Data-driven Model Reduction via Operator Inference for Coupled Aeroelastic Flutter

Benjamin G. Zastrow*, Anirban Chaudhuri†, Karen E. Willcox‡
University of Texas at Austin, Austin, TX, 78712

Anthony Ashley§, Mike Henson¶
Lockheed Martin, Fort Worth, TX, 76101

This paper presents a data-driven, physics-informed model reduction technique applied to two large-scale applications: parametric nonlinear aerodynamics and coupled aeroelastic flutter. We first develop a reduced-order model (ROM) parameterized by Mach number for the compressible Navier-Stokes equations solved via a computational fluid dynamics code. We use the non-intrusive operator inference scientific machine learning method as our model reduction technique. We then extend this technique to the coupled aeroelastic case by incorporating a modal decomposition model obtained from a finite element analysis code. We use NASA's FUN3D software to run the high-fidelity simulations. We demonstrate the parametric ROM for nonlinear subsonic fluid flow over the VAT aircraft and the coupled ROM for flutter of the AGARD wing. We show that compared to the high-fidelity FUN3D simulations, parametric operator inference speeds up computation time by approximately three orders of magnitude, and coupled operator inference speeds up flutter computation time by approximately two orders of magnitude. **

I. Introduction

Finite element analysis (FEA) and computational fluid dynamics (CFD) are widely used in aircraft design but their usefulness is limited by their high computational cost, which limits the degree to which the design space can be explored. CFD models can routinely have tens of millions of degrees of freedom, requiring significant simulation time on high performance computing clusters to make useful predictions. This precludes them from use in many-query problems, such as uncertainty quantification, flutter analysis, design optimization, and integration into digital twins for online control. Model reduction techniques lead to cheaper, physics-informed surrogates that can be used in many-query problems. This paper presents a data-driven model reduction technique applied to two large-scale applications in parametric nonlinear aerodynamics and coupled aeroelastic flutter analysis.

Model reduction aims to represent high-dimensional dynamics with a lower-dimensional system while maintaining a sufficient level of predictive accuracy. The reduced-order models (ROMs) can be run at orders of magnitude lower computational cost compared to the full-order models (FOMs). Projection-based approaches like the proper orthogonal decomposition (POD) [1–4] use the trajectory data from a FOM to derive a reduced basis, then the full-order operators of the governing equations are projected onto the reduced basis and the dynamics are integrated forward in time [5, 6]. Another technique, dynamic mode decomposition (DMD) [7, 8], uses the full-order trajectory data to construct a linear reduced operator to approximate the system, thus enabling analysis of eigenmodes and eigenvalues that often can reasonably characterize even nonlinear full-order dynamics. The governing equations of many physical systems of interest contain parametric dependencies, and thus they require a decision about whether to define the ROM using a local or global basis, with implications for both the accuracy and the computational performance of the ROM. In [5], the effectiveness of rational interpolation methods, balanced truncation, POD, and the reduced basis method are compared from a parametric ROM perspective.

*Graduate Student, Department of Aerospace Engineering and Engineering Mechanics, AIAA Student Member.
†Research Associate, Oden Institute for Computational Engineering and Sciences, AIAA Member.
‡Director, Oden Institute for Computational Engineering and Sciences, AIAA Fellow.
§Aeronautical Engineer Senior, CFD, AIAA Member
¶Aeronautical Engineer, Senior Staff, Transformational Solutions/AeroIT, AIAA Associate Fellow

**This version corrects a few minor errors in the manuscript published in the AIAA SciTech proceedings, notably in Figures 12 and 13, and it corrects Section IV.B to state that the AGARD fluid flow model is compressible.

Aeroelastic flutter was originally analyzed using linear methods that focused on the generation of V-g (velocity vs. damping) and V-f (velocity vs. frequency) plots to characterize the boundaries of the safe operational range for an aircraft [9, 10]. A number of model reduction techniques have been proposed as a way to bring higher fidelity aerodynamic and structural information into aeroelastic computations. In [11], a vortex lattice model of flow over a 2D airfoil is used to demonstrate that unsteady flow solutions can be characterized with a small number of dominant eigenmodes of a modal decomposition. A review of several modal reduction methods for unsteady aerodynamics, including eigenmodes, POD modes, and balanced modes, is given in [12], along with a discussion of implementing these methods for aeroelastic models. A parametric reduction method for flutter that interpolates between ROMs at different free stream Mach numbers using the tangent space of a Grassman manifold is applied to both the F-16 and F-18/A full-aircraft configurations in [13]. In another approach, Walsh functions are used to simultaneously excite multiple impulse responses in a CFD model [14]. The eigensystem realization algorithm (ERA) [15] is then used to convert these impulses into an unsteady aerodynamic ROM that can be coupled to a structural model to create a coupled aeroelastic ROM [16]. The Walsh function technique from [14] has been extended to a POD methodology in [17] that builds the basis via an incremental approach to avoid handling the full snapshot matrix all at once. In this work, we build non-intrusive ROMs for coupled multi-physics problems such as aeroelastic flutter by extending the operator inference method [18].

Operator inference is a scientific machine learning approach that replaces the intrusive Galerkin projection step with a non-intrusive linear least squares problem, thus combining the benefits of data-driven learning and physics-informed modeling [18]. An additional "lifting" step is formally suggested in [19] to expose the desired polynomial ordinary differential equation structure via introduction of auxiliary state variables. The parametric modeling capability of operator inference is detailed for an affine parametric structure in [20]. We build on the data-driven operator inference method [18] to learn non-intrusive ROMs for two real-world applications: a parametric fluid flow model of the Validation of Aeroelastic Tailoring (VAT) aircraft [21], and a coupled aerostructural flutter model of the Advisory Group for Aerospace Research and Development (AGARD) wing [22]. We generate high-dimensional simulation data by using NASA's FUN3D software [23] for the aerodynamic analysis of the VAT aircraft and the AGARD wing. For the flutter analysis, we use FUN3D's aeroelasticity capability to couple the structural and fluid dynamics and generate snapshot training data for the AGARD wing. The main contributions of this paper are:

(i) *developing a non-intrusive parametric ROM for large-scale aerodynamics analysis:* We develop a non-affine parametric data-driven operator inference ROM for aerodynamic analysis under varying Mach numbers. We build the ROMs via reduced-state solution interpolation with global basis vectors.

(ii) *developing a non-intrusive coupled ROM for flutter analysis:* We develop coupled data-driven operator inference ROMs for flutter analysis. We exploit existing knowledge of the structure of the coupled system by identifying basis functions for the aerodynamic states separately from the structural basis functions obtained via modal decomposition. We then combine the structural and fluid reduced states to learn a single coupled ROM.

The remainder of this paper is organized as follows. Section II describes the operator inference method as initially presented in [18], which will be used to construct the parametric reduced-order models. Section III applies the operator inference method to the coupled (aerostructural) setting. Section IV presents the application of the operator inference method to the VAT aircraft subsonic fluid flow problem and to the AGARD wing aerostructural flutter problem. Section V provides concluding remarks.

II. Non-intrusive learning of reduced-order models using operator inference

Operator inference is a scientific machine learning method that uses knowledge of the structure of the full-order governing equations to specify a matching structure for the reduced-order model. In Section II.A, we present a summary of the standard non-parametric operator inference framework. Section II.B presents an approach for parametric operator inference ROMs via reduced-state solution interpolation.

A. Non-parametric operator inference

This work targets systems governed by nonlinear partial differential equations. We consider a semi-discrete form of the governing equations after spatial discretization that results in a polynomial ordinary differential equation structure as

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{x}(t) = \mathbf{c} + \mathbf{A}\mathbf{x}(t) + \mathbf{H}\left(\mathbf{x}(t) \otimes \mathbf{x}(t)\right) + \mathbf{B}\mathbf{u}(t) \tag{1}$$

where $\mathbf{x}(t) \in \mathbb{R}^n$ is the semi-discrete state vector at time t, $\mathbf{u}(t) \in \mathbb{R}^m$ is the input vector at time t with m inputs, $\mathbf{c} \in \mathbb{R}^n$ are the constant terms, $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the discretized linear operator, $\mathbf{H} \in \mathbb{R}^{n \times n^2}$ is the discretized quadratic operator, and $\mathbf{B} \in \mathbb{R}^{n \times m}$ is the input operator. The dimension n of the state $\mathbf{x}$ is given by $n = dn_x$, with d as the number of state variables and a spatial discretization with n_x cells. For dynamical systems that are not initially in this polynomial form, we can expose such structure via lifting variable transformations as shown in [19].

We build a projection-based ROM that can preserve the polynomial structure of Eq. (1). To build a ROM, we need to derive a reduced basis which defines a low-dimensional subspace in which the dynamics are approximated. We obtain this basis from the full state trajectory via the proper orthogonal decomposition (POD). The POD basis vectors are obtained by taking the singular value decomposition (SVD) of the snapshot matrix $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_k] \in \mathbb{R}^{n \times k}$, where each column $\mathbf{x}_i = \mathbf{x}(t_i),\ i = 1, \ldots, k$ is the full-order state solution at a given time t_i, and k is the number of snapshots in $\mathbf{X}$. Typically $k \ll n$ for large-scale applications. The thin SVD of the snapshot matrix is $\mathbf{X} = \mathbf{V}\mathbf{\Sigma}\mathbf{W}^\top$, where $\mathbf{V} \in \mathbb{R}^{n \times k}$ and $\mathbf{W} \in \mathbb{R}^{k \times k}$ are orthogonal matrices, and $\mathbf{\Sigma} \in \mathbb{R}^{k \times k}$ is a square diagonal matrix consisting of the singular values of the snapshot matrix. Our reduced basis for projection $\mathbf{V}_r \in \mathbb{R}^{n \times r}$ consists of the first r columns of $\mathbf{V}$, where $r \ll n$ (and $r < k$). The approximation of the full state in terms of reduced state $\widehat{\mathbf{x}}(t) \in \mathbb{R}^r$ is given by $\mathbf{x}(t) \approx \mathbf{V}_r \widehat{\mathbf{x}}(t)$.

In an intrusive projection-based model reduction approach, Galerkin projection is used to obtain the reduced matrix operators and requires access to source code of the simulation software. However, operator inference is a non-intrusive projection method that does not require access to the full-order operators of Eqn. (1), instead learning the reduced operators via the solution of a linear least squares problem [18]. We wish to learn reduced-order dynamics that match the form of the full-order dynamics in Eq. (1), leading to the reduced system

$$\frac{\mathrm{d}}{\mathrm{d}t}\widehat{\mathbf{x}}(t) = \widehat{\mathbf{c}} + \widehat{\mathbf{A}}\widehat{\mathbf{x}}(t) + \widehat{\mathbf{H}}\big(\widehat{\mathbf{x}}(t) \otimes \widehat{\mathbf{x}}(t)\big) + \widehat{\mathbf{B}}\mathbf{u}(t), \tag{2}$$

where $\widehat{\mathbf{c}} \in \mathbb{R}^r, \widehat{\mathbf{A}} \in \mathbb{R}^{r \times r}, \widehat{\mathbf{H}} \in \mathbb{R}^{r \times r^2}$, and $\widehat{\mathbf{B}} \in \mathbb{R}^{r \times m}$ are the reduced operators.

The reduced operators are obtained by solving the regularized linear least squares problem

$$\min_{\widehat{\mathbf{c}}, \widehat{\mathbf{A}}, \widehat{\mathbf{H}}, \widehat{\mathbf{B}}} \left\{ \sum_{i=1}^{k} \left\| \widehat{\mathbf{c}} + \widehat{\mathbf{A}}\widehat{\mathbf{x}}_i + \widehat{\mathbf{H}}(\widehat{\mathbf{x}}_i \otimes \widehat{\mathbf{x}}_i) + \widehat{\mathbf{B}}\mathbf{u}_i - \dot{\widehat{\mathbf{x}}}_i \right\|_2^2 + \lambda_1 \left(\|\widehat{\mathbf{c}}\|_2^2 + \left\|\widehat{\mathbf{A}}\right\|_F^2 + \left\|\widehat{\mathbf{B}}\right\|_F^2 \right) + \lambda_2 \left\|\widehat{\mathbf{H}}\right\|_F^2 \right\}, \tag{3}$$

where $\lambda_1 > 0$ and $\lambda_2 > 0$ are the regularization hyperparameters, $\widehat{\mathbf{x}}_i$ is the reduced state vector of ith snapshot, and $\dot{\widehat{\mathbf{x}}}_i$ is the reduced state time derivative estimated via finite difference approximation. The regularization hyperparameters are split into two separate parts: (1) λ_1 penalizes the constant, linear, and input operators, and (2) λ_2 penalizes the quadratic operators. This was done because the quadratic operators are typically of significantly different order of magnitude, thus requiring a different scale of penalization. Regularization in the operator inference learning problem promotes stability of the ROM and helps avoid overfitting. Selecting good values for the hyperparameters λ_1 and λ_2 used in the learning step in Eqn. (3) is the key to effective operator inference regularization. In this work, we follow the regularization hyperparameter selection algorithm described in [20].

B. Parametric operator inference

For many systems of interest, a key question is how the dynamics are influenced by the variation of system parameters. This parametric dependence is incorporated into operator inference by modifying Eqn. (2) to yield

$$\frac{\mathrm{d}}{\mathrm{d}t}\widehat{\mathbf{x}}(t;\mu) = \widehat{\mathbf{c}}(\mu) + \widehat{\mathbf{A}}(\mu)\widehat{\mathbf{x}}(t;\mu) + \widehat{\mathbf{H}}(\mu)\big(\widehat{\mathbf{x}}(t;\mu) \otimes \widehat{\mathbf{x}}(t;\mu)\big) + \widehat{\mathbf{B}}(\mu)\mathbf{u}(t), \tag{4}$$

where $\mu \in \mathbb{R}^p$ is the vector of p parameters. The parametric case is addressed in the original operator inference paper [18] by learning a ROM for each parameter value and interpolating the reduced model operators $\widehat{\mathbf{c}}(\mu), \widehat{\mathbf{A}}(\mu), \widehat{\mathbf{H}}(\mu)$, and $\widehat{\mathbf{B}}(\mu)$ via a cubic spline, for a total of $r + r^2 + r^3 + rm$ interpolations.

We employ a related approach here, except rather than interpolating between reduced model operators, we directly interpolate between the reduced-state solutions. Let the training data for learning the ROM consist of q parameter values $\{\mu_1, \ldots, \mu_q\}$. Then the reduced-state interpolation for parametric operator inference consists of the following steps:

- Assemble the concatenated snapshot matrix $[\mathbf{X}(\mu_1), \ldots, \mathbf{X}(\mu_q)] \in \mathbb{R}^{n \times kq}$ that contains all the snapshots from the FOM results at all q training parameter values.
- Obtain the POD basis by taking the SVD of the concatenated snapshot matrix in the same manner as in Section II.A to get the global reduced basis V_r.

- Learn q local operator inference ROMs at training parameters $\mu_i, i = 1, \ldots, q$ using snapshot matrix $\mathbf{X}(\mu_i)$ and global reduced basis V_r.
- Integrate the q local ROMs to predict the reduced-state solutions $\widehat{\mathbf{X}}(\mu_i) \in \mathbb{R}^{r\times l}, i = i, \ldots, q$ for the desired number of time steps $l \geq k$ involving extrapolation in time.
- Interpolate over the dataset $\left\{\mu_i, \widehat{\mathbf{X}}(\mu_i)\right\}_{i=1}^{q}$ using rl cubic spline interpolations (other regression techniques can also be used) and use it to predict the reduced-state solution $\widehat{\mathbf{X}}(\mu) \in \mathbb{R}^{r\times l}$ at any new parameter value μ for l time steps.
- Use the global reduced basis to reconstruct to the full-order space $\mathbf{V}_r\widehat{\mathbf{X}}(\mu)$.

In a related work in [24], the reduced-state solution is fit using a Gaussian process regression that takes parameters and time as inputs. However, in this paper we do not have time as an input to our regression model. Instead we use the operator inference ROM to extrapolate in time at each training parameter sample, then we interpolate at each time step.

Our parametric ROM method has two specific differences compared to the operator interpolation method in [18]. First, reduced-state solution interpolation can be accomplished with a cubic spline *after* integration of the ODEs and thus without sacrificing stability. In contrast, attempting to interpolate via the reduced model operators directly tends to produce unstable ROMs for large-scale problems. We conclude that to correctly interpolate between operators, manifold interpolation techniques such as interpolation on a Grassman manifold [13, 25] should be investigated. Second, reduced-state solution interpolation never obtains the reduced model operators at the new parameter value, and thus we cannot test many typical stability metrics such as inspecting the sign of the real part of each eigenvalue of the linear operator.

III. Coupled operator inference for flutter

Systems with coupled physics present a particularly challenging problem for reduced-order modeling. We develop a non-intrusive coupled ROM using operator inference by taking advantage of prior knowledge of the form of the coupled system's governing equations. In this paper, we consider the form of the coupled aerostructural equations that arise from combining the governing equations of structural dynamics and fluid dynamics. For full-order solution generation, we use FUN3D's aeroelasticity capability [26] to model the fluid-structure interaction (FSI), and we summarize that method in this section. Then, we describe the process of learning a coupled operator inference ROM specifically for the flutter FSI problem.

The structural dynamics behavior is described by the semi-discrete second-order linear differential equation

$$M\ddot{\delta} + C\dot{\delta} + K\delta = F_{\mathrm{A}}, \tag{5}$$

where M is the mass matrix, C is the damping matrix, K is the stiffness matrix, δ is the vector of nodal displacements, and F_{A} is the aerodynamic forcing applied to the structure. Using modal decomposition to define a reduced basis for the structure uncouples the degrees of freedom. The structural state vector in the reduced space is then given by $\widehat{\mathbf{x}}_{\mathrm{S}} = \begin{bmatrix}\eta & \dot{\eta}\end{bmatrix}^{\top} \in \mathbb{R}^{2r_{\mathrm{s}}}$, where η and $\dot{\eta}$ are the modal displacements and modal velocities of the reduced-order structural system, and r_{s} is the number of reduced-order structural DOFs (modes) retained in the reduced basis. Then we rewrite the reduced structural dynamics equations in the block matrix format as a system of first-order equations,

$$\dot{\widehat{\mathbf{x}}}_{\mathrm{S}}(t) = \begin{bmatrix}0 & I\\ \Omega & Z\end{bmatrix}\begin{bmatrix}\eta(t)\\ \dot{\eta}(t)\end{bmatrix} + \begin{bmatrix}0\\ \widehat{F}_{\mathrm{A}}(t)\end{bmatrix}, \tag{6}$$

where $\Omega = \mathrm{diag}\left(\omega_1^2, \ldots, \omega_{r_{\mathrm{s}}}^2\right)$ is a diagonal matrix consisting of the squares of the natural frequencies ω_i for structural modes $i = 1, \ldots, r_{\mathrm{s}}$, $Z = \mathrm{diag}\left(2\omega_1\zeta_1, \ldots, 2\omega_{r_{\mathrm{s}}}\zeta_{r_{\mathrm{s}}}\right)$, and ζ_i is the damping ratio of structural mode i. The reader is referred to the FUN3D exposition of the aeroelastic modeling capability [26] and to standard texts on structural dynamics [27] for further details.

The fluid dynamics behavior is described by the viscous, compressible, three-dimensional Navier-Stokes equations. The Navier-Stokes equations are typically discretized using conservative variables as the states. Deriving a ROM using the conservative variables does not lead to a quadratic form. However, [28] details how we can use a lifting transformation by introducing the specific volume variable $\xi = \frac{1}{\rho}$, where ρ is the density, to expose quadratic structure for the Navier-Stokes equations. The lifted aerodynamic system has five state variables as defined by $\mathbf{x}_{\mathrm{A}} = \begin{bmatrix}\mathbf{p} & \mathbf{u} & \mathbf{v} & \mathbf{w} & \xi\end{bmatrix}^{\top} \in \mathbb{R}^{5n_x}$ for a spatial discretization with n_x cells, where $\mathbf{p} \in \mathbb{R}^{n_x}$ is the pressure, $\mathbf{u} \in \mathbb{R}^{n_x}$, $\mathbf{v} \in \mathbb{R}^{n_x}$, and $\mathbf{w} \in \mathbb{R}^{n_x}$ are the fluid velocities

in each direction, and $\xi \in \mathbb{R}^{n_x}$ is the specific volume. The non-intrusive nature of the operator inference method allows us to post-process the simulation data to the required form without needing to modify the high-fidelity solver.

In FUN3D, aeroelasticity is implemented by integrating the structural dynamics model in Eqn. (6) via a predictor-corrector scheme [26] coupled with the typical integration of the fluid dynamics. Note that the software assumes small deflections, an acceptable detail given that flutter detection only requires assessment of the damping of an initial transient, not accurate modeling of large-displacement dynamics.

We are now ready to consider the construction of a coupled ROM to represent this flutter model using operator inference. We generate training snapshots from a solve of the full-order FUN3D flutter simulation at a given operating condition. These snapshots consist of full-order fluid state data and reduced structural state (modal displacement) data because the FUN3D flutter simulation uses the reduced-order structural dynamics. After generation of the training snapshots, we project the full-order fluid state $\mathbf{x}_\mathrm{A}(t)$ to the reduced state, $\widehat{\mathbf{x}}_\mathrm{A}(t)$, using the reduced POD basis obtained via SVD of X_A.

The reduced state vector for the coupled FSI model, $\widehat{\mathbf{x}}_\mathrm{FSI}(t)$, is defined as the concatenation of the reduced fluid states $\widehat{\mathbf{x}}_\mathrm{A}(t)$ and the reduced structural states $\widehat{\mathbf{x}}_\mathrm{S}(t)$ as

$$\widehat{\mathbf{x}}_\mathrm{FSI}(t) = \begin{bmatrix} \widehat{\mathbf{x}}_\mathrm{A}(t) \\ \widehat{\mathbf{x}}_\mathrm{S}(t) \end{bmatrix}. \tag{7}$$

Then we learn a linear coupled operator inference ROM of the form

$$\frac{\mathrm{d}}{\mathrm{d}t}\widehat{\mathbf{x}}_\mathrm{FSI}(t) = \widehat{\mathbf{c}}_\mathrm{FSI} + \widehat{\mathbf{A}}_\mathrm{FSI}\widehat{\mathbf{x}}_\mathrm{FSI}(t) \tag{8}$$

where $\widehat{\mathbf{c}}_\mathrm{FSI}$ and $\widehat{\mathbf{A}}_\mathrm{FSI}$ are the reduced operators for the coupled aerostructural flutter system. We perform a single least squares calculation for the entire coupled operator inference ROM. In future work, we plan to extend this methodology by embedding the block sparsity of Eqn. (6) and our knowledge of the quadratic behavior of the fluid dynamics into Eqn. (8).

IV. Applications of parametric and coupled operator inference

We present the application of the operator inference method to the construction of a nonlinear aerodynamic ROM and a coupled aerostructural ROM. Section IV.A presents a parametric operator inference ROM for subsonic flow over the VAT aircraft. Section IV.B presents a coupled operator inference ROM for predicting flutter in the AGARD wing.

A. VAT aircraft: aerodynamic analysis with oscillating angle of attack

The validation of aeroelastic tailoring (VAT) aircraft, shown in Figure 1, is a full-aircraft, 3D model created by General Dynamics in the 1980s for carrying out fighter airplane research and testing [21, 29]. The VAT wing's airfoil transitions from the NACA 64A003.5 at the root to the NACA 64A004 at the tip. The wing has a root chord of 42 inches, a tip chord of 10.75 inches, and a half-span of 39.59 inches as shown in Figure 1c. We use an oscillating angle of attack around 2° with an amplitude of 0.25° as an input to the dynamical system to induce similar unsteady aerodynamics as seen in aircraft flutter. The angle of attack is prescribed by

$$\alpha(t) = 2 + 0.25\sin(200\pi * t),$$

where 200π rad s^{-1} corresponds to 100 Hz. This input is applied as a rigid body pitching motion of the VAT aircraft. The fluid flow is at a Reynolds number of 1.5 million. We study the problem for a range of Mach numbers in $[0.5, 0.8]$.

The full-order model of the VAT aircraft uses NASA's FUN3D software [23] for CFD, along with a mesh discretization generated by Pointwise v18.4R3 *. The mesh is constructed with a spherical far-field of increasing coarseness to reduce computational cost and with a radius of 350 times the root chord length. The unstructured mesh, shown in Figure 2, consists of $n_x = 2,321,731$ cells, each of which will represent five fluid flow state variables leading to $2,321,731 * 5 = 11,608,655$ degrees of freedom. We use NASA's FUN3D software to run a viscous, compressible CFD simulation with the one-equation Spalart-Almaras turbulence model. The time step size is 50 µs, providing a temporal resolution of 200 time steps per oscillation cycle of the angle of attack. Figure 3 shows a representative surface pressure contour over the VAT aircraft and an associated slice of the far-field pressure around the wing, demonstrating the kind of full-order snapshot data that will be generated for use in training and validating the operator inference ROM.

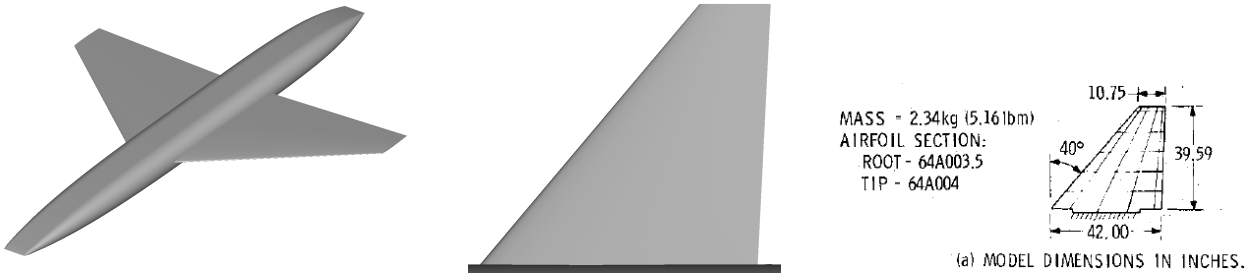

(a) VAT aircraft isometric view. **(b) VAT wing top view.** **(c) Dimensions of the VAT wing [29].**

Fig. 1 VAT aircraft geometry.

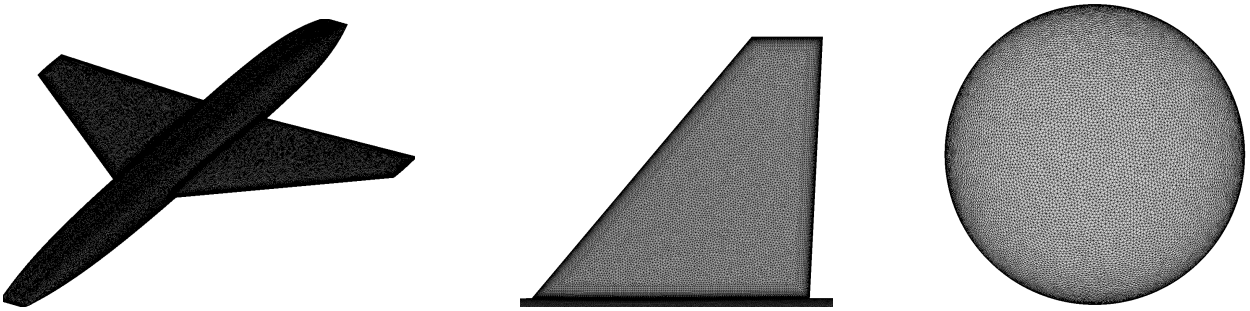

(a) VAT aircraft surface mesh. **(b) VAT wing mesh.** **(c) VAT far-field mesh.**

Fig. 2 VAT aircraft meshes.

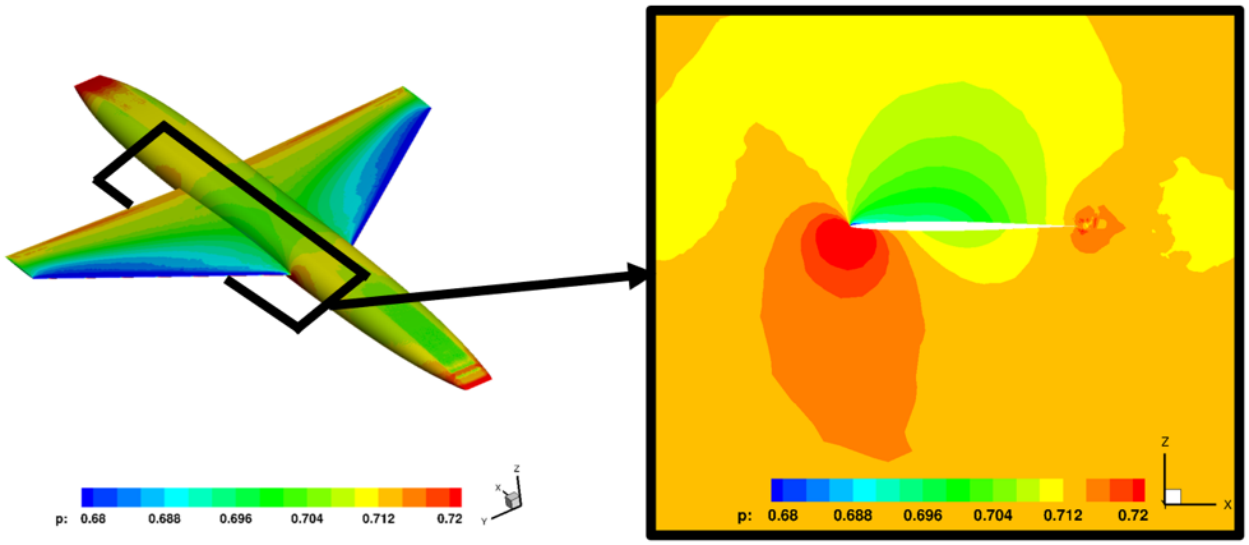

Fig. 3 Pressure contour from full-order FUN3D solution at Mach 0.6, 100 Hz oscillation of angle of attack.

We use the simulations at Mach 0.5, 0.7, and 0.8 as training data for the parametric ROM and the simulation at Mach 0.6 as testing data for assessing the ROMs' predictive capabilities. The snapshots are saved for the lifted state variables $\{p, u, v, w, \xi, \}$ as described in Section III. The training data consists of simulating from 0.25-0.27 s, capturing around two oscillations of the data, for a total of 400 snapshots for each parameter sample in the training set. The size of the global snapshot matrix constructed using concatenated data from the three training parameter samples is

*Pointwise User Manual, v18.4R3

$11,608,655 \times 1200$. The global POD basis vectors are obtained by taking the SVD of the global snapshot matrix. The singular value decay plot for the global snapshot matrix is shown in Figure 4a. We use a basis size of $r = 16$, capturing more than 99.999% of the total energy as shown in Figure 4b.

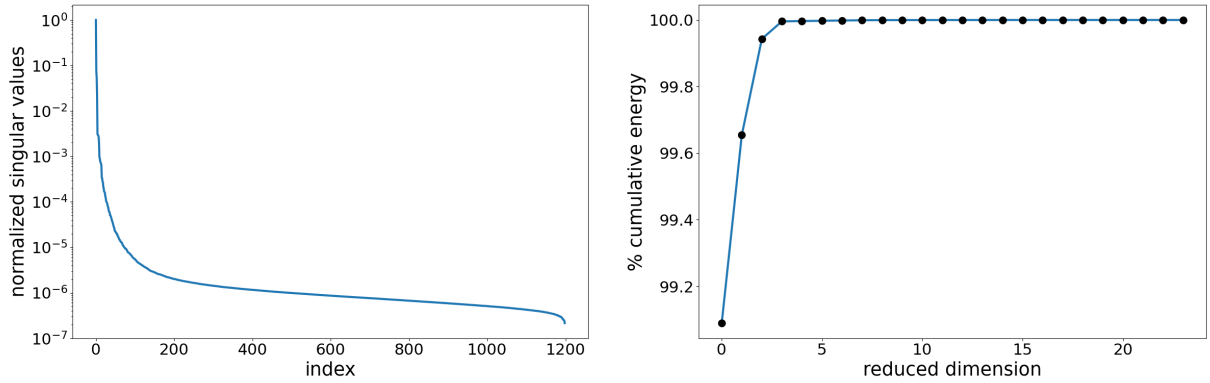

(a) Singular value decay, normalized by largest singular value (b) Cumulative energy retained for a given reduced dimension

Fig. 4 Singular value decay and cumulative energy of VAT global snapshot matrix for Mach $= \{0.5, 0.7, 0.8\}$**.**

We train a parametric operator inference ROM using the reduced-state solution interpolation method described in II.B. We first train the local operator inference ROMs at the three training parameter samples using the local snapshot matrix of size $11,608,655 \times 400$ consisting of 400 time steps. Note that the global reduced basis is used during the learning process. Then we integrate the three local ROMs to predict for a total of 1000 time steps from 0.25-0.3 s, capturing five oscillations of the data while extrapolating in time. The reduced-state solutions of size 16×1000 obtained from each of the three local ROMs are interpolated across the parameter space to capture predictions both *in parameter space and in time* for the parametric ROM. The prediction at the test parameter of Mach 0.6 for 1000 time steps is obtained from the trained parametric ROM and reconstructed to the full-order solution using the global reduced basis. We look at pressure traces at four different locations around the airfoil in one slice of the VAT wing as shown in Figure 5. Figure 6 compares the full-order and reduced-order pressure traces at the four different locations for Mach 0.6. The parametric operator inference ROM yields accurate predictions of the dominant periodic behavior of all four pressure traces for the unseen parameter with small discrepancy in the pressure oscillation amplitude showing good approximation of the FOM. This shows the usefulness of the scientific machine learning method that embeds physics into the learning process as compared to a black-box surrogate method that might not have the ability to learn the oscillatory behavior. Our parametric operator inference ROM actually learned the underlying dynamic feature of the responses in a non-intrusive way

To summarize, a single FUN3D VAT simulation takes about 23 hours to run on 64 cores, with an additional 26 hours needed to write the complete state outputs to file, for a total of about 2 days to complete a single full-order aerodynamic solution at a single operating condition. Once the training snapshots are generated from the FOM, the operator inference method requires approximately 1 hour to complete the POD and learning steps that produce the parametric ROM. After the ROM is built, an additional computational expense on the order of seconds is required to predict the reduced-order aerodynamic solution at a new parameter value (Mach number) over the same temporal domain as the FOM solutions. We see that the expense of generating full-order training snapshots can be on the order of days for large-scale CFD problems, but once the parametric ROM is built, we see speed-ups of at least three orders of magnitude for predictions of fluid state outputs at new parameter values.

B. AGARD wing: aeroelastic flutter

The AGARD wing originated from a NATO working group and is used in this paper to demonstrate coupled ROM flutter prediction capabilities. Flutter occurs when the structural dynamics (i.e. the resonance) of an aircraft's wings and the fluid flow over the surface of the wings become unstably coupled, typically leading to catastrophic vehicle failure. This makes avoiding flutter a critical consideration in aircraft design. However, it is expensive both to test and to model flutter accurately, due to the inherent multiphysics coupling of the problem. We show the potential of the coupled ROM method discussed in Section III to accurately predict the onset of flutter in an aerostructural model at much lower

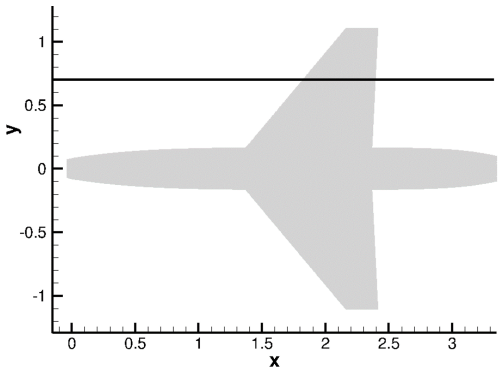

(a) Spanwise location.

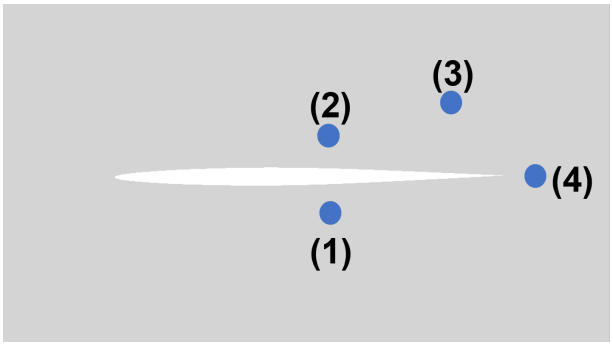

(b) Location relative to spanwise slice of airfoil.

Fig. 5 Locations of pressure traces for VAT aircraft model.

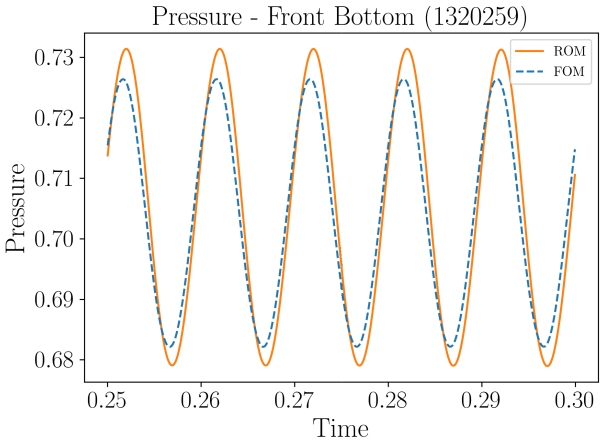

(a) Trace location 1

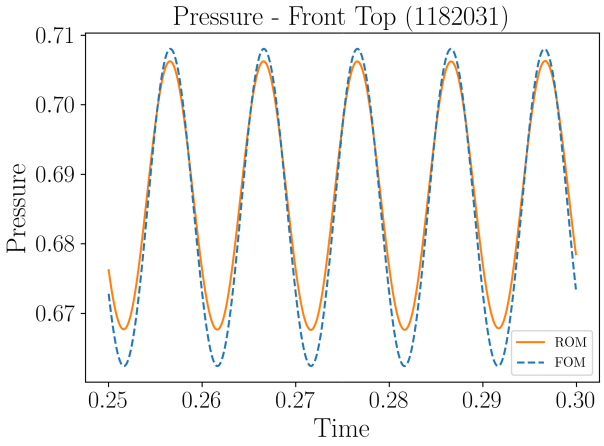

(b) Trace location 2

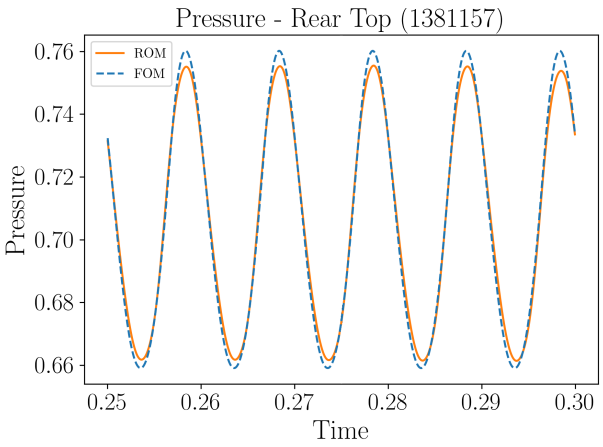

(c) Trace location 3

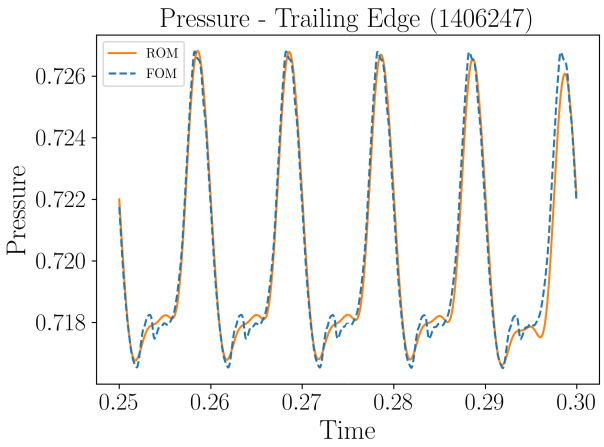

(d) Trace location 4

Fig. 6 Comparison of pressure traces through time for Mach 0.6 obtained from the FOM and ROM ($r = 16$) shown at four different positions in an airfoil slice of the wing.

computational cost than a full-order, nonlinear CFD solver (in our case, FUN3D).

The AGARD wing has an aspect ratio of 1.6525, a taper ratio of 0.6576, and a sweepback angle of 45 degrees [22]. Similar to the VAT aircraft wings, the AGARD wing has a NACA65A004 airfoil. We use the unstructured mesh of

the AGARD wing from the FUN3D example problem materials [†] as shown in Figure 7. FUN3D's built-in aeroelastic modeling feature is used to generate the full-order solutions. The aeroelastic simulation requires the results of a modal decomposition from a linear structural dynamics model to efficiently represent the structural behavior of the AGARD wing. Specifically, we use the first four AGARD wing mode shapes provided by the NASA FUN3D developers. These mode shapes are visualized in Figure 8 along with their associated natural frequencies. For more details about FUN3D's aeroelastic modeling capability, we refer the reader to the original exposition of these capabilities in [26].

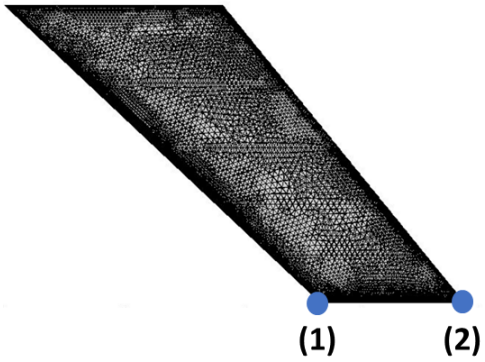

Fig. 7 AGARD wing CFD mesh with pressure trace locations.

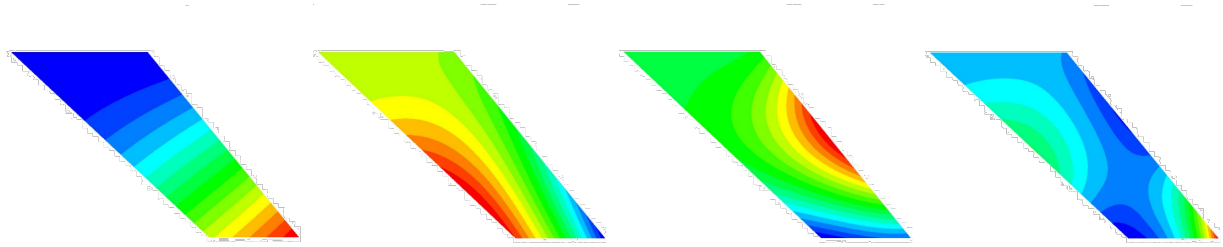

(a) Mode 1 (9.6 Hz) **(b) Mode 2 (38.2 Hz)** **(c) Mode 3 (48.3 Hz)** **(d) Mode 4 (91.5 Hz)**

Fig. 8 First four structural mode shapes of the fixed-boundary AGARD wing, scaled independently for visualization

To analyze a flutter case in FUN3D, we select a Mach number and a dynamic pressure, then we perturb the system by prescribing an nonzero initial generalized modal velocity for each of the structural mode shapes. Then the generalized modal displacement is monitored over time to see if the modal displacements settle to steady-state values or continue to increase and become unstable, which indicates that flutter has occurred. In this study, we use a viscous, compressible fluid flow model. The CFD problem is solved for five fluid flow state variables over $n_x = 492,778$ cells leading to $492,778 * 5 = 2,463,890$ degrees of freedom. The procedure is repeated at different Mach number - dynamic pressure combinations to characterize the flutter boundary. In Figure 9, we see the modal displacement of the first mode at Mach 0.9 and dynamic pressures $Q = \{65, 75, 85, 95\}$ psf for the full-order FUN3D results. Note that the orange curve for 75 psf is nearly in a limit cycle, indicating that this operating condition is near the flutter boundary. The green and red curves are growing in amplitude over time, indicating flutter.

Although it is not the key quantity of interest for flutter, it is informative to look at how well the ROM can replicate the fluid dynamics response of the system as well, since it is the coupling between the structural response and the fluid response that explains why a wing flutters under particular operating conditions. Figure 10 shows an example of the kind of surface pressure profile that occurs in the full-order FUN3D results.

This is an interesting demonstration case for our non-intrusive coupled ROM described in Section III because we do not have explicit access to the original mesh of the AGARD wing or even the structural dynamics code that was used to

[†]FUN3D v13.4 Training - Session 16: Aeroelastic Simulations: `https://fun3d.larc.nasa.gov/session16_2018.pdf`

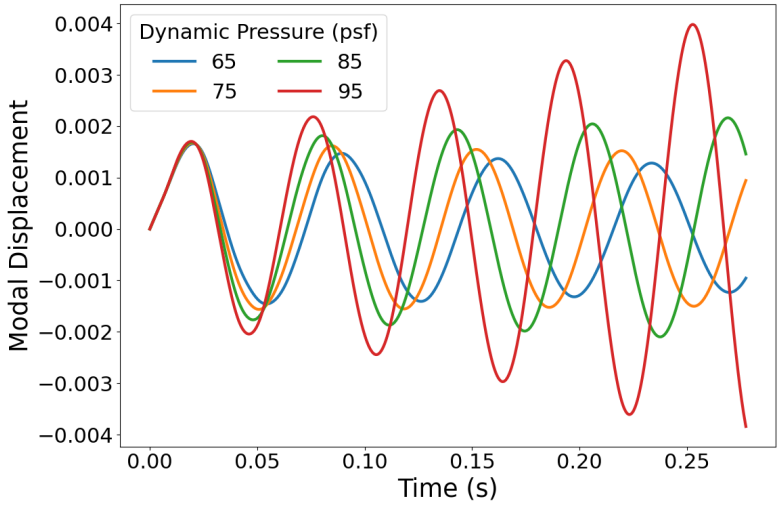

Fig. 9 Modal displacement of first AGARD mode shape for Mach 0.9 and dynamic pressures $Q = \{65, 75, 85, 95\}$ psf

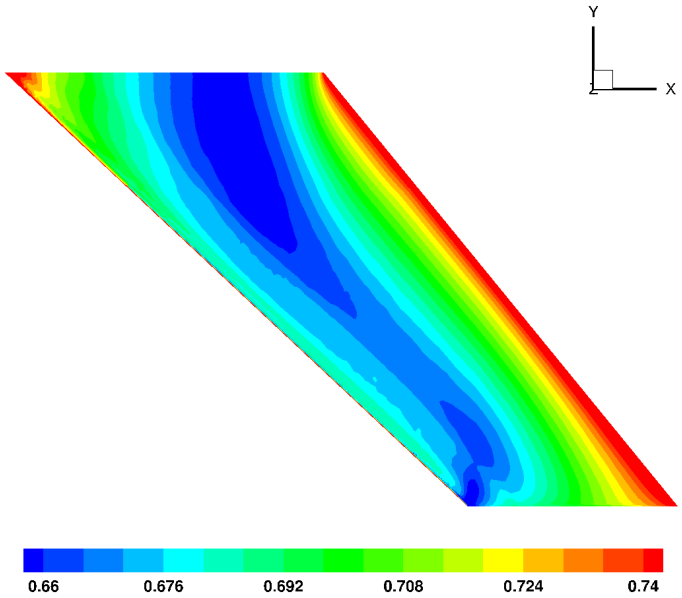

Fig. 10 Representative pressure contour on surface of AGARD wing at Mach 0.9, dynamic pressure 75 psf.

obtain the modal results. However, we can still learn a coupled ROM since it is completely non-intrusive and we can work with legacy data. The mode shapes can be used as additional reduced basis functions that we add directly to the set of bases we obtain from POD after running a FUN3D coupled solve. Thus we are able to very easily incorporate prior knowledge of the physics of the system into our model reduction method when this knowledge becomes available, without requiring any additional work on the part of the simulation analyst. We exploit the structure of the coupled aerostructural governing equations, which consist of linear structural dynamics equations and quadratic fluid dynamics (Navier-Stokes) equations, to learn the coupled ROM for flutter prediction. In this work, we learn a coupled operator inference ROM with linear form. We use basis functions from both the structural dynamics modal decomposition (i.e. an eigenvalue problem) and the proper orthogonal decomposition of the fluid state snapshot matrix. We will extend the coupled ROM to include quadratic terms for the fluid flow equations in a future effort. The fluid flow snapshot matrix consists of training data from 400 time steps with a time step size of 0.277 ms, leading to a training snapshot matrix of

size $2,463,890 \times 400$. The singular value decay plot for the fluid flow training snapshot matrix is shown in Figure 11a. We use a basis size of 18 for fluid flow POD, capturing 99.995% of the total energy as shown in Figure 11b . The total size of the reduced state vector for the coupled ROM ends up being $18 + 2 * 4 = 26$ after including four structural modes (each mode including a modal displacement and a modal velocity).

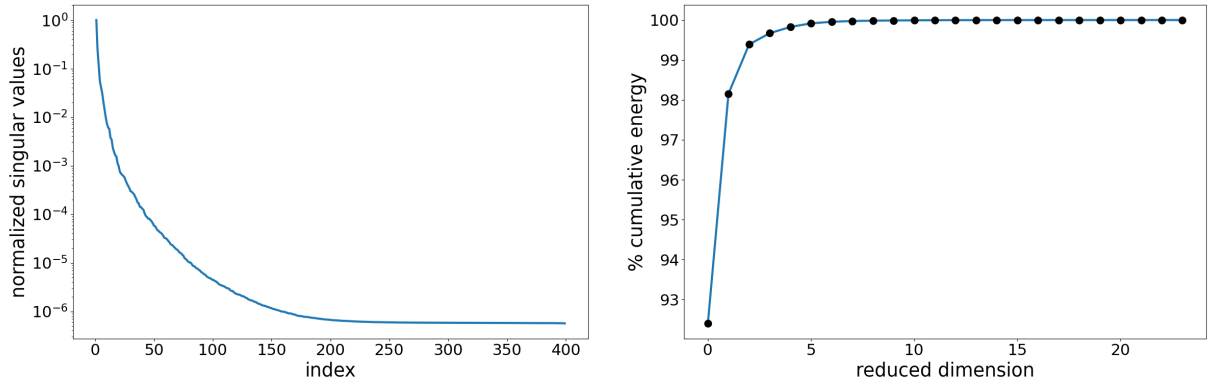

(a) Singular value decay, normalized by largest singular value (b) Cumulative energy retained for a given reduced dimension

Fig. 11 Singular value decay and cumulative energy of snapshot matrix.

Figure 12 compares the FOM and the coupled ROM predictions for a non-fluttering operating condition (Q = 75 psf) and a fluttering operating condition (Q = 95 psf). For the modal displacement plots in Figures 12a and 12c, the dashed black lines representing the coupled ROM predictions are overlaid almost exactly on the solid lines representing the full-order predictions for all four modes. In Figures 12b and 12d, we can see that absolute error between the FOM and the coupled ROM predictions is below 2% for the non-fluttering case of Q = 75 psf and below 5% for the fluttering case of Q = 95 psf. Typically, the final step of the flutter analysis would be to use a method such as the log decrement to numerically estimate the damping of the modal displacements at a given operating condition to decide if the case is fluttering. For that conventional workflow, the temporal extrapolation we see in Figure 12 using the coupled ROM is accurate for a sufficiently long duration to calculate damping and predict flutter. We also investigate the response of fluid states such as the pressure during aerostructural coupling. In Figure 13, we see the pressure trace at the leading edge and the trailing edge of the wing tip (see Figure 7) for the nonfluttering and fluttering cases. The coupled ROM is able to accurately track the higher frequency oscillations, especially in the trailing edge plots.

A single FUN3D flutter simulation takes 25 minutes to run on 32 cores, with an additional 3 minutes for writing data outputs to file, for a total of 28 minutes to complete a full-order flutter solution at a single operating condition. Once the training snapshots are generated from the FOM, the operator inference method requires approximately 1.5 minutes to complete the POD and learning steps that produce the ROM. After the ROM is built, an additional 3 seconds of computation are required to predict the reduced-order flutter solution over the same temporal domain as the FOM. Thus, we see that the operator inference ROM provides a run-time speed up from minutes for the FOM to seconds for the ROM, with the added ability to rapidly and accurately extrapolate beyond the temporal domain of the FOM results.

V. Conclusion

Operator inference is a data-driven, physics-informed model reduction method for dynamical systems. The non-intrusive nature of the operator inference method makes it an appealing strategy for generating ROMs for commercial solvers and for legacy solvers or data since it only requires access to the state solutions generated by the solvers and not to the discretized operators of the underlying governing equations. In this paper, we develop a parametric ROM for large-scale aerodynamic analysis and a coupled ROM for flutter analysis based on operator inference. We lift the Navier-Stokes equations governing the fluid flow problem to a quadratic form for operator inference and use a reduced-state solution interpolation method to build the parametric ROM. We show that the parametric ROM is able to accurately extrapolate in time and parameter space to predict the dynamical behavior for a large-scale aerodynamic problem using the VAT aircraft at unseen Mach numbers. Then we show the coupled ROM for coupled multiphysics applications, where we often have some type of physical insight into at least one of the governing physical regimes *a*

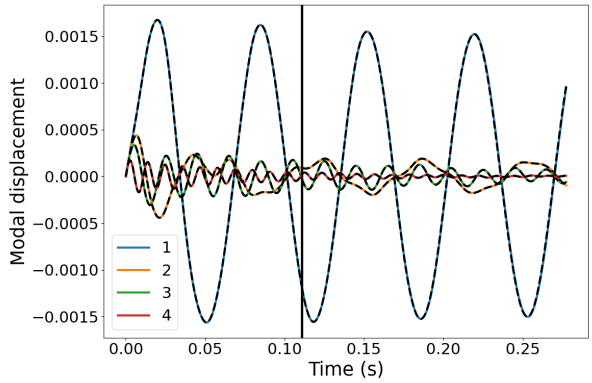

(a) Q=75, modal displacement (solid lines: FOM; dashed lines: ROM)

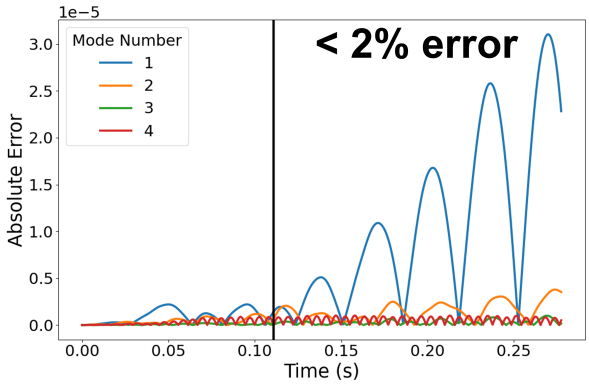

(b) Q=75, absolute error in modal displacement

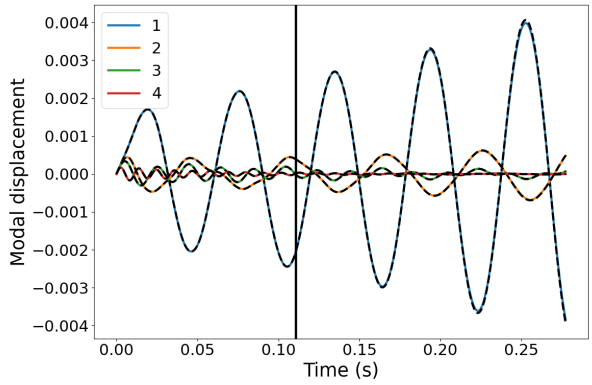

(c) Q=95, modal displacement (solid lines: FOM; dashed lines: ROM)

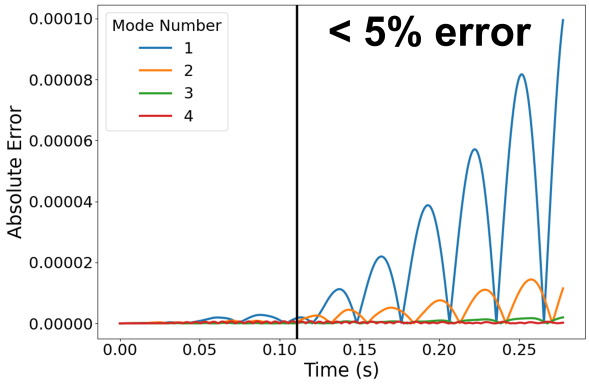

(d) Q=95, absolute error in modal displacement

Fig. 12 Comparison of modal displacement for four modes between the FOM and the coupled ROM at Mach=0.9. Training data and testing data split is shown by vertical black line.

priori. In the case of the flutter analysis, this prior knowledge is the modal decomposition information which, similar to many typical industry workflows, was provided to us without access to the initial full-order model. We are able to learn a non-intrusive coupled ROM, which is capable of predicting the key quantity of interest of generalized modal displacement to within less than 5% absolute error. Both the parametric ROM and the coupled ROM lead to significant speedups compared to the high-fidelity simulations. Our future work will combine these capabilities (parametric and coupled ROMs) for aeroelastic systems to create parametric coupled ROMs capable of predicting the onset of flutter at unseen dynamic pressures and Mach numbers. This effort would be supported by more robust development of the learning step of the operator inference method for coupled ROMs, in order to incorporate knowledge of the block-sparsity of the second-order coupled governing equations.

Acknowledgements

This work was supported in part by Lockheed Martin University Research Agreement MRA-16-005-RPP012-001 and by the AFOSR MURI on physics-based machine learning under grant FA9550-21-1-0084.

References

[1] Sirovich, L., "Turbulence and the dynamics of coherent structures. I. Coherent structures," *Quarterly of Applied Mathematics*, Vol. 45, No. 3, 1987, pp. 561–571.

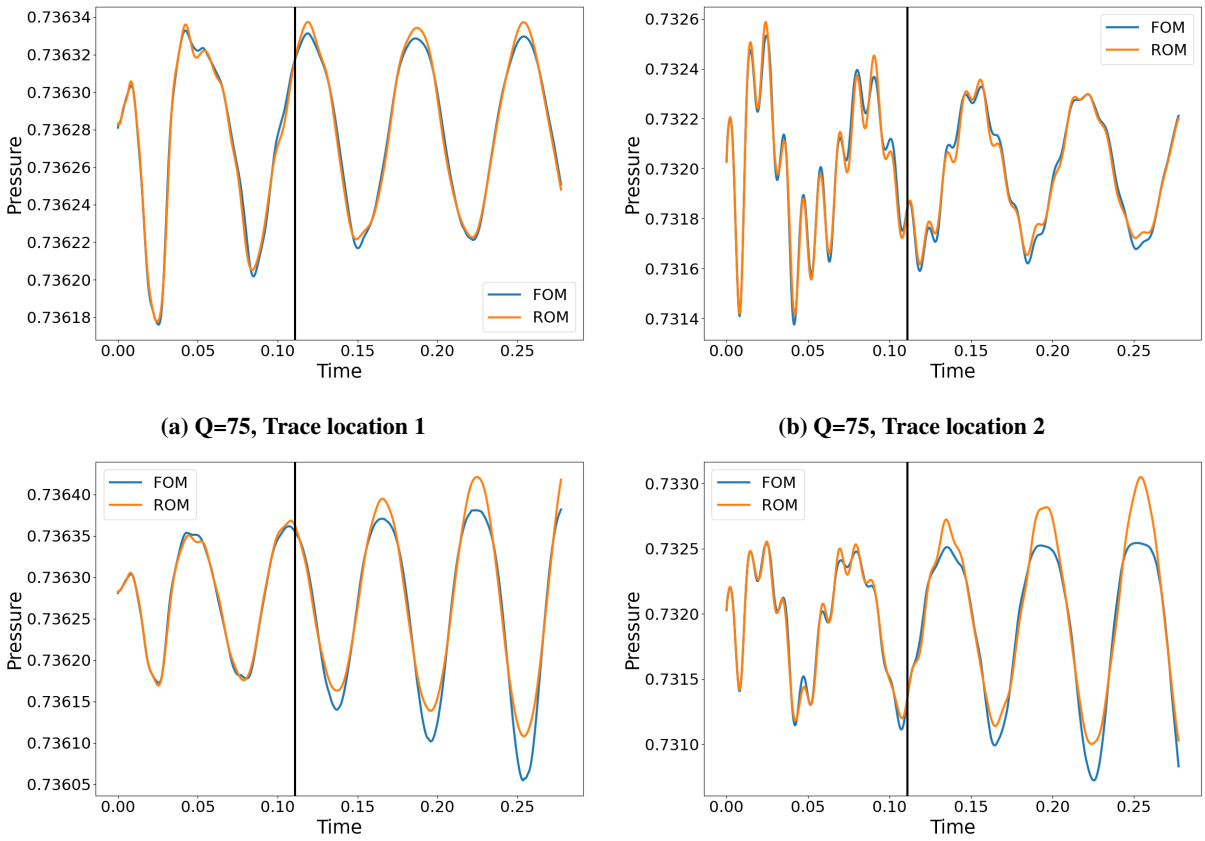

(a) Q=75, Trace location 1

(b) Q=75, Trace location 2

(c) Q=95, Trace location 1

(d) Q=95, Trace location 2

Fig. 13 Operator Inference pressure predictions at Mach=0.9 at the wing tip. Training data and testing data split is shown by vertical black line.

[2] Berkooz, G., Holmes, P., and Lumley, J. L., "The proper orthogonal decomposition in the analysis of turbulent flows," *Annual Review of Fluid Mechanics*, Vol. 25, 1993, pp. 539–575.

[3] Lumley, J. L., "The Structures of Inhomogeneous Turbulent Flow," *Atmospheric Turbulence and Radio Wave Propagation*, 1967, pp. 166–178.

[4] Rathinam, M., and Petzold, L., "A New Look at Proper Orthogonal Decomposition," *SIAM Journal on Numerical Analysis*, Vol. 41, No. 5, 2003, pp. 1893–1925. https://doi.org/10.1137/S0036142901389049.

[5] Benner, P., Güğercin, S., and Willcox, K., "A survey of projection-based model reduction methods for parametric dynamical systems," *SIAM Review*, Vol. 57, No. 4, 2015, pp. 483–531. https://doi.org/10.1137/130932715.

[6] Antoulas, A. C., *Approximation of Large-Scale Dynamical Systems*, SIAM, Philadelphia, PA, 2005.

[7] Schmid, P. J., "Dynamic mode decomposition of numerical and experimental data," *Journal of Fluid Mechanics*, Vol. 656, 2010, p. 5–28. https://doi.org/10.1017/S0022112010001217.

[8] Tu, J. H., Rowley, C. W., Luchtenburg, D. M., Brunton, S. L., and Kutz, J. N., "On dynamic mode decomposition: Theory and applications," *Journal of Computational Dynamics*, Vol. 1, No. 2, 2014, pp. 391–421. https://doi.org/10.3934/jcd.2014.1.391.

[9] Theodorsen, T., "General Theory of Aerodynamic Instability and the Mechanism of Flutter," *Annual Report of the National Advisory Committee for Aeronautics*, Vol. 268, 1935, p. 413.

[10] Dowell, E. H., *A Modern Course in Aeroelasticity*, 6th ed., Springer, Durham, NC, 2021.

[11] Hall, K. C., "Eigenanalysis of Unsteady Flows about Airfoils, Cascades, and Wings," *AIAA Journal*, Vol. 32, No. 12, 1994, pp. 2426–2432. https://doi.org/10.2514/3.12309.

[12] Dowell, E. H., and Hall, K. C., "Modeling of Fluid-Structure Interaction," *Annual Review of Fluid Mechanics*, Vol. 33, 2001, pp. 445–490. https://doi.org/10.1146/annurev.fluid.33.1.445.

[13] Amsallem, D., and Farhat, C., "Interpolation Method for the Adaptation of Reduced-Order Models to Parameter Changes and Its Application to Aeroelasticity," *AIAA Journal*, Vol. 46, No. 7, 2008, pp. 1803–1813. https://doi.org/10.2514/1.35374.

[14] Silva, W. A., "Simultaneous Excitation of Multiple-Input/Multiple-Output CFD-Based Unsteady Aerodynamic Systems," *Journal of Aircraft*, Vol. 45, No. 4, 2008, pp. 1267–1274. https://doi.org/10.2514/1.34328.

[15] Juang, J.-N., and Pappa, R. S., "An Eigensystem Realization Algorithm for Modal Parameter Identification and Model Reduction," *Journal of Guidance, Control, and Dynamics*, Vol. 8, No. 5, 1985, pp. 620–627. https://doi.org/10.2514/3.20031.

[16] Silva, W. A., "AEROM: NASA's Unsteady Aerodynamic and Aeroelastic Reduced-Order Modeling Software," *Aerospace*, Vol. 5, No. 2, 2018. https://doi.org/10.3390/aerospace5020041.

[17] Lowe, B. M., and Zingg, D. W., "Efficient Flutter Prediction Using Reduced-Order Modeling," *AIAA Journal*, Vol. 59, No. 7, 2021, pp. 2670–2683. https://doi.org/10.2514/1.J060006.

[18] Peherstorfer, B., and Willcox, K., "Data-driven Operator Inference for Nonintrusive Projection-Based Model Reduction," *Computer Methods in Applied Mechanics and Engineering*, Vol. 306, 2016, pp. 196–215. https://doi.org/10.1016/j.cma.2016.03.025.

[19] Qian, E., Kramer, B., Peherstorfer, B., and Willcox, K., "Lift & Learn: Physics-informed machine learning for large-scale nonlinear dynamical systems," *Physica D: Nonlinear Phenomena*, Vol. 406, 2020, p. 132401. https://doi.org/https://doi.org/10.1016/j.physd.2020.132401.

[20] McQuarrie, S. A., Huang, C., and Willcox, K. E., "Data-driven reduced-order models via regularised Operator Inference for a single-injector combustion process," *Journal of the Royal Society of New Zealand*, Vol. 51, No. 2, 2021, pp. 194–211. https://doi.org/10.1080/03036758.2020.1863237.

[21] Rogers, W., Braymen, W., Murphy, A., Graham, D., and Love, M., "Validation of aeroelastic tailoring by static aeroelastic and flutter tests," Tech. Rep. AFWAL-TR-81-3160, General Dynamics Corporation, Fort Worth, Texas, 1982.

[22] Yates, E. C. J., "AGARD Standard Aeroelastic Configurations for Dynamic Response. Candidate Configuration I.-Wing 445.6," Tech. rep., NASA Langley Research Center, Hampton, Virginia, 08 1987.

[23] Biedron, R. T., Carlson, J.-R., Derlaga, J. M., Gnoffo, P. A., Hammond, D. P., Jones, W. T., Kleb, B., Lee-Rausch, E. M., Nielsen, E. J., Park, M. A., Rumsey, C. L., Thomas, J. L., Thompson, K. B., and Wood, W. A., *FUN3D Manual 13.6*, NASA Langley Research Center, Hampton, Virginia, 2019.

[24] Guo, M., and Hesthaven, J. S., "Data-driven reduced order modeling for time-dependent problems," *Computer Methods in Applied Mechanics and Engineering*, Vol. 345, 2019, pp. 75–99. https://doi.org/10.1016/j.cma.2018.10.029.

[25] Absil, P.-A., Mahony, R., and Sepulchre, R., "Riemannian Geometry of Grassmann Manifolds with a View on Algorithmic Computation," *Acta Applicandae Mathematica*, Vol. 80, No. 2, 2004, pp. 199–220. https://doi.org/10.1023/B:ACAP.0000013855.14971.91.

[26] Biedron, R., and Thomas, J., "Recent enhancements to the FUN3D flow solver for moving-mesh applications," *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, 2009, p. 1360. https://doi.org/10.2514/6.2009-1360.

[27] Ginsberg, J. H., *Mechanical and Structural Vibrations*, John Wiley & Sons, Inc., New York, NY, 2001.

[28] Swischuk, R., Kramer, B., Huang, C., and Willcox, K., "Learning Physics-Based Reduced-Order Models for a Single-Injector Combustion Process," *AIAA Journal*, Vol. 58, No. 6, 2020, pp. 2658–2672. https://doi.org/10.2514/1.J058943.

[29] Ruhlin, C. L., Watson, J. J., Ricketts, R. H., and Doggett, R. V., "Evaluation of Four Subcritical Response Methods for On-Line Prediction of Flutter Onset in Wind Tunnel Tests," *Journal of Aircraft*, Vol. 20, No. 10, 1983, pp. 835–840. https://doi.org/10.2514/3.44951.