AIAA SciTech Forum
8–12 January 2018, Kissimmee, Florida
2018 AIAA Aerospace Sciences Meeting

10.2514/6.2018-0274

Dimension Reduction for Shape Design Insight

Zachary del Rosario*, Aaron Towne†, Gianluca Iaccarino‡
Stanford University, Stanford, CA, 94305, USA

Conceptual design requires conceptual understanding, what we call *design insight*. Novel designs require novel insights. Such understanding is usually hard-won through difficult analysis. We propose *algorithmic* approaches to developing insight, utilizing dimension reduction techniques. This work generalizes techniques for scalar quantities to entire fields, enabling the identification of low-dimensional structure of more general objects. We present applications to aerodynamic and structural shape design.

Nomenclature

Dimension Reduction

$\cdot_a$	Active subscript
$\cdot_i$	Inactive subscript
$\mathbb{E}[\cdot]$	Expectation
$\hat{\cdot}$	Sample mean
$\langle \cdot, \cdot \rangle$	Inner product
$\boldsymbol{C}$	Average outer product of gradients
$\boldsymbol{R}$	Correlation tensor
$\boldsymbol{\Lambda}$	Eigenvalues
$\boldsymbol{W}$	Eigenvectors
$\boldsymbol{\phi}$	Proper Orthogonal Decomposition mode
$\boldsymbol{y}$	Active variables
$\boldsymbol{z}$	Inactive variables
k	Dimensionality of reduced space

Design Problem

$\mathcal{X}$	Input parameter space
$\boldsymbol{W}$	Matrix weight function
∇f	Gradient
ρ	Weight function
$\boldsymbol{s}$	Independent parameters
$\boldsymbol{\xi}$	Stochastic parameters
f	Quantity of interest
m	Dimensionality of input space
n	Number of samples

I. Introduction

Original designs require new knowledge. Pahl and Beitz[1] define one such case of new knowledge as being "often based on the application of the latest scientific knowledge and insights." In this work, we seek such *design insight* – deep understanding of a physical system which aids the design process. Such insight is a key part of developing new solutions, and defining procedures and techniques to aid this development is useful to accelerate design.

A classical technique is the so-called 'method of concentrated staring'.[2] While time-honored, such a procedure is not guaranteed to converge in finite time. Contemporary efforts have succeeded in deriving *scientific insights* – such as governing ordinary differential equations[3] and conservation laws[4] – from data. Closer to engineering practice, recent advances in fluid dynamics/turbulence consider identifying flow features using solution data, perhaps most famously the Proper Orthogonal Decomposition (POD).[5] While POD is usually employed to understand the output of a dynamical system, the active subspace (AS) is used to understand the input space, and identify low-dimensional structure. Our previous work considered the active subspace as it relates to physical problems, with a particular eye towards developing insights suitable for engineering design.[6] In this work we further develop these ideas, drawing connections between POD and the AS in order to handle more general problems of interest. In particular, we seek to handle spatial and temporal variation of some quantity of interest (qoi), performing dimension reduction over a set of stochastic variables, and interpreting the resulting structure in a physical context.

*PhD Candidate, Aeronautics and Astronautics, AIAA Student Member
†Postdoctoral Scholar, Center for Turbulence Research, AIAA Member
‡Professor, Department of Mechanical Engineering, AIAA Associate Fellow

Section II reviews the active subspace, and Section III reviews the Proper Orthogonal Decomposition. Section IV draws connections between active subspaces and POD, informing a dimension reduction technique which considers variation in additional variables. Sections V and VI demonstrate this approach (respectively) on the design of airfoils and structural members. Section VII provides some concluding remarks.

II. Active Subspaces

This section summarizes and motivates the active subspace.[7] In what follows, unbolded symbols denote scalars, bolded lowercase letters denote vectors, bolded uppercase letters are matrices, and caligraphic characters represent sets. Suppose we have some qoi $f(\boldsymbol{\xi})$ where $\boldsymbol{\xi}$ takes values in a domain $\boldsymbol{\xi} \in \mathcal{X} \subseteq \mathbb{R}^m$. As m increases, the computational cost of studying our qoi grows exponentially. This *curse of dimensionality* is the motivating issue for active subspaces. Informally, the active subspace yields a linear subspace decomposition of $\mathbb{R}^m$, split into *active* and *inactive* directions. The idea is to either ignore or treat with lower fidelity the inactive directions, and retain the active ones. This effectively reduces the dimension of the input space that must be considered.

Formally, we assume the gradient of our qoi $\nabla_{\boldsymbol{\xi}} f$ exists, and define a weight function $\rho(\boldsymbol{\xi})$. This ρ may represent a probabilistic description of the input variables (i.e. a joint density function), or simply a function which weights regions of parameter space according to the designer's interest[a]. In principle $\mathcal{X}$ may be unbounded; in the cases below we restrict attention to finite hyperrectangle bounds $\mathcal{X} = \{\boldsymbol{\xi} \,|\, \boldsymbol{\xi}_l \leq \boldsymbol{\xi} \leq \boldsymbol{\xi}_u\}$. We then compute the following matrix

$$\boldsymbol{C} = \int_{\mathbb{R}^m} \nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}) \nabla_{\boldsymbol{\xi}}^T f(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) \, d\boldsymbol{\xi}, \tag{1}$$

where $d\boldsymbol{\xi} = \prod_{i=1}^m d\xi_i$. Since $\boldsymbol{C}$ is symmetric positive-semidefinite, it admits an eigenvalue decomposition $\boldsymbol{C} = \boldsymbol{W}\boldsymbol{\Lambda}\boldsymbol{W}^T$. We then choose a decomposition based on the eigenvalues of $\boldsymbol{C}$, that is

$$\boldsymbol{W} = [\boldsymbol{W}_a \, \boldsymbol{W}_i], \; \boldsymbol{\Lambda} = \begin{bmatrix} \boldsymbol{\Lambda}_a & \\ & \boldsymbol{\Lambda}_i \end{bmatrix} \tag{2}$$

where $\boldsymbol{\Lambda}_a = \text{diag}(\lambda_1, \ldots, \lambda_k)$ with $\lambda_i \geq \lambda_{i+1}$ for all $i \in \{1, \ldots, m-1\}$ and $k < m$. The columns of $\boldsymbol{W}_a$ are the first k eigenvectors of $\boldsymbol{C}$. With a decomposition of $\boldsymbol{C}$ chosen, we define the active ($\boldsymbol{y}$) and inactive ($\boldsymbol{z}$) variables

$$\boldsymbol{y} = \boldsymbol{W}_a^T \boldsymbol{\xi} \in \mathbb{R}^k, \; \boldsymbol{z} = \boldsymbol{W}_i^T \boldsymbol{\xi} \in \mathbb{R}^{m-k}. \tag{3}$$

In this way, the active directions $\boldsymbol{W}_a$ define linear combinations of the input parameters. One can then treat the Inactive Variables $\boldsymbol{z}$ with lower fidelity, focusing computational effort on the active directions. As an illustrative example, consider the function $f(\boldsymbol{\xi}) = \frac{1}{2}(.7\xi_1 + .3\xi_2)^2$. Clearly, the quantity of interest f changes along the direction $[.7, .3]^T$, and does not change at all along $[.3, -.7]^T$ (Fig. 1). A procedure which estimates the AS will identify these directions, both allowing us to reduce the dimensionality of f, and providing some insight into how changes in parameter space contribute to variability in f.

Previous work has found a low-dimensional active subspace in a number of different scientific and engineering applications, including scramjets,[8] hydrologic models,[9] and car aerodynamics.[10] Our previous work investigated the connections between aerodynamic shape design and the active subspace, which is briefly summarized below.[6]

A. Example: Transonic airfoil design

Here we consider the effect of shape perturbations to an airfoil, and their effects on the associated lift coefficient. We begin with a baseline NACA0012 geometry and modify the geometry through the Free Form Deformation strategy.[11] Our qoi is the lift coefficient, and our design variables are the vertical displacement of 38 control points. We define reasonable hypercube bounds on the design variables, and approximate Equation 1 by an estimated $\hat{\boldsymbol{C}}$ matrix computed via Monte Carlo sampling ($n = 10^3$). Figure 2 presents diagnostic results to assess the estimator of $\hat{\boldsymbol{C}}$, while Figure 3 summarizes the insights derived from the estimated active subspace.

[a]Subject to the constraint $\rho(\boldsymbol{\xi}) > 0$ for $\boldsymbol{\xi} \in \mathcal{X}$, and $\rho = 0$ for $\boldsymbol{\xi} \notin \mathcal{X}$.

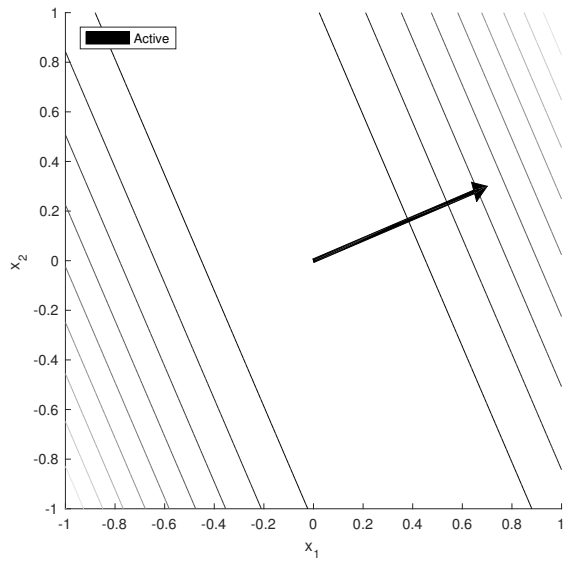

Figure 1: Isolines depicting sets of constant $f(x_1, x_2)$, with an arrow depicting the corresponding active direction. Note that on average, the function changes more along the active directions than the inactive directions; in this case, the change in the inactive direction is exactly zero.

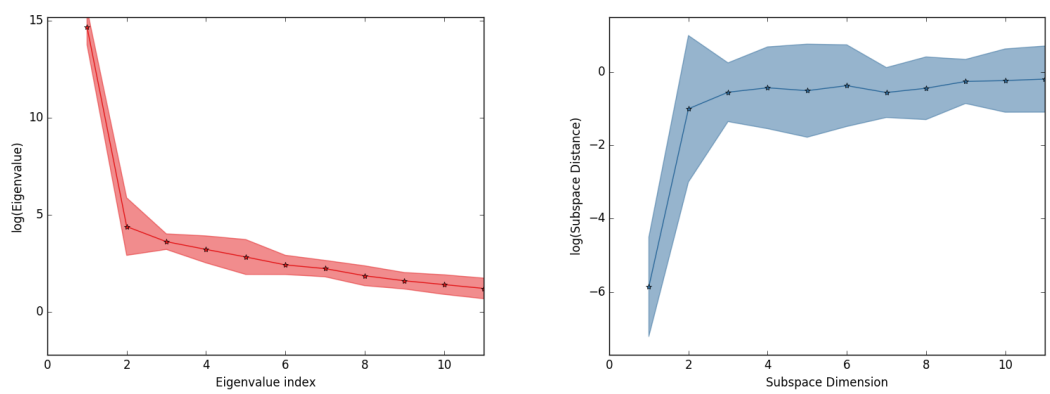

Figure 2: Bootstrap intervals of the eigenvalues of $\hat{\boldsymbol{C}}$ and subspace distance for NACA0012 test case. The left image shows bootstrap intervals of the eigenvalues, while the right displays those of the subspace distance. Stars denote the estimates derived from the entire sample set, while curves and shaded bounds show the bootstrap mean and bounds at one standard error, respectively. The dimension of the active subspace is chosen by seeking eigenvalue gaps; thus, the spectrum of $\hat{\boldsymbol{C}}$ suggests a one-dimensional active subspace. The subspace distance is computed between the active subspace estimate $\hat{\boldsymbol{W}}_a$ and a bootstrap estimate of the same quantity. This quantity is bounded above by 1; the smaller this value, the less variable our estimate $\hat{\boldsymbol{W}}_a$.

Figure 3 shows that in our example, the leading active subspace direction corresponds to the camber line of the airfoil. Since the active subspace is one-dimensional, this implies that the camber line accounts for the majority of variability of the lift coefficient, among the whole space of possible shape perturbations. The dependence of the lift coefficient on the camber line is already well known from Thin Airfoil Theory.[12] However, note that the conclusion above is derived not from an analytic approach, but *through an algorithm operating on data*. Such data-driven approaches have the potential to supplement a designer's intuition and suggest novel insights, especially in the context of complicated or less-studied systems.

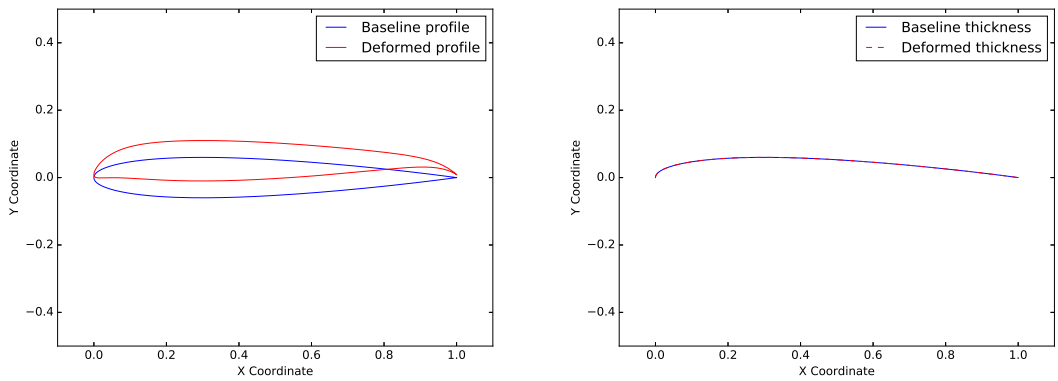

Figure 3: Comparison of baseline and perturbed geometry. The left image compares the baseline geometry against the airfoil deformed along the leading active subspace direction. This deformation corresponds to the camber line, as evidenced by the thickness distributions plotted in the right image. Since an airfoil may be uniquely decomposed into a linear combination of its thickness and camber, and since both profiles have the same thickness distribution, we conclude the deformation is modifying the camber of the airfoil alone. Taking the L1 norm of the difference between thickness distributions yields 5.6948×10^{-7}.

III. Proper Orthogonal Decomposition

Proper orthogonal decomposition is a standard technique used to extract energetic structures from stochastic data.[5] The basic objective is to find a deterministic function $\boldsymbol{\phi}(\boldsymbol{s})$ that best approximates a stochastic function $\boldsymbol{u}(\boldsymbol{\xi}, \boldsymbol{s})$ on average. Here, $\boldsymbol{s}$ is a set of independent variables and $\boldsymbol{\xi}$ parameterizes the stochastic function. This goal is accomplished by maximizing the expected value of the normalized projection of the stochastic function onto the deterministic one,[5] i.e.,

$$\mathbb{E}\left[|\langle \boldsymbol{u}(\boldsymbol{\xi}, \boldsymbol{s}), \boldsymbol{\phi}(\boldsymbol{s})\rangle|^2\right] / \langle \boldsymbol{\phi}(\boldsymbol{s}), \boldsymbol{\phi}(\boldsymbol{s})\rangle, \tag{4}$$

where $\langle \cdot, \cdot \rangle$ is an inner product on the space of $\boldsymbol{u}$ and $\mathbb{E}[\cdot]$ is the expectation operator over the probability space.

The function $\boldsymbol{\phi}(\boldsymbol{z})$ that maximizes (4) must satisfy the eigenvalue problem

$$\langle \boldsymbol{R}(\boldsymbol{s}, \boldsymbol{s}'), \boldsymbol{\phi}^*(\boldsymbol{s}')\rangle = \lambda \boldsymbol{\phi}(\boldsymbol{s}), \tag{5}$$

where

$$\boldsymbol{R}(\boldsymbol{s}, \boldsymbol{s}') = \mathbb{E}\left[\boldsymbol{u}(\boldsymbol{\xi}, \boldsymbol{s})\boldsymbol{u}^*(\boldsymbol{\xi}, \boldsymbol{s}')\right] \tag{6}$$

is the two-point correlation tensor. If each member of the stochastic ensemble is square integrable, then $\boldsymbol{R}(\boldsymbol{s}, \boldsymbol{s}')$ is a nuclear operator and Hilbert-Schmidt theory guarantees that the eigenmodes satisfying equation (5) have a number of special properties. There is a countably infinite set of eigenmodes, $\{\boldsymbol{\phi}_j, \lambda_j\}$, that can be ranked according to their eigenvalue, $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$, which gives the energy captured by each mode. The eigenvectors are orthogonal, $\langle \boldsymbol{\phi}_j, \boldsymbol{\phi}_i \rangle = \delta_{ji}$, and complete. Accordingly, any member of the stochastic ensemble can be expanded as

$$\boldsymbol{u}(\boldsymbol{\xi}; \boldsymbol{s}) = \sum_{j=1}^{\infty} a_i(\boldsymbol{\xi}) \boldsymbol{\phi}_i(\boldsymbol{s}) \tag{7}$$

with $a_i(\boldsymbol{\xi}) = \langle \boldsymbol{u}(\boldsymbol{\xi}, \boldsymbol{s}), \boldsymbol{\phi}_i(\boldsymbol{s})\rangle$. This expansion is optimal in the sense that an approximation obtained by truncating the series at order k captures more energy than any other possible expansion using k basis functions. The eigenmodes also provide a diagonal representation of the correlation tensor,

$$\boldsymbol{R}(\boldsymbol{s}, \boldsymbol{s}') = \sum_{j=1}^{\infty} \lambda_j \boldsymbol{\phi}_j(\boldsymbol{s}) \boldsymbol{\phi}_j^*(\boldsymbol{s}'), \tag{8}$$

so they are its principle components. If the independent variable $\boldsymbol{s}$ includes dimensions that are homogeneous or stationary, the ensemble members will not be square-integrable, but a well-behaved eigenvalue problem can be recovered by first expanding these dimensions into Fourier modes and solving the Fourier-transformed version of equation (5).

POD has been used extensively within the fluid mechanics community to extract coherent structures from data.[13,14] As an example, Figure (4) shows several structures extracted from large-eddy simulation data for a Mach 0.9 jet.[15] The jet is round and stationary, so each POD modes corresponds to a particular azimuthal Fourier mode and temporal frequency, and the modes are functions of the streamwise and radial coordinates x and r, respectively. The axisymmetric components of the pressure at four frequencies are shown in the figure.[16] Each mode takes the form of a wavepacket, and these structures have been shown to contribute significantly to the production of jet noise.[17]

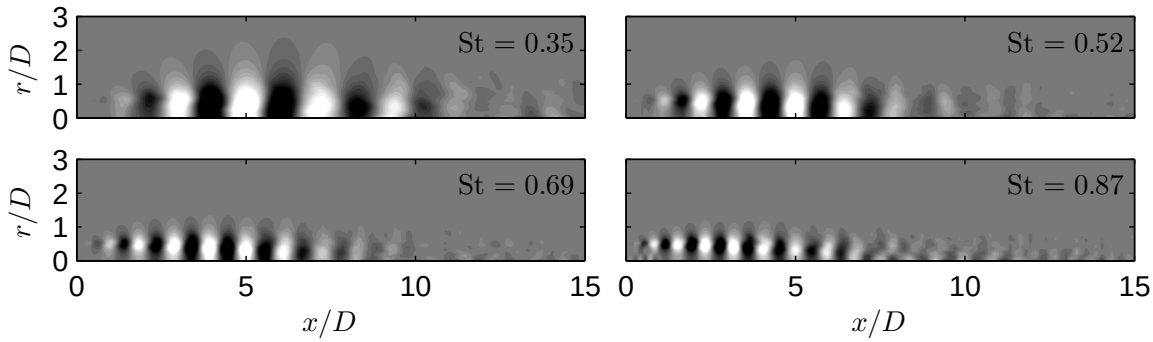

Figure 4: Pressure component of axisymmetric spectral POD mode for a Mach 0.9 turbulent jet at four frequencies. The frequencies are reported in terms of the Strouhal number $St = fD/U$, where f is the frequency, U is the jet velocity, and D is the jet nozzle diameter.

IV. Handling Independent Variables

When pursuing input parameter dimension reduction of the stochastic variables $\boldsymbol{\xi}$ in the presence of independent variables $\boldsymbol{s}$, there are a number of approaches one could adopt. In this section, we detail two approaches which make assumptions on the functional form of f, synthesize a POD-inspired approach, then show that the synthesized approach automatically recovers the previous two as special cases when their assumptions are satisfied. The two-point gradient correlation tensor will play a role in every approach which follows; it is defined here for convenience:

$$\boldsymbol{C}(\boldsymbol{s}, \boldsymbol{s}') \equiv \int \nabla_{\boldsymbol{\xi}} \boldsymbol{f}(\boldsymbol{\xi}, \boldsymbol{s}) \nabla_{\boldsymbol{\xi}}^T \boldsymbol{f}(\boldsymbol{\xi}, \boldsymbol{s}') d\boldsymbol{\xi}. \tag{9}$$

A. Integral approach

The basic idea of this approach is to integrate out the independent variables, and study the integrated quantity using the standard active subspace. We call this the integrated Active Subspace (iAS). This is a sensible approach when the integrated quantity is of practical interest, or when the gradient is well-correlated throughout the independent variables. Define

$$\boldsymbol{F}(\boldsymbol{\xi}) \equiv \int \boldsymbol{f}(\boldsymbol{\xi}, \boldsymbol{x}) d\boldsymbol{s}. \tag{10}$$

The gradient of $\boldsymbol{F}$ with respect to the stochastic parameters is

$$\nabla_{\boldsymbol{\xi}} \boldsymbol{F}(\boldsymbol{\xi}) = \int \nabla_{\boldsymbol{\xi}} \boldsymbol{f}(\boldsymbol{\xi}, \boldsymbol{s}) d\boldsymbol{s}. \tag{11}$$

Then the active subspace for $\boldsymbol{F}$ is defined by the eigenvalue problem

$$\boldsymbol{C}_{\mathrm{INT}}\boldsymbol{\phi}_{\mathrm{INT}} = \lambda_{\mathrm{INT}}\boldsymbol{\phi}_{\mathrm{INT}}, \tag{12}$$

where

$$\begin{aligned}\boldsymbol{C}_{\mathrm{INT}} &\equiv \int \nabla_{\boldsymbol{\xi}}\boldsymbol{F}(\boldsymbol{\xi})\nabla_{\boldsymbol{\xi}}^T\boldsymbol{F}(\boldsymbol{\xi})d\boldsymbol{\xi}, \\ &= \int\int\int \nabla_{\boldsymbol{\xi}}\boldsymbol{f}(\boldsymbol{\xi},\boldsymbol{s})\nabla_{\boldsymbol{\xi}}^T\boldsymbol{f}(\boldsymbol{\xi},\boldsymbol{s}')d\boldsymbol{s}d\boldsymbol{s}'d\boldsymbol{\xi}, \\ &= \int\int \boldsymbol{C}(\boldsymbol{s},\boldsymbol{s}')d\boldsymbol{s}d\boldsymbol{s}'.\end{aligned} \tag{13}$$

B. Local approach

The basic idea of this approach is to compute a separate active subspace for each value of the independent variable $\boldsymbol{s}$. We call this the local active subspace (lAS). This is a sensible approach when the gradient is independent throughout independent variable space. The corresponding eigenvalue problem is then

$$\boldsymbol{C}_{\mathrm{LOC}}(\boldsymbol{s})\boldsymbol{\phi}_{\mathrm{LOC}}(\boldsymbol{s}) = \lambda_{\mathrm{LOC}}(\boldsymbol{s})\boldsymbol{\phi}_{\mathrm{LOC}}(\boldsymbol{s}), \tag{14}$$

with

$$\begin{aligned}\boldsymbol{C}_{\mathrm{LOC}}(\boldsymbol{s}) &\equiv \int \nabla_{\boldsymbol{\xi}}\boldsymbol{f}(\boldsymbol{\xi},\boldsymbol{s})\nabla_{\boldsymbol{\xi}}^T\boldsymbol{f}(\boldsymbol{\xi},\boldsymbol{s})d\boldsymbol{\xi}, \\ &= \boldsymbol{C}(\boldsymbol{s},\boldsymbol{s}).\end{aligned} \tag{15}$$

C. Synthesized approach

Usually, the active subspace considers input parameter spaces, while POD considers output feature spaces. While the two techniques are often applied and interpreted differently, they are mathematically equivalent. To see this, note that to define an active subspace without independent variables, we select a qoi $f(\boldsymbol{\xi})$ and a weight function $\rho(\boldsymbol{\xi})$. Recall Equation (1)

$$\boldsymbol{C} = \int_{\mathbb{R}^m} \nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi})\nabla_{\boldsymbol{\xi}}^T f(\boldsymbol{\xi})\rho(\boldsymbol{\xi})\, d\boldsymbol{\xi}. \tag{1}$$

The leading eigenvectors of $\boldsymbol{C}$ form a basis for the active subspace. In the POD framework, we consider a case where there are no independent variables $\boldsymbol{s}$. We select a vector of features $\boldsymbol{u}(\boldsymbol{\xi})$ and a weight function $\rho(\boldsymbol{\xi})$. The goal is then to find $\boldsymbol{\phi}$ which maximizes the quantity

$$\lambda = \mathbb{E}[|\langle \boldsymbol{u}(\boldsymbol{\xi}), \boldsymbol{\phi}\rangle|^2]. \tag{16}$$

The solution to Equation 16 is the leading eigenvector of the eigenvalue problem

$$\boldsymbol{R}\boldsymbol{\phi} = \lambda\boldsymbol{\phi}, \tag{17}$$

where

$$\boldsymbol{R} = \int_{\mathbb{R}^m} \boldsymbol{u}(\boldsymbol{\xi})\boldsymbol{u}(\boldsymbol{\xi})^T\rho(\boldsymbol{\xi})d\boldsymbol{\xi}, \tag{18}$$

is the correlation tensor. Subsequent orthogonal solutions may be found by considering additional eigenvectors of $\boldsymbol{R}$. Note that if we select $\boldsymbol{u}(\boldsymbol{\xi}) = \nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi})$, we see that $\boldsymbol{C} = \boldsymbol{R}$. That is, the subspace arising from POD is the active subspace.

While the subspaces arising from the two formulations are identical, the distinct perspective between the two frameworks is both meaningful and useful. A POD interpretation of the modes arising from (17) would lead to a low-rank representation of the gradient; that is

$$\nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}) \approx \sum_{i=1}^{k} \psi_i(\boldsymbol{\xi})\boldsymbol{\phi}_i, \tag{19}$$

with $k < m$. The active subspace does not consider such an expansion, but rather considers the subspace spanned by the first k modes $\boldsymbol{\phi}_i$. This subtle distinction is one of intent; the POD perspective seeks to represent $\nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi})$, while the active subspace uses the same information to make statements about $f(\boldsymbol{\xi})$.

The observation above suggests an immediate generalization to the active subspace. It is straightforward to account for independent variables in the POD framework. For example, consider a qoi $f(\boldsymbol{\xi}, \boldsymbol{s})$ which depends on stochastic parameters $\boldsymbol{\xi} \in \mathbb{R}^m$ and an additional set of independent variables $\boldsymbol{s} \in \mathbb{R}^p$. In POD, the objective is to maximize

$$\lambda = \mathbb{E}[|\langle \nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}, \boldsymbol{s}), \boldsymbol{\phi}(\boldsymbol{s})\rangle|^2], \tag{20}$$

where the inner product is defined as

$$\langle \boldsymbol{a}(\boldsymbol{s}), \boldsymbol{b}(\boldsymbol{s})\rangle = \int_{\mathbb{R}^p} \boldsymbol{b}^*(\boldsymbol{s})\boldsymbol{M}(\boldsymbol{s})\boldsymbol{a}(\boldsymbol{s})d\boldsymbol{s}, \tag{21}$$

where $\boldsymbol{M}(\boldsymbol{s})$ is a user-defined weight function. The function $\boldsymbol{\phi}(\boldsymbol{s})$ which maximizes Equation 20 is the solution to the following eigenvalue problem

$$\int_{\mathbb{R}^p} \boldsymbol{R}(\boldsymbol{s}, \boldsymbol{s}')\boldsymbol{M}(\boldsymbol{s}')\boldsymbol{\phi}_{\mathrm{POD}}(\boldsymbol{s}')d\boldsymbol{s}' = \lambda \boldsymbol{\phi}_{\mathrm{POD}}(\boldsymbol{s}), \tag{22}$$

where

$$\boldsymbol{R}(\boldsymbol{s}, \boldsymbol{s}') = \int_{\mathbb{R}^m} \nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}, \boldsymbol{s}) \nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}, \boldsymbol{s})^T \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} = \boldsymbol{C}(\boldsymbol{s}, \boldsymbol{s}') \tag{23}$$

is the $\boldsymbol{s}$ -dependent two-point correlation tensor of the gradient. The choice of $\boldsymbol{M}(\boldsymbol{s})$ may be selected to emphasize different components of $\nabla_{\boldsymbol{\xi}} f$ and/or different regions of $\boldsymbol{s} \in \mathbb{R}^p$. Note that if $\boldsymbol{M}(\boldsymbol{s}) = \mathrm{Diag}(\delta(\boldsymbol{s} - \boldsymbol{\tau}_0))$, then this generalized approach reduces to computing the usual active subspace at $\boldsymbol{\tau}_0$.

D. Limiting cases

As noted above, the synthesized POD method reduces to the integral and local approaches automatically, based on properties of f. We consider these two special cases:

1. *Perfectly correlated gradients*

Suppose the gradient is perfectly correlated in $\boldsymbol{s}$; this may be expressed mathematically as

$$\boldsymbol{C}(\boldsymbol{s}, \boldsymbol{s}') = \boldsymbol{C}_{\mathrm{const}}. \tag{24}$$

To relate to the integral method, we first note that from (13) we have $\boldsymbol{C}_{\mathrm{INT}} = S^2 \boldsymbol{C}_{\mathrm{const}}$, where S^2 is the hyper-volume of the $\boldsymbol{s}$ domain. Inserting (24) into (22) yields

$$\frac{1}{S^2}\boldsymbol{C}_{\mathrm{INT}} \int \boldsymbol{\phi}(\boldsymbol{s}')d\boldsymbol{s}' = \lambda \boldsymbol{\phi}(\boldsymbol{s}). \tag{25}$$

Since the matrix is independent of $\boldsymbol{s}$, so too are its eigenvectors, thus (25) becomes

$$\frac{1}{S}\boldsymbol{C}_{\mathrm{INT}}\boldsymbol{\phi} = \lambda \boldsymbol{\phi} \tag{26}$$

This is identical to the eigenproblem for the integrated approach (12), up to a scaling factor $1/S$ which has no impact on the eigenvectors, and uniformly scales all eigenvalues.

2. Uncorrelated gradients

Suppose the gradient is uncorrelated in $\boldsymbol{s}$; this may be expressed as

$$\boldsymbol{C}(\boldsymbol{s}, \boldsymbol{s}') = \boldsymbol{C}_{\mathrm{LOC}}(\boldsymbol{s})\delta(\boldsymbol{s} - \boldsymbol{s}'). \tag{27}$$

This form is identical to (15); thus inserting (27) into (22) yields

$$\int \boldsymbol{C}_{\mathrm{LOC}}(\boldsymbol{s})\delta(\boldsymbol{s} - \boldsymbol{s}')\boldsymbol{\phi}(\boldsymbol{s}')d\boldsymbol{s}' = \lambda\boldsymbol{\phi}(\boldsymbol{s}), \tag{28}$$

and evaluating the integral leads to (14).

These two limits demonstrate that the synthesized POD method recovers the two approaches above automatically, **based on the properties of $\boldsymbol{f}$**. Intuitively, the POD approach handles the boundary cases appropriately, and so one would expect the approach recovers an appropriate dimension reduction for cases intermediate to these limits.

E. Illustrative test case

Here we consider the generalized dimension reduction technique suggested above on a simple test case. Consider the function

$$f(\boldsymbol{\xi}, s) = (\xi_1 \cos(\pi s/2) + \xi_2 \sin(\pi s/2))^3, \tag{29}$$

with uniform $\rho(\boldsymbol{\xi}, s)$ over the unit hypercube $(\xi_1, \xi_2, s) \in [0, 1]^3$. Figure 5 depicts the resulting POD modes for this example, as well as the locally computed active subspace results. The modes are depicted as *traces*; plots of the individual vector components as they vary in the independent variable s. The results of Figure 5 demonstrate that while the modes from the POD and local AS formulations are similar, the local AS modes have jumps which make them inconvenient to interpret. The differences could be reduced by manually changing the sign of the lAS modes (as the sign is arbitrary), but it would not be clear that this is the 'proper' correction without a priori knowledge of the POD modes, as there is no reason to expect smoothness in the lAS framework. In contrast, the relative sign between different s values is an intrinsic property of the POD modes.

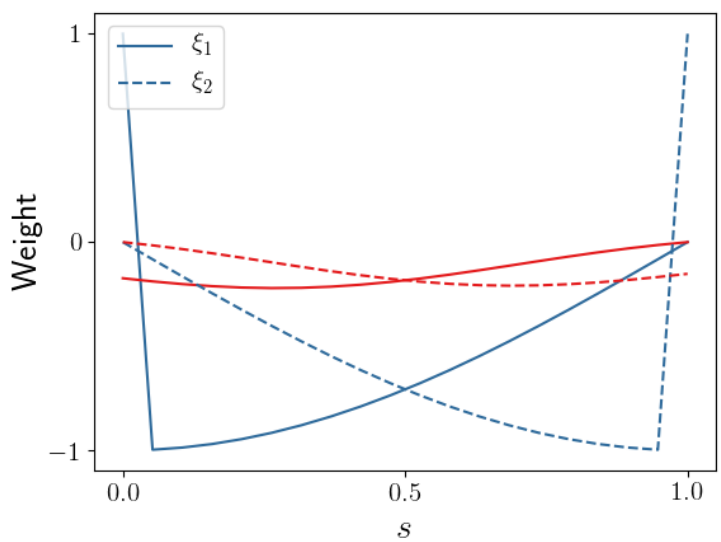

Figure 5: Leading mode for the lAS and POD modes. Each colored curve is the s-varying value for each entry of the depicted vector mode, with Blue depicting the lAS, and Red for POD. Note that the two techniques are in qualitative agreement. The discontinuities in the active subspace are inconvenient, but not intrinsically meaningful as eigenvectors are only defined up to scalar multiples.

When applying an (estimated) active subspace, it is common to study the *summary plots*,[18] defined by the pairs $(f_i, \boldsymbol{W}_a^T \boldsymbol{\xi}_i)$. In the case where $k \leq 2$, it is possible to depict these pairs graphically. We may define a generalization which accounts for additional variability in $\boldsymbol{s}$ by considering the pairs $(f_i, \boldsymbol{\Phi}(\boldsymbol{s}_i)^T \boldsymbol{\xi}_i)$, where $\boldsymbol{\Phi}(\boldsymbol{s}_i)$ are the modes resulting from the POD approach. Figure 6 demonstrates that the POD modes successfully account for all the variability in our qoi, leading to a complete dimension reduction.

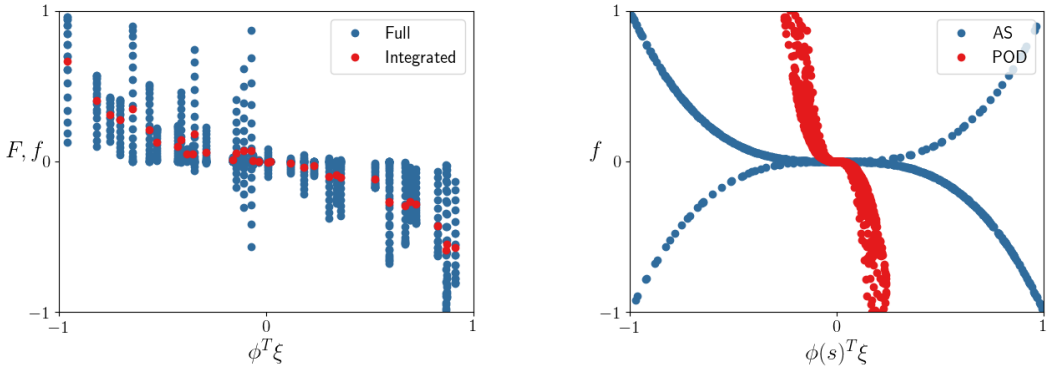

Figure 6: Summary plots using the leading iAS mode (Left) and leading lAS and POD modes (Right). On the Left, note that the averaged (in s) quantity (Red dots) collapses under the iAS mode, but the full variability is not captured. On the Right, the sign ambiguity of the lAS prevents a full collapse of the data (Blue), while the POD modes resolve this issue (Red).

V. Airfoil Design

In this section, we consider the shape design of a subsonic ($M = 0.5, \alpha = 1.25°$) airfoil. Here, we introduce the physical problem, compare the modes arising from different dimension reduction techniques, the illustrate a use of the POD modes to perform spatially-localized sensitivity analysis. Note that our objective in this work is *not* to accomplish a single design task, e.g. retrieve a single airfoil geometry which optimizes the lift, but rather to derive *insight*, e.g. determine how a particular shape deformation affects airfoils on average.

A. Problem description

We consider the physical problem of the geometric design of an airfoil. Our quantity of interest is the vertical contribution of the entire pressure distribution $C_p(s)$. As with Subsection II A, we consider a baseline NACA0012 profile and apply $d = 38$ Hicks-Henne bump functions in order to parameterize the geometry.[11] We use the open-source CFD solver SU2[11] to solve the Euler equations, draw a low-discrepancy set of design points via a Halton sequence,[19] and utilize finite-diferencing to compute the gradient of our qoi. Figure 7 gives the baseline flow about the airfoil, and a schematic view of the airfoil coordinate system s used below.

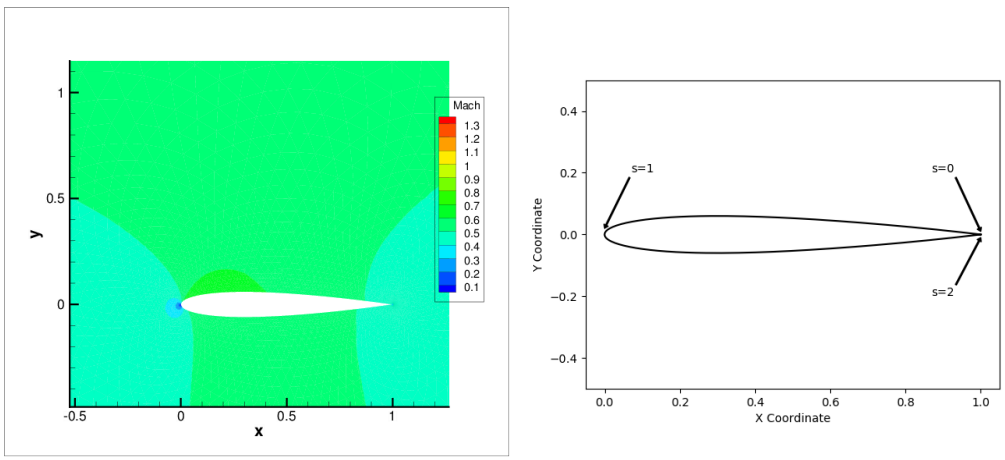

Figure 7: Baseline flow for NACA0012 airfoil at $M = 0.5$ and $\alpha = 1.25°$ (Left), and schematic view of airfoil coordinate system (Right).

B. Dimension reduction modes

In this section, we investigate the modes arising from lAS, iAS, and POD approaches to dimension reduction. Figure 8 gives the individual eigenvalues and their cumulative sum; this spectrum suggests that the leading POD mode captures 94% of the average variability of the gradient, as measured by the $l2$ norm. Thus, we compare only the leading POD mode against the leading modes of the lAS and iAS approaches.

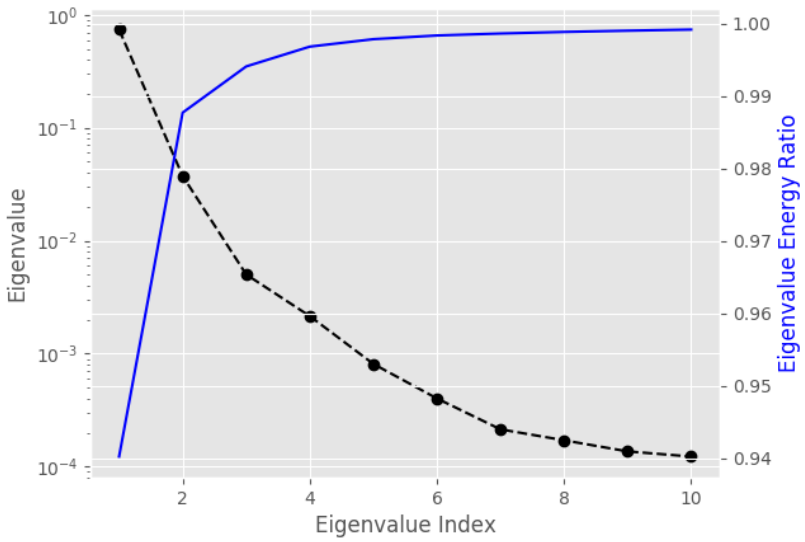

Figure 8: Eigenvalues arising from the POD computation. The leading POD mode accounts for the vast majority of the eigenvalue energy.

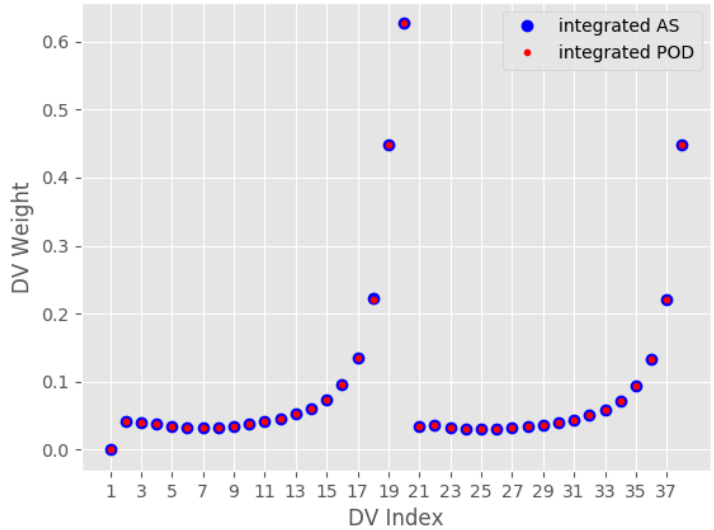

Figure 9: Comparison of leading integrated active subspace and integrated POD modes. Note that the modes are in close agreement. While the iAS and POD modes are fundamentally different objects, in this case their leading modes share some information. This is not always the case, as we will see below.

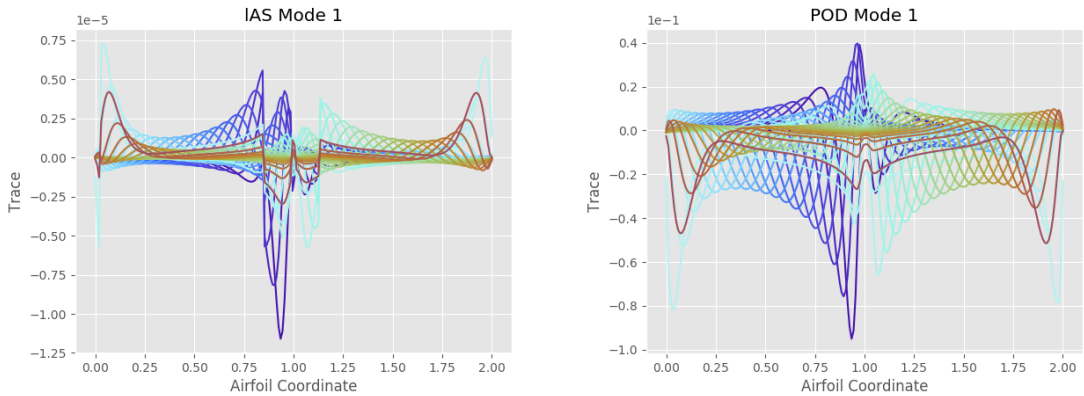

Figure 10: Comparison of leading local active subspace mode (Left) and leading POD mode (Right). Each color denotes a different design variable. Note that the active subspace directions are eigenvectors, and thus are defined only up to a scalar multiple. To aid in comparing the lAS with the POD modes, we have scaled the former with the local eigenvalue $\lambda(s)$ arising from the lAS computation. Note that the lAS modes exhibit jumps, a consequence of the local eigenvalue problem for each value of s.

C. Sensitivity

In this section, we suggest a practical application of the POD modes to a form of s -localized sensitivity analysis. We consider a deformation to the airfoil corresponding to the leading iAS mode; as suggested above in Subsection A, this is related to a change in camber of the airfoil. Note that Figure 8 implies the leading POD mode captures $\sim 94\%$ of the average gradient energy, as measured by the $l2$ norm. This suggests that taking the inner product of an arbitrary vector $\boldsymbol{v}$ defining a shape deformation with this leading POD mode will describe the variability of the qoi, localized in s and averaged in ξ.

Figure 11 depicts this sensitivity analysis, whose results suggest that the leading iAS mode causes the greatest change to the pressure distribution at the fore and aft of the airfoil. We corroborate this finding in Figure 12, which applies the deformation to different baseline geometries.

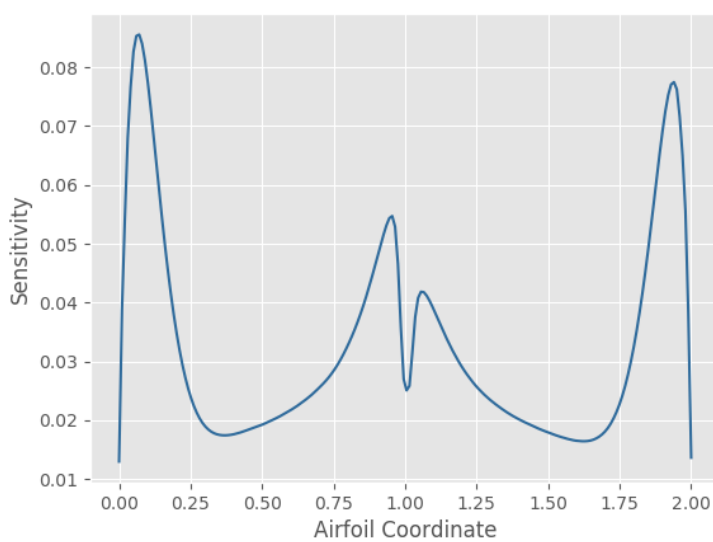

Figure 11: Localized sensitivity of the chosen shape deformation. Recall that $s = 1$ corresponds to the fore of the airfoil, while $s = 0, 2$ correspond to the aft. Note that these sensitivity results suggest the pressure distribution changes most at the fore and aft, under the studied shape deformation.

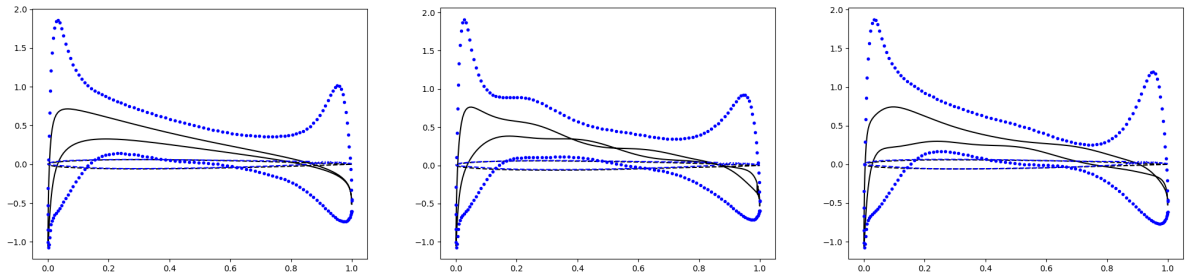

Figure 12: Comparison of baseline (Black) and perturbed (Blue) pressure distributions. Each of the baseline geometries are taken from the random ensemble of designs, and perturbed along the leading iAS mode. As expected, the same perturbation affects different points in design space in the same fashion predicted by the sensitivity analysis; that is, the pressure distribution is most perturbed at the fore and aft of the airfoil.

VI. Structural Design

In this section, we consider the structural design of a thin-walled box-beam, and consider the field of shear stress as our qoi. Here, we introduce the physical problem, compare the modes arising from different dimension reduction techniques, then illustrate a use of the POD modes to affect spatially-targeted changes to the qoi.

A. Problem description

Figure 13 gives a schematic view of the geometry for the thin-walled box-beam, as well as an example shear distribution for a beam of uniform thickness. The beam is assumed to have a square cross-section of perimeter 4, with a thickness that varies in the *beam coordinate* $s \in [0, 4]$. An external shear loading is assumed to be applied through the shear center, such that no twist is generated. We use the method of *shear flow* to determine the shear distribution in the box-beam, making standard assumptions about the out-of-plane shear components.[20]

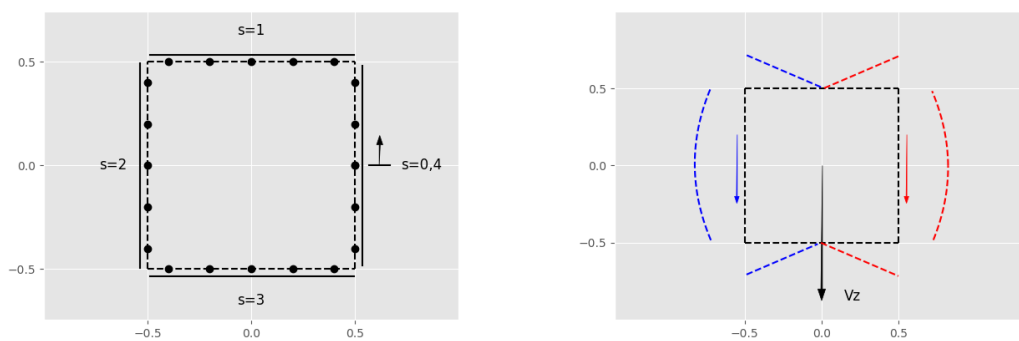

Figure 13: Schematic depiction of box-beam geometry (Left) and shear distribution for a uniform thickness (Right). The coordinate s is used to depict locations on the box-beam, starting with the center of the right edge, and moving counter-clockwise. The dotted black line is the beam profile, and the control points for the thickness distribution are depicted as black dots. On the Left, the solid black line depicts an illustrative uniform thickness distribution. On the Right, the Vz arrow depicts the shear loading. The colors blue and red depict shear in the positive and negative s direction, respectively. For elucidation, the blue and red arrows depict the shear direction, while the corresponding blue and red curves depict the magnitude of the shear stress; the magnitude is greater where the curve is further from the beam profile

Our qoi is the entire distribution of shear $\tau(s) = q(s)/t(s)$. Note that this physical problem exhibits both local and global behavior. Since the shear flow $q(s) = q'(s) + q_0$ is dependent on a local $q'(s)$ and global q_0

component, there is obviously correlation between different beam coordinates. Furthermore, since the shear itself $\tau(s) = q(s)/t(s)$ is dependent on the local thickness, the shear exhibits some local behavior as well. Thus, we have a combination of both local and global behaviors, and expect that the ultimate behavior of the shear distribution will not lie at either of the limiting cases described above.

B. Dimension reduction modes

In this section, we investigate the modes arising from the lAS, iAS, and POD approaches to dimension reduction on the structural problem. Figure 14 reports results from the POD approach, and gives the individual eigenvalues and their cumulative percentage. This spectrum suggests the leading POD mode accounts for a large portion of the gradient variability ($\sim 40\%$), but is not sufficient to describe most of the variation.

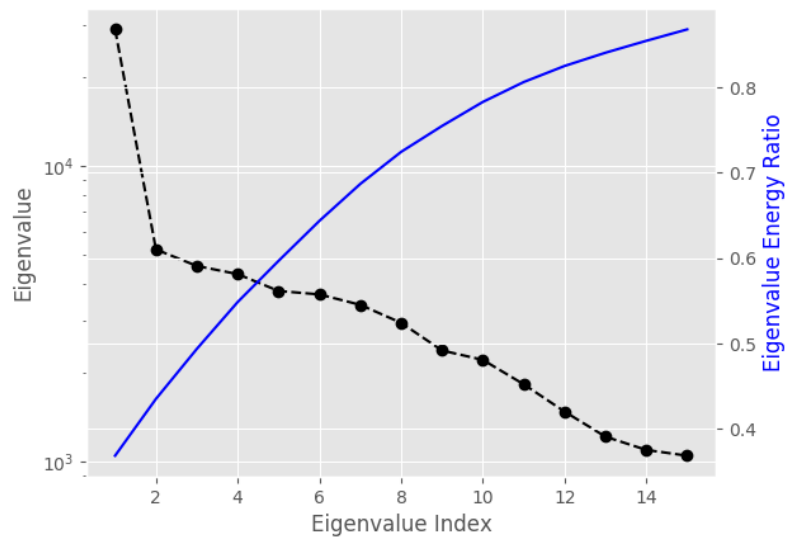

Figure 14: Eigenvalues arising from the POD computation. The leading POD mode accounts for a modest portion of the eigenvalue energy.

C. Manipulating the shear

In this section, we suggest a practical application of the POD modes to locally modify the shear distribution. One may regard the POD modes as localizing in $\boldsymbol{s}$ the sensitivity of the qoi to the design variables, accounting for structure in the independent variables. We first consider taking the leading POD mode at a single point $s = 0$, applying this deformation to the baseline geometry of uniform thickness, and observing the results (Fig. 17). The results lead to a reduction in shear at $s = 0$, but undesirable behavior outside the studied point. We then consider averaging the POD modes over the sidewalls ($s \in [0, 0.5] \cap [1.5, 2.5] \cap [3.5, 4.0]$), and apply this composite deformation (Fig. 18). This leads to a reduction in shear over both sidewalls. The modes themselves correspond to a thickening of the beam – an intuitive conclusion, but here arising not from intuition, but from data.

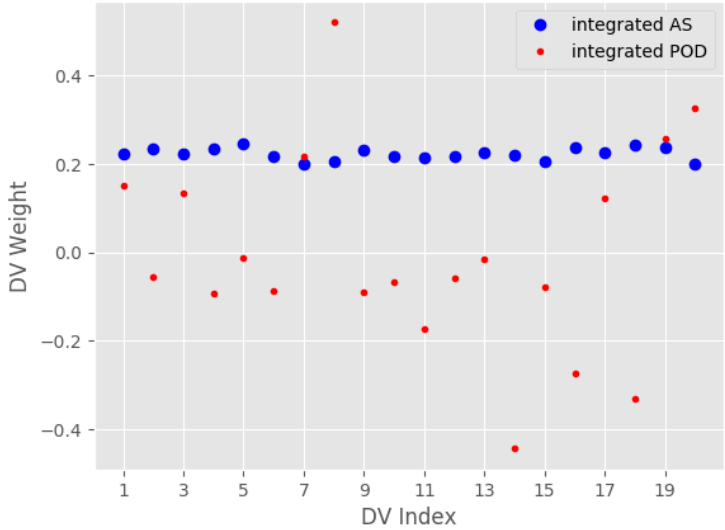

Figure 15: Comparison of leading integrated active subspace and integrated POD modes. Note that the modes are in disagreement; unlike the airfoil problem, the modes arising from different computational procedures do not share information. Note that the leading iAS mode corresponds to a uniform thickening of the structure, which is what we would expect for affecting the average shear in the beam.

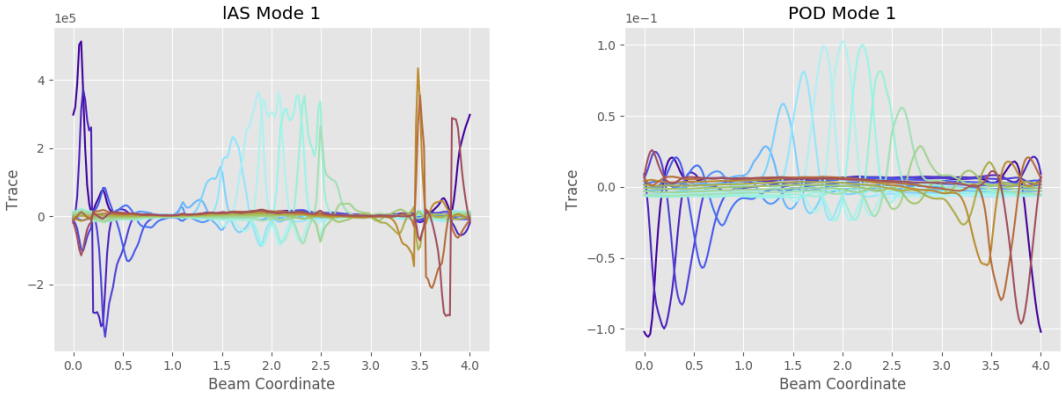

Figure 16: Comparison of leading local active subspace mode (Left) and leading POD mode (Right). As before, the lAS mode has been scaled by its spatially-varying eigenvalue $\lambda(s)$ to aid in comparing these different objects.

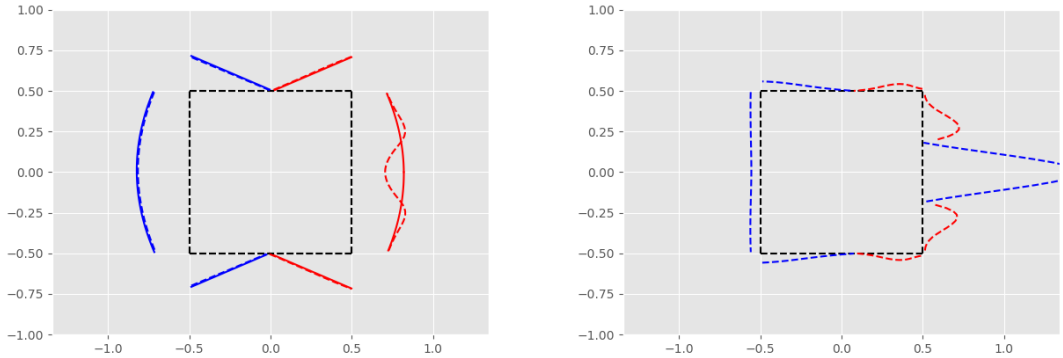

Figure 17: Shear distributions (Left) and thickness perturbation (Right) using the POD mode at $s = 0$. On the Left, the shear distribution for the baseline profile of uniform thickness is shown in solid curves, while that of the perturbed geometry is shown in dashes. Note that the shear at the studied location is reduced, but adjacent regions are affected differently. On the Right, the geometry perturbation is visualized; Blue denotes adding material, while Red denotes removal. The geometry perturbation corresponds to thickening the walls near the point $s = 0$, but thinning on the adjacent region.

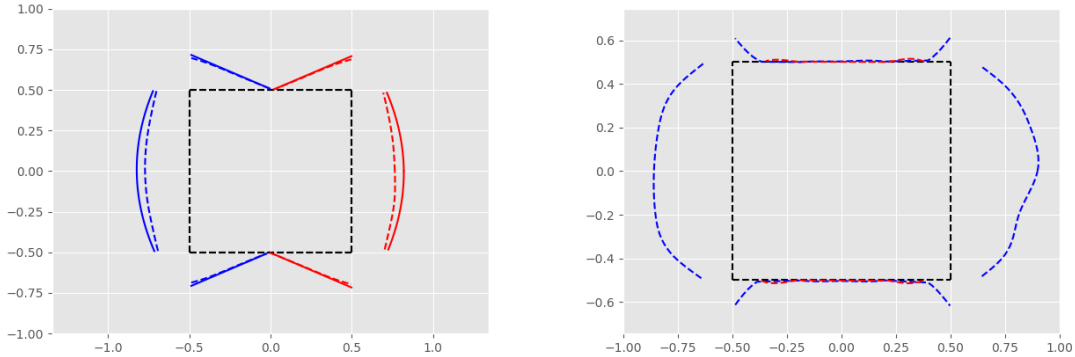

Figure 18: Shear distributions (Left) and thickness perturbation (Right) using the POD modes averaged over the sidewalls. Note that the shear across both sidewalls is reduced, while the rest of the shear distribution is largely unaffected. This is accomplished by thickening both walls in a smooth fashion. While certainly an intuitive way to reduce shear, this particular deformation was derived not from intuition, but from data. Note that symmetry is not explicitly enforced in this problem – the lack of perfect symmetry in the deformation is due to a lack of perfect convergence in the estimation procedure.

VII. Conclusion

In this work, we sought automated, data-driven techniques for deriving exploitable understanding of a physical problem – so-called design insight. We reviewed previous work with scalar quantities of interest, using dimension reduction techniques to identify usable structure within a set of stochastic variables $\boldsymbol{\xi}$. We presented a mathematical framework to handle quantities which vary in a set of additional independent variables $\boldsymbol{s}$, and compared our POD-inspired method with other dimension reduction techniques. Finally, we introduced a pair of design problems of interest in aerospace applications, and illustrated the usage of dimension reduction to probing these settings, with an eye towards developing understanding. The result for these well-studied problems was the recovery of sensible insights; however, applying these same techniques to novel problems suggests the possibility of discovering new knowledge.

The first author was supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-114747. The second author was was funded by NASA grant no. NNX15AU93A and by the NNSA Predictive Science Academic Alliance Program II, grant no. DE-NA0002373.

References

[1]Pahl, G., Beitz, W., Feldhusen, J., and Grote, K., *Engineering Design: A Systematic Approach*, Springer, 3rd ed., 2007.

[2]Shvartsman, S., personal communication.

[3]Brunton, S. L., Proctor, J. L., and Kutz, J. N., "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proceedings of the National Academy of Science*, Vol. 113, No. 15, 2016, pp. 3932–3937.

[4]Schmidt, M. and Lipson, H., "Distilling Free-Form Natural Laws from Experimental Data," *Science*, Vol. 324, No. 5923, 2009, pp. 81–85.

[5]Lumley, J. L., *Stochastic tools in turbulence*, Courier Corporation, 2007.

[6]del Rosario, Z., Constantine, P. G., and Iaccarino, G., "Developing Design Insight Through Active Subspaces," *19th AIAA Non-Deterministic Approaches Conference*, 2017, p. 1090.

[7]Constantine, P. G., *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*, SIAM Philadelphia, 2015.

[8]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.

[9]Jefferson, J., Gilbert, J., Constantine, P. G., and Maxwell, R., "Active subspaces for sensitivity analysis and dimension reduction of an integrated hydrologic model," *Computers & Geosciences*, Vol. 83, 2015, pp. 127–138.

[10]Othmer, C., Lukaczyk, T. W., Constantine, P. G., and Alonso, J. J., "On Active Subspaces in Car Aerodynamics," *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2016, p. 4294.

[11]Palacios, F., Colonno, M., Aranake, A., Campos, A., Copeland, S., Economon, T., Lonkar, A., Lukaczyk, T., Taylor, T., and Alonso, J., "Stanford University Unstructured (SU2): An open-source integrated computational environment for multi-physics simulation and design," *AIAA Aerospace Sciences Meeting and Exhibit*, 2013.

[12]Anderson, J. J., *Fundamentals of Aerodynamics*, McGraw-Hill, 3rd ed., 2001.

[13]Berkooz, G., Holmes, P., and Lumley, J. L., "The proper orthogonal decomposition in the analysis of turbulent flows," *Annu. Rev. Fluid Mech.*, Vol. 25, No. 1, 1993, pp. 539–575.

[14]Towne, A., Schmidt, O. T., and Colonius, T., "Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis," *arXiv:1708.04393*, 2017.

[15]Bres, G., Jaunet, V., Le Rallic, M., Jordan, P., Towne, A., Schmidt, O., Colonius, T., Cavalieri, A., and Lele, S., "Large eddy simulation for jet noise: azimuthal decomposition and intermittency of the radiated sound," *22nd AIAA/CEAS Aeroacoustics Conference*, 2016.

[16]Towne, A., *Advancements in jet turbulence and noise modeling: accurate one-way solutions and empirical evaluation of the nonlinear forcing of wavepackets*, Ph.D. thesis, California Institute of Technology, 2016.

[17]Jordan, P. and Colonius, T., "Wave Packets and Turbulent Jet Noise," *Annu. Rev. Fluid Mech.*, Vol. 45, 2013, pp. 173–195.

[18]Cook, R. D., *Regression graphics: Ideas for studying regressions through graphics*, Vol. 482, John Wiley & Sons, 2009.

[19]Owen, A., "A randomized Halton algorithm in R," *ArXiv e-prints*, jun 2017.

[20]Sun, C.-T., *Mechanics of aircraft structures*, J. Wiley, 2006.