliakbari, K., and Fadaee, F., "Sensitivity analysis and multi-objective optimization of energy consumption and thermal comfort by using interior light shelves in residential buildings," *Renewable Energy*, Vol. 159, 2020, pp. 736–755. https://doi.org/10.1016/j.renene.2020.05.127.

[40] Grey, Z. J., Mosleh, S., Rezac, J. D., Ma, Y., Coder, J. B., and Dienstfrey, A. M., "Optimizing unlicensed band spectrum sharing with subspace-based Pareto tracing," *ICC 2021-IEEE International Conference on Communications*, IEEE, 2021, pp. 1–7. https://doi.org/10.1109/ICC42927.2021.9500533.

[41] Kolb, J., and Hameyer, K., "Sensitivity analysis of manufacturing tolerances in permanent magnet synchronous machines with stator segmentation," *IEEE Transactions on Energy Conversion*, Vol. 35, No. 4, 2020, pp. 2210–2221. https://doi.org/10.1109/TEC.2020.3017279.

[42] Liu, B., Liu, J., Yu, X., and An, G., "A novel decomposition method for manufacture variations and the sensitivity analysis on compressor blades," *Aerospace*, Vol. 9, No. 10, 2022, p. 542. https://doi.org/10.3390/aerospace9100542.

[43] Sobieszczanski-Sobieski, J., "The case for aerodynamic sensitivity analysis," *Sensitivity Analysis in Engineering*, 1987, pp. 77–96.

[44] Wu, X., Zhang, W., and Song, S., "Uncertainty quantification and sensitivity analysis of transonic aerodynamics with geometric uncertainty," *International Journal of Aerospace Engineering*, Vol. 2017, No. 1, 2017, p. 8107190. https://doi.org/10.1155/2017/8107190.

[45] Blonigan, P. J., Wang, Q., Nielsen, E. J., and Diskin, B., "Least-squares shadowing sensitivity analysis of chaotic flow around a two-dimensional airfoil," *AIAA Journal*, Vol. 56, No. 2, 2018, pp. 658–672. https://doi.org/10.2514/1.J055389.

[46] Liu, C., and Lee, S., "Parametric airfoil design and sensitivity analysis for turbulent boundary-layer trailing-edge noise reduction," *AIAA Journal*, Vol. 60, No. 4, 2022, pp. 2324–2341. https://doi.org/10.2514/1.J060096.

[47] PM, M. A. S., and Raj, L. P., "Sensitivity analysis of geometric parameters on the aerodynamic performance of a multi-element airfoil," *Aerospace Science and Technology*, Vol. 132, 2023, p. 108074. https://doi.org/10.1016/j.ast.2022.108074.

[48] Shum, J. G., and Lee, S., "Computational Study of Airfoil Design Parameters and Sensitivity Analysis for Dynamic Stall," *AIAA Journal*, Vol. 62, No. 4, 2024, pp. 1611–1617. https://doi.org/10.2514/1.J063474.

[49] Ohashi, M., Fukagata, K., and Tokugawa, N., "Adjoint-based sensitivity analysis for airfoil flow control aiming at lift-to-drag ratio improvement," *AIAA Journal*, Vol. 59, No. 11, 2021, pp. 4437–4448. https://doi.org/10.2514/1.J060415.

[50] Quick, J., King, R. N., de Frahan, M. T. H., Ananthan, S., Sprague, M., and Hamlington, P. E., "Field Sensitivity Analysis of Turbulence Model Parameters for Flow Over a Wing," *International Journal for Uncertainty Quantification*, Vol. 12, No. 1, 2022. https://doi.org/10.1615/Int.J.UncertaintyQuantification.2021036467.

[51] Robertson, A., Sethuraman, L., and Jonkman, J. M., "Assessment of wind parameter sensitivity on extreme and fatigue wind turbine loads," *2018 Wind Energy Symposium*, 2018, p. 1728. https://doi.org/10.2514/6.2018-1728.

[52] Velarde, J., Kramhøft, C., and Sørensen, J. D., "Global sensitivity analysis of offshore wind turbine foundation fatigue loads," *Renewable energy*, Vol. 140, 2019, pp. 177–189. https://doi.org/10.1016/j.renene.2019.03.055.

[53] Grey, Z. J., and Constantine, P. G., "Active subspaces of airfoil shape parameterizations," *AIAA Journal*, Vol. 56, No. 5, 2018, pp. 2003–2017. https://doi.org/10.2514/6.2017-0507.

[54] Doronina, O., Grey, Z., and Glaws, A., "Grassmannian Shape Representations for Aerodynamic Applications," *AAAI 2022 Workshop on AI for Design and Manufacturing (ADAM)*, 2022, pp. 1–5.

[55] Pennec, X., "Probabilities and statistics on Riemannian manifolds: Basic tools for geometric measurements," *NSIP*, Vol. 3, 1999, pp. 194–198. URL https://hal.inria.fr/inria-00615833.

[56] Fletcher, P., Lu, C., and Joshi, S., "Statistics of shape via principal geodesic analysis on Lie groups," *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, Vol. 1, 2003, pp. I–I. https://doi.org/10.1109/CVPR.2003.1211342.

[57] Glaws, A., Hokanson, J., King, R., and Vijayakumar, G., "Regularizing invertible neural networks for airfoil design through dimension reduction," *AIAA SCITECH 2022 Forum*, 2022, p. 1098. https://doi.org/10.2514/6.2022-1098.

[58] Jolliffe, I., "Principal component analysis," *Encyclopedia of Statistics in Behavioral Science*, 2005. https://doi.org/10.1002/0470013192.bsa501.

[59] Bendokat, T., Zimmermann, R., and Absil, P.-A., "A Grassmann manifold handbook: Basic geometry and computational aspects," *Advances in Computational Mathematics*, Vol. 50, No. 1, 2024, p. 6. https://doi.org/10.1007/s10444-023-10090-8.

[60] Ceze, M., Hayashi, M., and Volpe, E., "A study of the CST parameterization characteristics," *27th AIAA applied aerodynamics conference*, 2009, p. 3767. https://doi.org/10.2514/6.2009-3767.

[61] Lane, K., and Marshall, D., "Inverse airfoil design utilizing CST parameterization," *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2010, p. 1228. https://doi.org/10.2514/6.2010-1228.

[62] Marques, S., and Hewitt, P., "Aerofoil optimisation using CST parameterisation in SU2," *Royal Aeronautical Society Applied Aerodynamics Group Conference*, 2014.

[63] Jung, Y. S., Govindarajan, B., and Baeder, J., "Turbulent and Unsteady Flows on Unstructured Grids using Line-Based Hamiltonian Paths and Strands," *AIAA Journal*, Vol. 55, No. 6, 2017, pp. 1986–2001. https://doi.org/10.2514/1.J055141.

[64] Jiang, G.-S., and Shu, C.-W., "Efficient Implementation of Weighted ENO Schemes," *Journal of Computational Physics*, Vol. 126, No. 1, 1996, pp. 202–228. https://doi.org/10.1006/jcph.1996.0130.

[65] Roe, P. L., "Approximate Riemann solvers, parameter vectors, and difference schemes," *Journal of Computational Physics*, Vol. 43, No. 2, 1981, pp. 357–372. https://doi.org/10.1016/0021-9991(81)90128-5.

[66] Behr, M., and Tezduyar, T., "Finite element solution strategies for large-scale flow simulations," *Computer Methods in Applied Mechanics and Engineering*, Vol. 112, No. 1, 1994, pp. 3–24. https://doi.org/10.1016/0045-7825(94)90016-7.

[67] Spalart, P. R., and Allmaras, S. R., "A One-Equation Turbulence Model for Aerodynamics Flow," *La Recherche Aerospatiale 1*, Vol. 1, No. 3, 1994, pp. 5–21. https://doi.org/10.2514/6.1992-439.

[68] Medida, S., "Correlation-based Transition Modeling for External Aerodynamic Flows," Ph.D. thesis, University of Maryland, 2014.

[69] Jung, Y. S., Vijayakumar, G., Ananthan, S., and Baeder, J., "Local correlation-based transition models for high-Reynolds-number wind-turbine airfoils," *Wind Energy Science*, Vol. 7, No. 2, 2022, pp. 603–622. https://doi.org/10.5194/wes-7-603-2022.

[70] Cacuci, D. G., *Sensitivity & Uncertainty Analysis, Volume 1*, Chapman and Hall/CRC, 2003, Chap. Local Sensitivity and Uncertainty Analysis of Nonlinear Systems. https://doi.org/10.1201/9780203498798.

[71] Sobol, I. M., "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates," *Mathematics and computers in simulation*, Vol. 55, No. 1-3, 2001, pp. 271–280. https://doi.org/10.1016/S0378-4754(00)00270-6.

[72] Constantine, P. G., and Diaz, P., "Global sensitivity metrics from active subspaces," *Reliability Engineering & System Safety*, Vol. 162, 2017, pp. 1–13. https://doi.org/10.1016/j.ress.2017.01.013.

[73] Constantine, P. G., *Active subspaces: Emerging ideas for dimension reduction in parameter studies*, SIAM, 2015.

[74] Hokanson, J. M., and Constantine, P. G., "Data-driven polynomial ridge approximation using variable projection," *SIAM Journal on Scientific Computing*, Vol. 40, No. 3, 2018, pp. A1566–A1589. https://doi.org/10.1137/17M1117690.

[75] Constantine, P. G., Zaharatos, B., and Campanelli, M., "Discovering an active subspace in a single-diode solar cell model," *Statistical Analysis and Data Mining: The ASA Data Science Journal*, Vol. 8, No. 5-6, 2015, pp. 264–273. https://doi.org/10.1002/sam.11281.

[76] Constantine, P. G., Eftekhari, A., Hokanson, J., and Ward, R. A., "A near-stationary subspace for ridge approximation," *Computer Methods in Applied Mechanics and Engineering*, Vol. 326, 2017, pp. 402–421. https://doi.org/10.1016/j.cma.2017.07.038.

[77] Quayle, M., "BigFoil.com Airfoil Database," , 2022. https://www.bigfoil.com.

[78] Selig, M. S., "UIUC airfoil data site," , 1996. URL m-selig.ae.illinois.edu/ads/coord_database.html.

[79] Hepperle, M., "JavaFoil – Analysis of Airfoils," , 2018. URL https://www.mh-aerotools.de/airfoils/javafoil.htm.

[80] Abbott, I. H., Von Doenhoff, A. E., and Stivers Jr, L., "Summary of Airfoil Data," Tech. Rep. NACA-TR-824, National Advisory Committee for Aeronautics, Langley Field, VA, 1945.

[81] Jonkman, J., Butterfield, S., Musial, W., and Scott, G., "Definition of a 5-MW reference wind turbine for offshore system development," Tech. Rep. NREL/TP-500-38060, National Renewable Energy Lab. (NREL), Golden, CO (United States), 2009. https://doi.org/10.2172/947422.

[82] Gaertner, E., Rinker, J., Sethuraman, L., Zahle, F., Anderson, B., Barter, G., Abbas, N., Meng, F., Bortolotti, P., Skrzypinski, W., Scott, G., Feil, R., Bredmose, H., Dykes, K., Sheilds, M., Allen, C., and Viselli, A., "IEA Wind TCP Task 37: Definition of the IEA 15-Megawatt Offshore Reference Wind Turbine," Tech. Rep. NREL/TP-5000-75698, National Renewable Energy Lab. (NREL), Golden, CO (United States), 2020. https://doi.org/10.2172/1603478.

[83] Doronina, O., Glaws, A., Grey, Z., and Lee, B., "G2Aero Database of Airfoils - Curated Airfoils," , 2024. https://doi.org/10.25984/2448331, URL https://data.openei.org/submissions/6198.

[84] Fletcher, P. T., Lu, C., and Joshi, S., "Statistics of shape via principal geodesic analysis on Lie groups," *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, Vol. 1, IEEE, 2003, pp. I–I. https://doi.org/10.1109/CVPR.2003.1211342, URL https://doi.org/10.1109/CVPR.2003.1211342.

[85] Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A., "A kernel two-sample test," *The Journal of Machine Learning Research*, Vol. 13, No. 1, 2012, pp. 723–773.

[86] Ramos, D., Glaws, A., King, R., Lee, B., Doronina, O., Baeder, J., Vijayakumar, G., and Grey, Z., "Airfoil Computational Fluid Dynamics - 9k shapes, 2 AoA's," , 2023. https://doi.org/10.25984/2222587.

[87] Herman, J., and Usher, W., "SALib: An open-source Python library for sensitivity analysis," *Journal of Open Source Software*, Vol. 2, No. 9, 2017, p. 97. https://doi.org/10.21105/joss.00097.

[88] Constantine, P. G., Howard, R., Salinger, A. G., Grey, Z., Diaz, P., and Fletcher, L., "Python Active-subspaces Utility Library," *J. Open Source Softw.*, Vol. 1, No. 5, 2016, p. 79. https://doi.org/10.21105/joss.00079.

[89] Hokanson, J. M., "PSDR: Parameter Space Dimension Reduction Toolbox," https://github.com/jeffrey-hokanson/PSDR, 2019.

[90] Glaws, A., Constantine, P. G., and Cook, R. D., "Inverse regression for ridge recovery: a data-driven approach for parameter reduction in computer experiments," *Statistics and Computing*, Vol. 30, No. 2, 2020, pp. 237–253. https://doi.org/10.1007/s11222-019-09876-y.

[91] Constantine, P. G., Emory, M., Larsson, J., and Iaccarino, G., "Exploiting active subspaces to quantify uncertainty in the numerical simulation of the HyShot II scramjet," *Journal of Computational Physics*, Vol. 302, 2015, pp. 1–20. https://doi.org/10.1016/j.jcp.2015.09.001.

[92] Sethuraman, L., Glaws, A., Skinner, M., and Parans Paranthaman, M., "Advanced multimaterial shape optimization methods as applied to advanced manufacturing of wind turbine generators," *Wind Energy*, 2024. https://doi.org/10.1002/we.2911.