Copyright
by
Stefanie Joyce Salinger
2021

The Thesis Committee for Stefanie Joyce Salinger
certifies that this is the approved version of the following thesis:

Toward Predictive Digital Twins for Self-Aware Unmanned Aerial Vehicles: Non-Intrusive Reduced Order Models and Experimental Data Analysis

APPROVED BY

SUPERVISING COMMITTEE:

Karen Willcox, Supervisor

James Koch

Toward Predictive Digital Twins for Self-Aware Unmanned Aerial Vehicles: Non-Intrusive Reduced Order Models and Experimental Data Analysis

by

Stefanie Joyce Salinger

THESIS

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

MASTER OF SCIENCE IN ENGINEERING

THE UNIVERSITY OF TEXAS AT AUSTIN

May 2021

Acknowledgments

I wish to thank the multitude of people who helped me in the process of completing this thesis.

First, I want to thank my advisor, Dr. Karen Willcox, for bringing me into her research group and providing continual guidance, expertise, and support. I would also like to thank Dr. James Koch for his consistent assistance and availability throughout the CFD and reduced order modeling process. Additionally, I would like to thank Michael Kapteyn for his consistent support throughout the data collection and analysis process and for his expertise on digital twins. I am also grateful to Jaco Pretorius for his help on the hardware and experimental side, and Shane McQuarrie for his help and expertise on the operator inference approach. I also wish to acknowledge Cory Kays and Aurora Flight Sciences for the development of the testbed aircraft.

This work was supported in part by a Cockrell School of Engineering graduate fellowship. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1610403. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

I want to also express my thanks to my friends and colleagues in the

Aerospace Engineering Department and the Willcox Research Group for all of their assistance and support throughout my degree.

Finally, I am immensely thankful to my parents, Jeanne and Jeff Salinger, and to the rest of my family for their unending love, support, and encouragement throughout my educational endeavours.

Toward Predictive Digital Twins for Self-Aware Unmanned Aerial Vehicles: Non-Intrusive Reduced Order Models and Experimental Data Analysis

Stefanie Joyce Salinger, M.S.E.
The University of Texas at Austin, 2021

Supervisor: Karen Willcox

The concept of the Digital Twin describes the use of comprehensive and authoritative digital models tailored to a unique physical asset that dynamically adapt as the asset evolves over time and are able to inform valuable decisions. A key challenge is to make a digital twin truly predictive so that it can be used to drive high-consequence decisions with quantified confidence. Currently, this can only be achieved through high-fidelity physics-based models. These models are computationally expensive to solve and prohibitive to use in a real-time context. Reduced-order modeling methods have emerged as a powerful tool for enabling high-fidelity simulations with computationally efficient models. This thesis aims to advance the framework for a predictive digital twin for unmanned aerial vehicles using physics-based models and scientific machine learning, as well as hardware experimentation. In particular, this work develops and demonstrates non-intrusive projection-based reduced-order modeling strategies for aerodynamic loading that can be applied to an unmanned aerial vehicle (UAV). In addition, this research presents an experimental data collection and analysis methodology to further the evolution of a digital-twin-enabled self-aware UAV. Self-awareness in this context refers to the ability of the vehicle to collect information about itself and its surroundings and to use this information to alter the way it completes missions via on-board dynamic decision-making. Emulation of wing damage states on the hardware testbed produces data sets that can be used in conjunction with previously developed computational methods in order to enable classification of the UAV structural state in flight. The two-way coupling between estimation of the UAV structural state and dynamic mission replanning is a capability that is critical for realizing the self-aware UAV concept.

Table of Contents

Acknowledgments **iv**

Abstract **vi**

List of Tables **x**

List of Figures **xi**

Chapter 1. Introduction **1**

1.1 Motivation . . . 1
1.2 Reduced Order Modeling . . . 3
1.3 Predictive Digital Twins . . . 5
1.4 Scientific Machine Learning via Operator Inference . . . 7
1.5 Hardware Testbed . . . 9
1.6 Research Objectives . . . 13

Chapter 2. Reduced Order Models via Operator Inference **15**

2.1 High Angle of Attack Unsteady Flow Case . . . 15
2.2 Creation of Reduced Order Model . . . 18
2.2.1 Setup . . . 18
2.2.2 Determining ROM Dimension . . . 19
2.2.3 Regularization . . . 21
2.3 Results . . . 23

Chapter 3. Experimental Data Collection and Analysis **31**

3.1 Experimental Setup and Sensor Technology . . . 31
3.2 Proof-of-Concept Results . . . 33

Chapter 4. Conclusions **38**
4.1 Research Summary 38
4.2 Future Research Directions 39

Appendices **41**

Appendix A. Additional Snapshot Results **42**
A.1 Pressure Contours 42
A.2 Pressure Distributions 47

References **49**

List of Tables

2.1 Properties of Unsteady Flow Case 17
2.2 Optimized regularization parameters for increasing ROM sizes. 22

3.1 Material properties of different access panel cases. 34

List of Figures

1.1 Digital twin approach: data-driven adaptation of component-based reduced-order models (Figure credit: M. Kapteyn). . . . 6

1.2 Schematic of the designed and manufactured wing structure. . 11

1.3 The custom-built self-aware UAV hardware testbed. 11

1.4 Flight path recorded during the maiden test flight. 12

2.1 NACA 0012 airfoil at angle of attack of 45°. 16

2.2 CFD domain mesh showing (a) the mesh out to the farfield and (b) the mesh zoomed in on the airfoil. 17

2.3 Rank of the data matrix for increasing ROM sizes. 20

2.4 Cumulative energy of the singular values. 21

2.5 Predicted and actual pressure contours at snapshots. 25

2.6 Pressure absolute error contours at snapshots. 26

2.7 Absolute and relative ℓ^2 errors for the state variables. 27

2.8 Predicted and actual pressure distribution with errors at snapshots. 28

2.9 Predicted and actual lift (same as induced drag) over time. . . 29

2.10 Predicted and actual lift (same as induced drag) over time; solid black vertical line indicates boundary between training and testing regimes. 30

3.1 Experimental setup and wireless sensor used for data collection. 32

3.2 Preliminary experimental results. Top: Raw vibration sensor output for each access panel case. Middle: High-pass filtered vibration sensor output. Vertical offsets are added to better show the difference between cases. Bottom: Integrated high-pass filtered sensor data (two repeated trials for each case). . . 35

3.3 Information model for the proposed architecture. 37

A.1 Continued below. 42

A.1 Continued below. 43

A.1 Continued below. 44
A.1 Continued below. 45
A.1 Predicted and actual pressure contours at additional snapshots 46
A.2 Continued below. 47
A.2 Predicted and actual pressure distribution with errors at additional snapshots. 48

Chapter 1

Introduction

This chapter introduces the motivation behind research supporting digital twins and their potential applications, and provides background on reduced-order modeling concepts. Then, a description of the existing work that is furthered in this thesis is given. These topics include the predictive digital twin framework, the operator inference scientific machine learning approach, and the developed hardware testbed. Based on this work, the research objectives for this thesis are presented.

1.1 Motivation

A digital twin is a comprehensive and authoritative model of a physical asset that is tailored to a unique physical asset and adapts as the asset evolves over time [1]. Predictive digital twin models for aircraft systems are required to be predictive, reliable, and explainable. They must simulate previously unseen scenarios, obey the laws of physics, have quantifiable uncertainty, and utilize real-world operational parameters and quantities. High-fidelity multi-physics air vehicle models substantiated by experimental ground and flight test data serve as the basis for digital twin development. However, high-fidelity models and large experimental databases are too complex and computationally expensive for direct incorporation into a digital twin or for use in near real-time decision making.

A recent approach [2] has shown that reduced-order modeling provides a mathematical foundation for creating a predictive digital twin. A predictive digital twin must be built upon a combination of data and physics based models. In the digital twin setting, models must scale efficiently to the full system level, admit an expressive parameterization of system properties (geometry, material properties, etc.), and be fast to update. The high-fidelity physics-based models used in the system design phase, such as computational fluid dynamics (CFD) and structural finite element method (FEM) models, are too complex and computationally expensive to meet these requirements. Reduced-order modeling is a mathematical approach to deriving surrogates that retain much of the accuracy of the high-fidelity models, but via a lower dimensional model that is much faster to solve. Thus, reduced-order modeling is a key computational technology to operationalizing complex physics-based models in the digital twin context.

The Digital Twin paradigm has seen increasing attention in recent years. Digital twins can serve as an authoritative information source for a unique physical asset, which facilitates accurate coupling and fusion of multiple models [3, 4] and ensures that current relevant information is available to stakeholders. In addition, traditional approaches to aircraft certification, fleet management, and sustainment rely on statistical distributions of material properties, heuristic design philosophies, and assumed similarities between physical testing and operational conditions [4]. Thus, digital twin applications have been proposed as a replacement and/or complement to conventional engineering practices in structural health monitoring and aircraft sustainment [4, 5], simulation-based vehicle certification [4, 3], and fleet management [4, 6]. Additionally, outside of aerospace engineering, promising digital twin application areas include healthcare [7], manufacturing [8], smart infrastructure [9], and education [10]. Specifically in aircraft structural health monitoring, proposed digital twin systems would be capable of mitigating damage or degradation by activating self-healing mechanisms or by recommending changes in mission profile to decrease loadings and increase both the life span and probability of mission success [4]. Inspired by all of these works, an approach has recently been proposed to enable the concept of a self-aware unmanned aerial vehicle (UAV) by constructing a predictive digital twin of the vehicle [2]. A self-aware aerospace vehicle is one that can leverage online sensor data to dynamically gather information about its structural health, and respond intelligently by replanning its mission [11, 12].

1.2 Reduced Order Modeling

Accurately modeling a complete engineering system, as is needed to create a digital twin, is possible with high-fidelity computational models. Physics-based models based on discretized partial differential equations (PDEs), such as CFD and finite-element analysis (FEA), can achieve this. In contrast with purely data-driven models, like neural networks, physics-based models offer a greater degree of interpretability, reliability, and predictive capability. However, a CFD and FEA based digital twin then results in large computational models that require significant high performance computing resources to solve. In many applications, digital twins are required to provide near real-time insights in order for them to be used effectively for operational decision making. This requires the ability to rapidly adapt the computational model in the face of changing model parameters, and rapidly evaluate the model to provide analysis and prediction. The use of purely high-fidelity physics-based models quickly becomes computationally intractable in this type of real-time, many-query context.

Model order reduction [13, 14, 15, 16] provides a mathematical foundation for accelerating complex computational models so that they may be operationalized in the digital twin context. Reduced-order modeling involves investing upfront computational time during a training phase and the resulting reduced-order models (ROMs) can then be rapidly evaluated over parametric sweeps or during an online operational phase. ROMs combine the rich information embedded in high-fidelity simulations with the efficiency of low-dimensional surrogate models. Projection-based model reduction considers the class of problems for which the governing equations are known and for which a high-fidelity model is available [13, 14]. The goal is to derive a ROM that has lower complexity and yields accurate solutions with reduced computation time. Projection-based approaches define a low-dimensional manifold on which the dominant dynamics evolve. However, the intrusive nature of these methods has limited their impact in practical applications. Non-intrusive model reduction methods instead derive the ROMs by fitting to simulation data without requiring explicit access to the high-fidelity model operators. These methods are powerful and often yield good results, but since the approximations are based on generic data-fit representations, they lack the mathematical theory that would enable a user to determine whether a ROM can issue reliable predictions in areas outside of the training data. Therefore, recent work has formulated the ROM task through the lens of projection-based model reduction, but with a non-intrusive operator inference that learns the ROMs directly from simulation output data without needing access to source code [17, 18, 19].

1.3 Predictive Digital Twins

The current digital twin in this approach is based on a high-fidelity component-based reduced-order structural model of the airframe described in detail in [2]. This model is capable of simulating the structural response of the airframe and characterizing the structural limits of the aircraft in a range of different structural damage states. The digital twin is enabled by dynamically updating this structural model based on structural sensor data, and then using the updated structural model for rapid analysis and prediction. Adaptation of the digital twin is achieved by training interpretable machine learning classifiers on the output of the structural model, and then using these classifiers online to rapidly infer which structural state best explains the observed sensor data [20].

As shown in Figure 1.1, a library of component-based ROMs is constructed during an offline phase. Building these ROMs requires training snap-

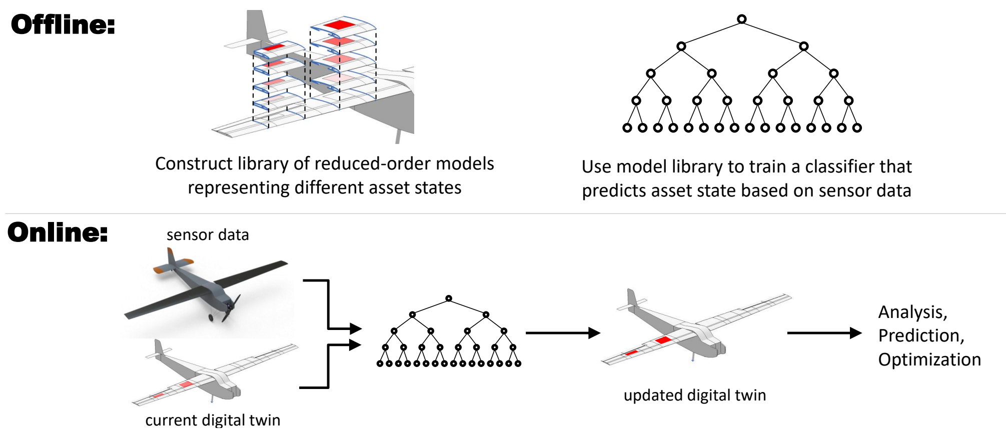

Figure 1.1: Digital twin approach: data-driven adaptation of component-based reduced-order models (Figure credit: M. Kapteyn).

shots generated from high-fidelity FEM and CFD solvers. The ROMs then provide a rapid simulation capability which is used to sample many different scenarios (different vehicle health conditions, different operating conditions, etc.). The ROM predictions from these scenarios form a large set of training data, which can be used to train a machine learning classifier. In particular, optimal classification trees are used because they are scalable and lead to highly interpretable machine learning outputs.

This model library can then be used during an online phase to rapidly create, adapt, and evaluate reduced-order models in support of analysis, prediction, and optimization. The online phase takes place when the vehicle is operating. Data from the onboard sensors are fed into the classifier. The classifier identifies which ROMs in the model library best explain the sensed data. These ROMs are then selected to comprise an updated digital twin, and used to issue rapid predictions in support of a decision. In [20], it was shown virtually via simulation that this approach could be used to diagnose a compromised portion of the wing structure and adapt the vehicle's trajectory according to a dynamically updated flight envelope.

1.4 Scientific Machine Learning via Operator Inference

Reduced-order modeling is a crucial element in the creation of a predictive digital twin. It provides the rapid simulation capability that allows the generation of sufficient data to train a machine learning classifier. Using high-fidelity tools to sample enough scenarios would be computationally intractable. Additionally, the ROMs are used to issue rapid predictions in support of the online dynamic assessment of vehicle operation. The challenge here is that there is a limited time (a fraction of a second) to issue the prediction.

Academically, reduced-order modeling has made considerable advances in the past two decades, with many successes shown over a broad range of problems in fluid dynamics, structures, acoustics, and thermal modeling. A large class of ROM methods are projection based - in the training phase a low-dimensional subspace is identified, and the ROM is derived by projecting the high-fidelity model operators onto a low-dimensional subspace. In doing so, the physics of the problem is embedded in the reduced-order representation. The drawback is that most projection-based ROM methods are intrusive, meaning the projection of the high-fidelity operators onto the low-dimensional subspace requires intrusive access and modification of the high-fidelity source code. For this reason, non-intrusive model reduction methods have emerged as alternatives that learn a model based on training data, without requiring explicit access to the high-fidelity model operators. Machine learning methods are non-intrusive and black-box, meaning that they operate only on simulation output data without knowledge of the internals of the underlying high-fidelity model. These approaches are limited in applicability to complex problems because generating enough training data is computationally expensive and the complexity of problems with multi-physics interactions and high-dimensional parameter spaces can be difficult to handle with generic approximations.

Recent work has developed a scientific machine learning approach called operator inference that targets the combination of the convenience and flexibility of non-intrusive machine learning together with the physics-based structure exploitation of reduced-order modeling [17, 18, 19, 21]. This creates the capability to learn predictive reduced-order models that provide approximate predictions of complex physical phenomena while exhibiting several orders of magnitude computational speedup over high-fidelity simulations. This work formulates the reduced-order modeling task through the lens of projection-based model reduction, but with a non-intrusive operator inference that learns the ROMs directly from simulation output data without needing access to source code [17]. By maintaining the connection to projection-based model reduction, some of the mathematical properties that give the ROMs their predictive capabilities are retained. This approach has been demonstrated for simple academic examples such as linear elasticity [17] and the compressible Euler equations [21], as well as a complex rocket combustion application [18, 19] that follows the implementation of the General Equation and Mesh Solver (GEMS) CFD code [22]. The steps of the approach are: 1) Run simulations of the high-fidelity code over representative scenarios to generate solutions (each solution is called a "snapshot" and represents the predicted physical quantities over the vehicle); 2) Identify low-dimensional structure in the snapshots, using the proper orthogonal decomposition (POD) or other dimensionality reduction technique, in combination with the physical insight drawn from the governing equations; 3) Project the snapshots onto the identified low-dimensional subspace, to obtain reduced snapshots; 4) Solve a linear least-squares regression problem to identify the low-dimensional operators that define the ROM that best match the reduced snapshots. All of these steps can be done non-intrusively—that is, without requiring access to any of the high-fidelity model source code.

1.5 Hardware Testbed

This section describes the aircraft testbed system created to begin working through the integration challenges presented by the self-aware vehicle technology. The testbed aircraft consists of a Telemaster aircraft kit (fuselage, landing gear, empennage), but outfitted with custom-designed and manufactured carbon fiber wings with custom sensors and avionics (Pixhawk autopilot, custom-built sensor boards, off-the-shelf power hardware). The fuselage-wing joint consists of a metal tube fitting so that different wings can be swapped onto the aircraft with minimal effort and in a rapid succession, such that multiple wings can be tested within a single flight test session. It was envisioned that this capability would enable several wings – from a pristine 'baseline' wing configuration through progressively more damaged wings – to be tested over the course of a single flight test, so that the same flight conditions and platform could be used to both collect sensor data and to test the dynamic data-driven algorithms in real flight conditions.

The driving design requirement of the testbed aircraft was that the wing provide a structurally similar response to a larger, more advanced Low-Cost Attritable Aircraft Technology (LCAAT) wing structure. Thus, even at the smaller scale, the preferred structural design of the wing used techniques similar to that of the larger LCAAT wing (albeit with reduced flight performance). The design criteria for the wings included: hollow carbon-fiber construction, 12-foot wingspan, plywood wing tip and root ribs, and inclusion of ailerons and flaps. The wing uses a constant 9% thick airfoil section representative of typical profiles at the mid-subsonic speed range (NACA 2309). The wing structure is split into 4 bays where the boundaries of each are designated by the plywood ribs. A split flaperon (carbon laid up over a foam core) is located in the outer two bays. The wing spar was sized for a maximum tip deflection in a 4G pull-up maneuver and the wing skins were sized for buckling in a 4G

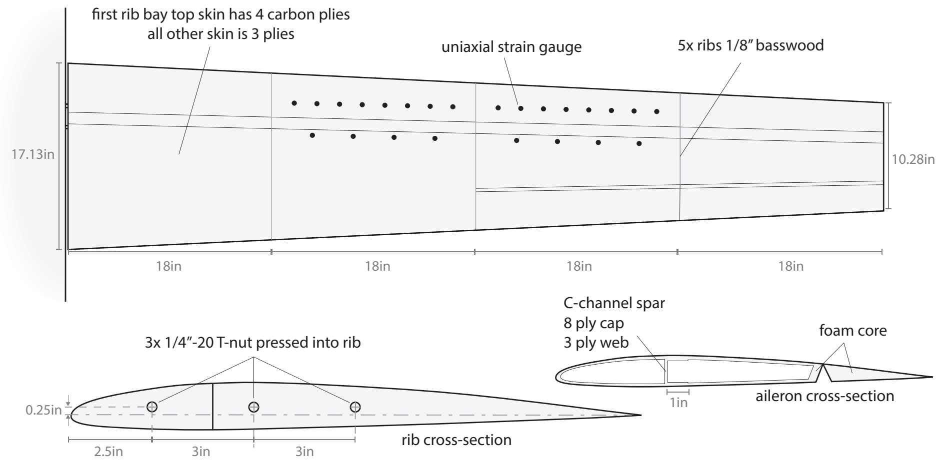

Figure 1.2: Schematic of the designed and manufactured wing structure.

pull-up maneuver. Finally, the wing was designed to have access panels on the bottom skin so that any sensors, wiring, or other hardware could be placed or modified after the wing has been constructed and assembled. Figure 1.2 contains a schematic of the wing structure.

Figure 1.3 shows the final manufactured aircraft during a series of flight tests. The first flight test conducted for the testbed aircraft consisted of a full

Figure 1.3: The custom-built self-aware UAV hardware testbed.

system assembly and checkout at the field, as well as a maiden flight. The maiden flight of the testbed aircraft was primarily used to verify that all flight hardware functioned properly in-flight and to verify that the aircraft is fully controllable and flies as anticipated. Thus, a relatively benign flight path was flown – takeoff and climb to 300 ft above ground level, fly general racetrack patterns with banks limited to 30 degrees for up to 6 minutes of flight time, and a typical descent and landing.

Figure 1.4 shows the flight path of the aircraft during its maiden flight. For this maiden flight the testbed aircraft was outfitted with a preliminary

Figure 1.4: Flight path recorded during the maiden test flight.

sensor suite consisting of twenty-four uniaxial strain gauges mounted on the top surface of the right wing, as shown in Figure 1.2. However, analysis of this data showed that the measured strain had a low signal-to-noise ratio, and thus did not correlate well with z-acceleration (i.e., aircraft maneuver) data as one would expect. Finite element simulations of the wing during the design phase showed that differences in wing deflection caused by damage would be much smaller than any differences due to varying aircraft maneuvers. Thus, this result suggests that a more advanced sensing architecture is necessary in order to detect damage in-flight.

1.6 Research Objectives

The objectives of this research aim to advance the framework for a predictive digital twin for UAVs using physics-based models and scientific machine learning, and leverage a custom-built hardware testbed to support the development of a digital-twin-enabled self-aware UAV.

Currently, the development of the digital twin has been focused on purely structural models, but the concepts underlying this structural digital twin can extend to aerodynamics as well. Information about the aerodynamic state of the vehicle would improve the predictive capabilities of the digital twin by enabling an accurate, physics-based prediction of the in-flight structural loads. Therefore, on the modeling side of this work, the research objectives are to: 1) obtain high-fidelity aerodynamic data for an appropriate test case; 2) establish a robust non-intrusive reduced-order modeling methodology and demonstrate its utility in creating ROMs for the high-fidelity models without access to source code; and 3) validate, evaluate, and demonstrate the efficient simulation capabilities of the ROMs.

To date, the digital twin models have shown promising results in simulation. However, experimental investigations into the type of sensor-driven damage detection and characterization required to achieve this functionality have shown limited success. In particular, it has been shown that damage detection and characterization place high demand on sensing capability and robustness as well as computational efficiency of the data assimilation process, even for limited damage cases on simplified wing structures [23, 24, 25]. Thus, the need remains to validate the proposed digital-twin-enabled self-aware UAV concept experimentally. To this end, this work also presents an experimental methodology for data collection and demonstration of the digital twin concept. Through experimentation with the hardware testbed, the research objectives in this area are to: 1) develop a data acquisition architecture that produces high-quality data capable of enabling the self-aware capability; 2) identify challenges and limitations that might hinder the success of these computational methods when applied to experimental data; 3) develop strategies for adapting and integrating the various computational methods to overcome these challenges; and 4) successfully implement and validate these approaches on the testbed in order to demonstrate the effectiveness of the end-to-end dynamic data-driven application system.

Chapter 2

Reduced Order Models via Operator Inference

This chapter describes the construction of reduced order models for unsteady flow over an airfoil using the operator inference approach applied to NASA's FUN3D CFD solver [26]. First, the airfoil test case problem being investigated is described in detail. Then, the necessary considerations for creating the reduced order model are discussed. Finally, the results comparing the reduced order model prediction to the full order model data are presented.

2.1 High Angle of Attack Unsteady Flow Case

This benchmark case [27] simulates flow over a 2D NACA 0012 airfoil at a 45° angle of attack, as shown in Figure 2.1. The free stream Reynolds number (4.8E6), Mach number (0.6), and speed of sound (340 m/s) are such that the flow separates and creates an oscillatory wake. In this example, lift oscillations occur at 450 Hz. To determine an appropriate time step, a characteristic time, $\mathrm{t}^*_{\mathrm{chr}}$, is identified using the frequency of oscillation as $\mathrm{t}^*_{\mathrm{chr}} = 1/\mathrm{f}^*_{\mathrm{chr}} = 1/450$ Hz $= 0.002222$ s. The nondimensional characteristic time in FUN3D is related to physical characteristic time by $\mathrm{t}_{\mathrm{chr}} = \mathrm{t}^*_{\mathrm{chr}}\mathrm{a}^*_{\mathrm{ref}}(\mathrm{L}_{\mathrm{ref}}/\mathrm{L}^*_{\mathrm{ref}})$, for compressible flow, where $\mathrm{a}^*_{\mathrm{ref}}$ is the reference (in this case, free stream)

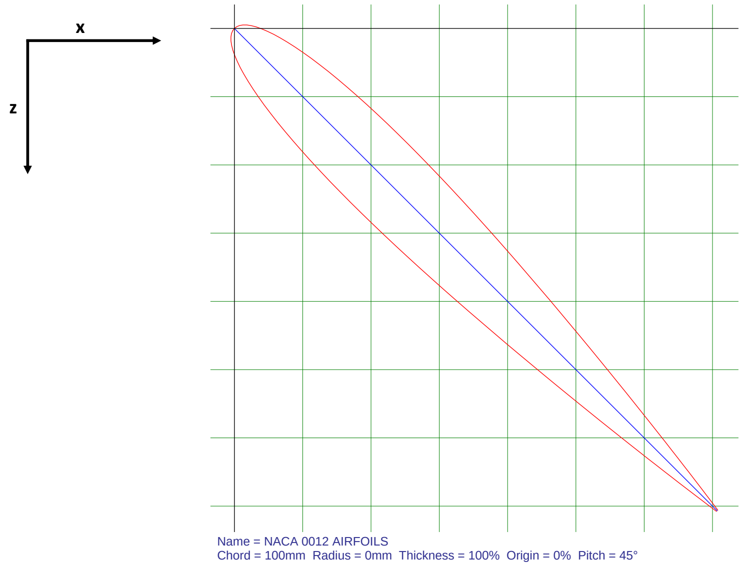

Figure 2.1: NACA 0012 airfoil at angle of attack of 45°.

speed of sound, L^*_{ref} is the reference length of the physical problem (in this case, the chord), and L_{ref} is the corresponding length in the grid (considered nondimensional). In this example, the chord is 0.1 m and the corresponding chord-in-grid is 1.0, so $L_{ref}/L^*_{ref} = 1.0/0.1 = 10\ m^{-1}$. Thus, the nondimensional characteristic time is $t_{chr} = (0.002222)(340)(10) = 7.555$. If N time steps are desired within the characteristic time, the time step size is given by $\Delta t = t_{chr}/N$, and thus a time step size of 0.07555 seconds was chosen for 100 time steps per lift cycle. The properties of this test case are summarized in Table 2.1.

Figure 2.2 shows the structured hex mesh used for CFD simulations. The mesh includes 9728 hex cells, 19824 nodes, 49192 edges, and 6 boundaries (3 Farfield Riemann, 2 Y=constant Symmetry Planes, and 1 viscous surface).

Table 2.1: Properties of Unsteady Flow Case

Airfoil	NACA 0012
Angle of Attack	45°
Reynolds Number	4.8E6
Mach Number	0.6
Speed of Sound	340 m/s
Lift Oscillation Frequency	450 Hz
Chord Length	0.1 m
Time Steps per Lift Cycle	100
Nondimensional Time Step Size	0.07555 sec

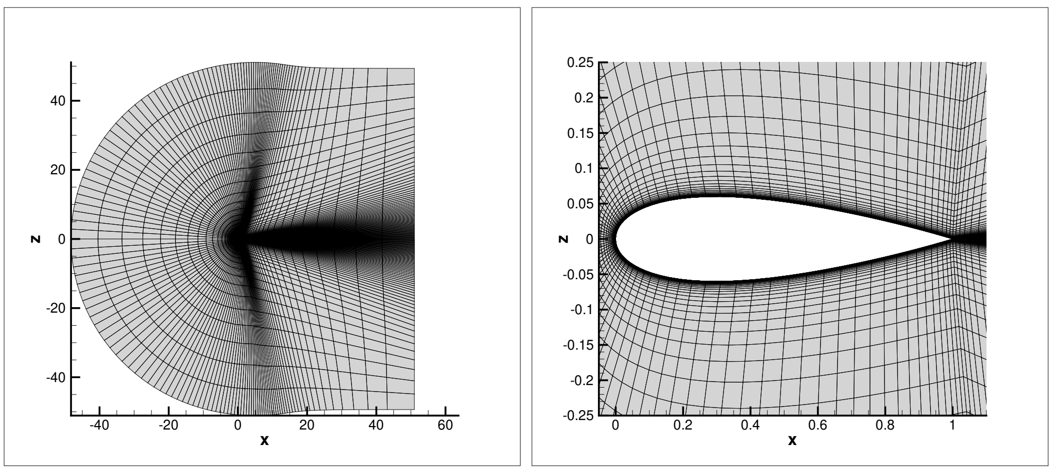

(a) (b)

Figure 2.2: CFD domain mesh showing (a) the mesh out to the farfield and (b) the mesh zoomed in on the airfoil.

This is a 2D problem, but the grid is adapted to a 3D domain with one cell spanning the y-dimension to be compatible with the solver. The CFD simulation was performed using NASA's FUN3D suite of tools using a Spalart-Allmaras turbulence model [28], and the time history data was exported to use in constructing the reduced order model.

2.2 Creation of Reduced Order Model

2.2.1 Setup

The operator inference approach to reduced order modeling targets problems governed by systems of partial differential equations. A projection-based ROM preserves the polynomial structure of a full-order model, and as such, it is desirable for the governing equations of the system being modeled to have this polynomial structure. A key feature of the operator inference approach is that it gives complete flexibility in the set of physical variables that define the ROM. Thus, the physical variables can be chosen to expose the desired polynomial structure in the governing equations, and then the snapshot data can be transformed into those variables.

The governing equations in this airfoil case are the Navier-Stokes equations for compressible flow, defined by the unsteady equations for conservation of mass, momentum, and energy. In this example, the state variables are chosen to be the specific volume variables: pressure, x-velocity, z-velocity, and specific volume. The motivation for using these variables is that the conservation equations for mass, momentum, and energy for viscous flow with no source term become quadratic in these variables when written in the specific volume form [18], and thus a ROM with a quadratic polynomial structure can be expected to yield a good approximation to the dynamics [21].

The snapshot data was collected at uniform time intervals of $\Delta t = .07555$ for $k = 100$ time steps, which results in the snapshot matrix $\mathbf{X} \in \mathbb{R}^{n \times k}$, where the jth column is the solution trajectory at time t_j and $n = 39648$ is the number of degrees of freedom in each CFD solution ($n = \frac{\text{\# of nodes}}{2} \times$ # of state variables $= \frac{19824}{2} \times 4$). Operator inference also requires $\dot{\mathbf{X}}$, the time derivative of the data, which is approximated with finite differences. Since the data comes from a quadratic polynomial model, the reduced model will have the form $\dot{\hat{\mathbf{x}}}(t) = \hat{\mathbf{c}} + \hat{\mathbf{A}}\hat{\mathbf{x}}(t) + \hat{\mathbf{H}}(\hat{\mathbf{x}}(t) \otimes \hat{\mathbf{x}}(t))$, where $\otimes$ is the column-wise Kronecker product and $\hat{\mathbf{c}}$, $\hat{\mathbf{A}}$, and $\hat{\mathbf{H}}$ are the reduced-order operators defining the ROM.

2.2.2 Determining ROM Dimension

There is a maximum ROM size that should be chosen. If the size r is too large, an "information limit" will be hit where more information (snapshots) is needed because increases in r do not add information. This maximum can be determined by examining the rank of the data matrix, $\mathbf{D} \in \mathbb{R}^{k \times d(r)}$ where $d(r) = 1 + r + \frac{r(r+1)}{2}$ [19] for a quadratic ROM of order r. The data matrix is the known data, projected to the r-dimensional subspace. In this quadratic polynomial case, the data matrix is $\mathbf{D} = [1_{\mathrm{k}} \quad \hat{\mathbf{X}}^{\mathrm{T}} \quad (\hat{\mathbf{X}} \otimes \hat{\mathbf{X}})^{\mathrm{T}}]$. $\hat{\mathbf{X}}$ is the projected data given by $\hat{\mathbf{X}} = \mathbf{V}_{\mathrm{r}}^{\mathrm{T}}\mathbf{X}$, where $\mathbf{V}_{\mathrm{r}}$ is the POD basis of rank r corresponding to the snapshot matrix $\mathbf{X}$. As r increases, the data matrix will reach a point at which it is no longer full column rank. The rank of the data matrix for increasing ROM sizes can be seen in Figure 2.3, and from this study it was determined that the maximum ROM size for this problem is $r = 12$.

To determine an appropriate dimension for the reduced model, the singular values of the snapshot matrix are computed. An indicator for the

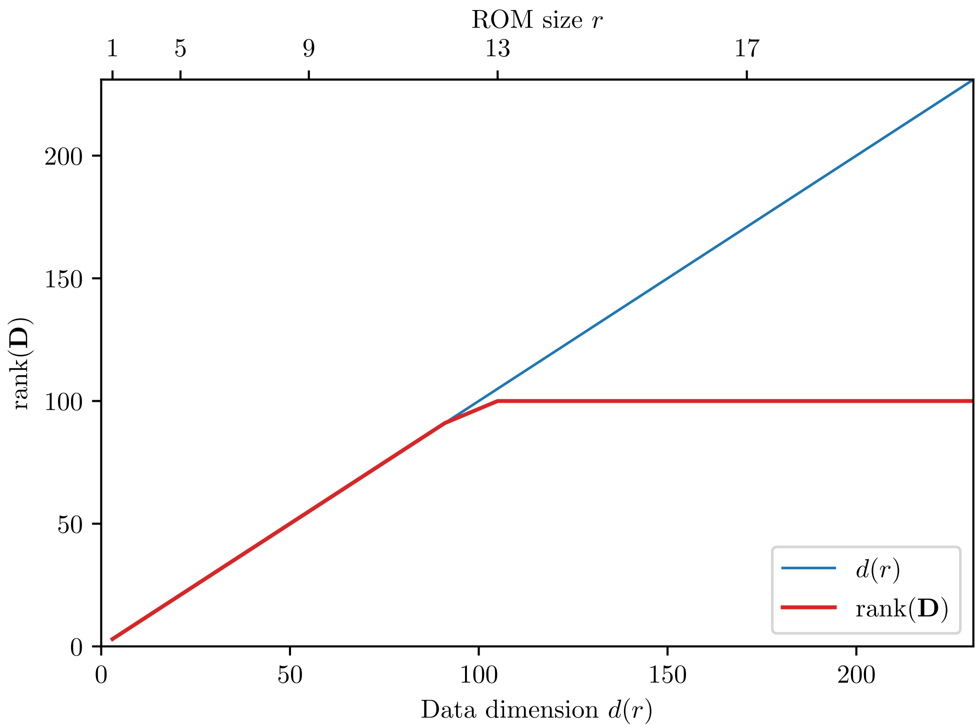

Figure 2.3: Rank of the data matrix for increasing ROM sizes.

reduced dimension is the amount of energy captured by the first j singular values, defined by $\kappa_j = \frac{\sum_{i=1}^{j} \sigma_i^2}{\sum_{i=1}^{n} \sigma_i^2}$. Figure 2.4 shows this cumulative energy for the singular value indices and computes the smallest indices j such that $\kappa_j > 99\%$, $\kappa_j > 99.9\%$, and $\kappa_j > 99.99\%$. These values are reached at $j = 4$, $j = 6$, and $j = 9$ respectively. This suggests a size $r = 6$ would capture 99.9% of the energy in the system, which is sufficient. Thus, a reduced dimension of $r = 6$ is chosen and the POD basis of rank $r = 6$ corresponding to the snapshot matrix $\mathbf{X}$ is computed as the 6 leading left singular vectors.

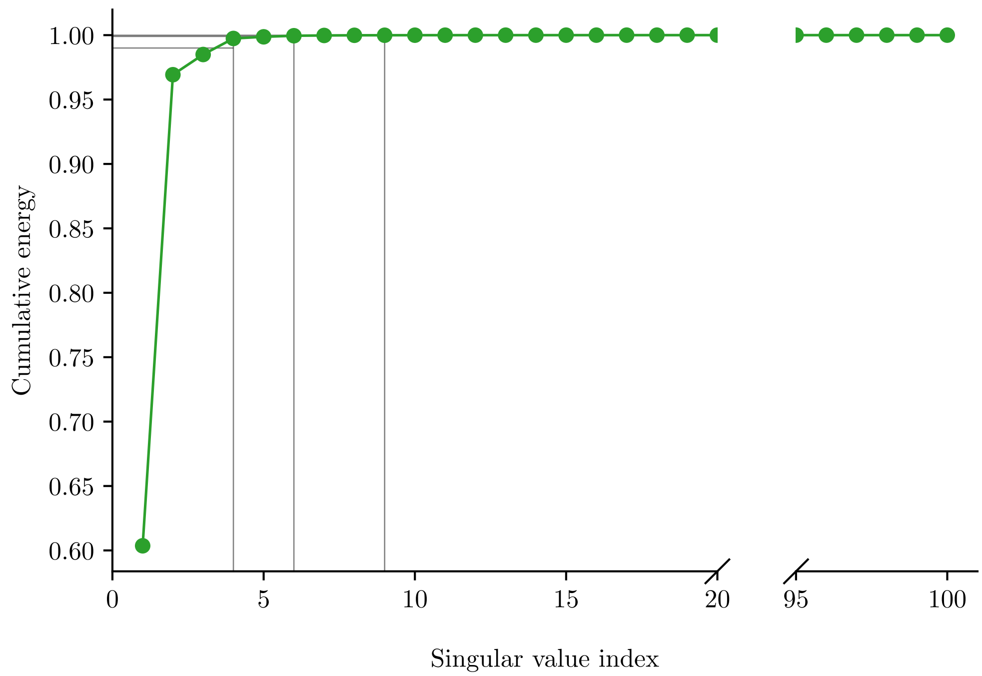

Figure 2.4: Cumulative energy of the singular values.

2.2.3 Regularization

In the operator inference approach, the operators are inferred by solving the data-driven minimization problem: $\min_{\mathbf{O}} \left\| \mathbf{D}\mathbf{O}^{\mathrm{T}} - \mathbf{R}^{\mathrm{T}} \right\|_{\mathrm{F}}^2$, where $\mathbf{D}$ is the data matrix defined in Section 2.2.2, $\mathbf{O} = [\hat{\mathbf{c}} \quad \hat{\mathbf{A}} \quad \hat{\mathbf{H}}] \in \mathbb{R}^{r \times d(r)}$ is the matrix of unknown operators, and $\mathbf{R} = [\dot{\hat{\mathbf{x}}}_0 \quad \dot{\hat{\mathbf{x}}}_1 \quad \ldots \quad \dot{\hat{\mathbf{x}}}_{k-1}] \in \mathbb{R}^{r \times k}$ is the matrix of projected time derivatives. This problem decouples into r independent linear least-squares problems, one for each of the rows of the matrix $\mathbf{O}$ of unknown operators. Each sub-problem is typically well-posed, but can also be noisy due to error in the numerical estimation of time derivatives, the truncation of the POD basis, or model misspecification. Thus, the ROMs can

Table 2.2: Optimized regularization parameters for increasing ROM sizes.

ROM size	Regularization Parameter
2	0.3711
4	1.2329
6	1.6730
8	2.7389
10	10.3598
12	12.4622

suffer from overfitting the operators to the data and exhibit poor predictive performance over the time domain. To combat this, a Tikhonov regularization [29, 30] is introduced to the sub-problems where a scalar hyperparameter λ penalizes each entry of the inferred ROM operators, driving the ROM to global stability. Thus, the minimization problem becomes: $\min_{\mathbf{o}_i} \|\mathbf{D}\mathbf{o}_i - \mathbf{r}_i\|_2^2 + \|\Gamma\mathbf{o}_i\|_2^2$, $i = 1, ..., r$, where $\mathbf{o}_i$ is row i of $\mathbf{O}$, $\mathbf{r}_i$ is row i of $\mathbf{R}$, and $\Gamma = \lambda\mathbf{I}^{d(r)\times d(r)}$. The hyperparameter λ is chosen so that the resulting ROM minimizes error over the training domain while maintaining a bound on the integrated POD coefficients. This optimization problem is given by: $\min\|\hat{\mathbf{X}} - \tilde{\mathbf{X}}\|$ subject to $\max_{i,j}|\tilde{\mathbf{X}}_{i,j}| \leq B$, $B > 0$, where $\tilde{\mathbf{X}}$ is obtained by solving the least-squares minimization problem for the unknown operators with regularizer $\lambda\mathbf{I}$ and then integrating the reduced model over the time domain, and B is the bound on the integrated POD coefficients. B is chosen as a multiple (in this case, 1.5) of the maximum absolute entry of the projected training data $\hat{\mathbf{X}}$. The results of this optimization process for increasing ROM sizes for this problem can be seen in Table 2.2. Since the ROM size chosen was $r = 6$, the regularization parameter $\lambda = 1.6730$ was used in the construction of the reduced model.

2.3 Results

The ROM was constructed and simulated over the time domain of 100 time steps and compared to the results from the full order model (FOM). Figure 2.5 shows the predicted (ROM) and actual (FOM) pressure contours at a few snapshots of the lifting period (additional snapshot pressure contours are given in Appendix A.1). Pressure error contours are provided in Figure 2.6.

For the error analysis, the ℓ^2-norm errors between the snapshot data and the predicted data are computed. This is broken down into the absolute ℓ^2 error, given by $e_{\ell^2,\text{ abs}} = ||X - X_{\text{ROM}}||_2$, and the relative ℓ^2 error, given by $e_{\ell^2,\text{ rel}} = \frac{||X - X_{\text{ROM}}||_2}{||X||_2}$, where X is the FOM state variable data and X_{ROM} is the state variable data predicted by the ROM. The absolute and relative ℓ^2 errors were computed for each of the state variables and are shown in Figure 2.7.

The predicted and actual pressure distributions around the airfoil were also compared. The FOM and ROM data for the state on the airfoil surface was extracted, and the coefficient of pressure was plotted against the nondimensional chord location. The FOM and ROM pressure distributions for a few snapshots are shown in Figure 2.8, along with the absolute error on both the upper surface and the lower surface (additional snapshot pressure distributions are given in Appendix A.2).

It is difficult to see many differences between the predicted and actual results in Figure 2.5, a testament to the performance of the ROM. The differences between the FOM and ROM data in Figure 2.8 are nearly invisible, also showing how well the ROM is able to predict over the time domain.

The pressure on the airfoil surface was then converted from the nondimensional FOM and ROM output values to the dimensionalized values in kPa and integrated over the airfoil surface to compute lift and induced drag for both the FOM and ROM. As the airfoil is at an angle of attack of $45°$, the lift and induced drag are equivalent. The comparison of the FOM and ROM lift over time is shown in Figure 2.9.

Finally, the training data of 100 snapshots was used to build a ROM, and then predict over a time domain of 300 time steps. Additional snapshot data from the FOM was collected to compare to the ROM prediction. Figure 2.10 shows the lift over time for these 300 steps (the induced drag again would be the same plot). The vertical black line indicates the boundary between the training and testing data. This shows that the ROM predicts well in the training regime, and is able to predict out past the training data sufficiently well for another lifting cycle. However, on the third lifting cycle, the error between the FOM and ROM lift starts to increase. Thus, the further from the training regime, the poorer the predictive performance of the ROM.

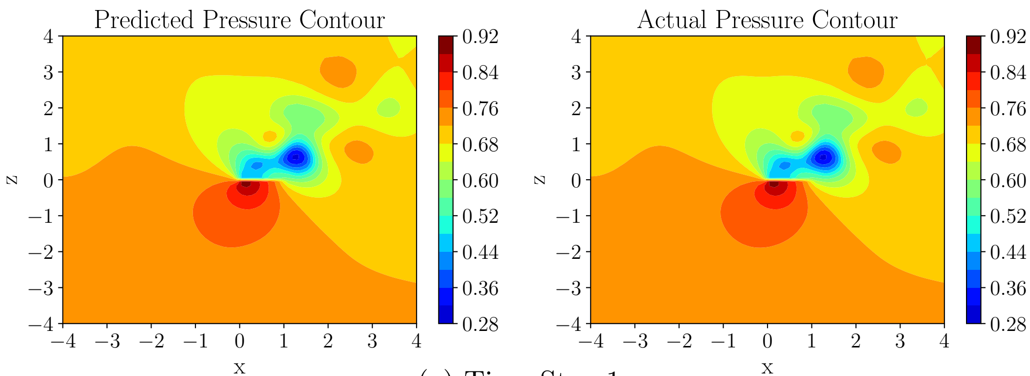

(a) Time Step 1

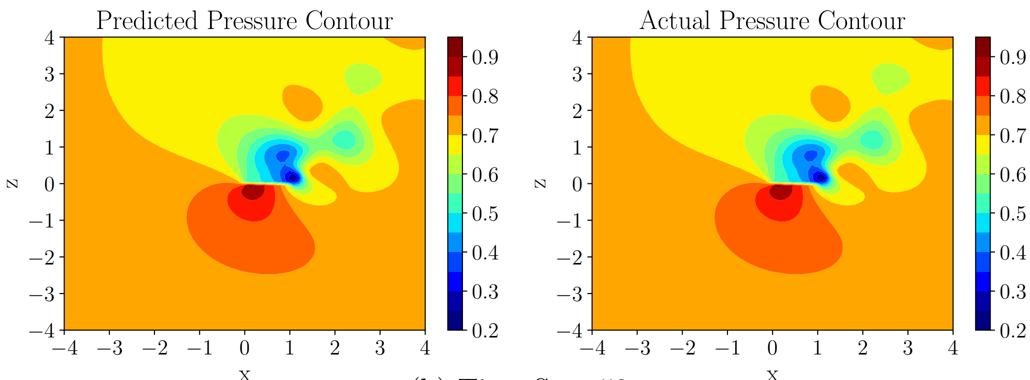

(b) Time Step 50

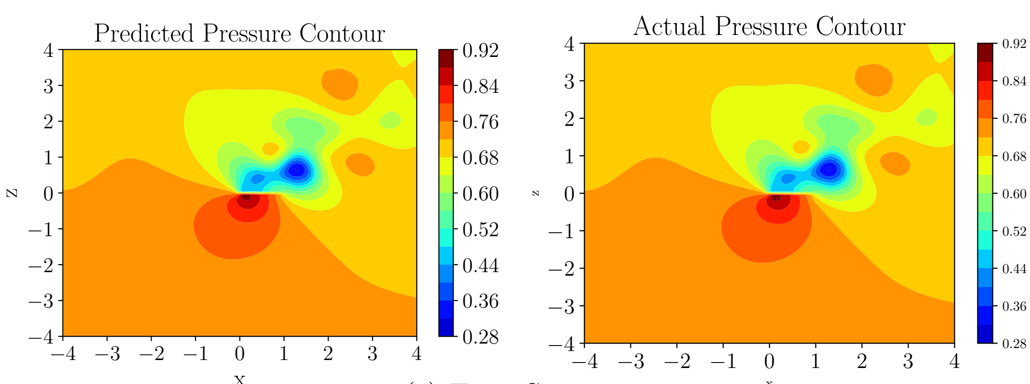

(c) Time Step 100

Figure 2.5: Predicted and actual pressure contours at snapshots.

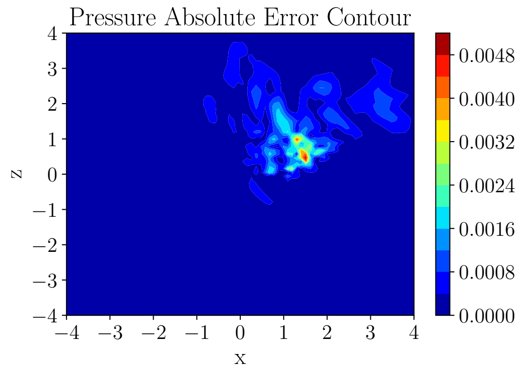

(a) Time Step 1

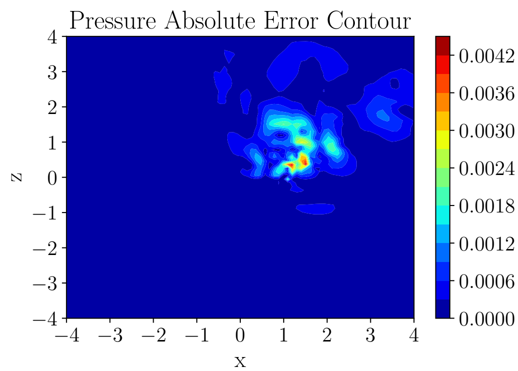

(b) Time Step 50

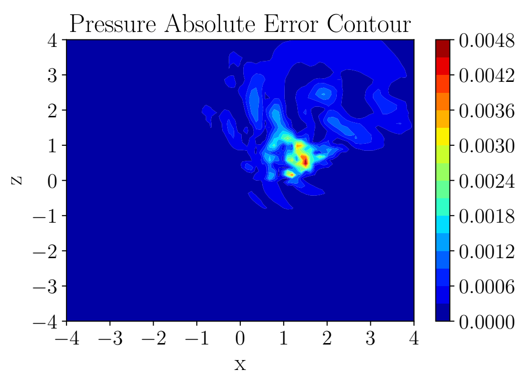

(c) Time Step 100

Figure 2.6: Pressure absolute error contours at snapshots.

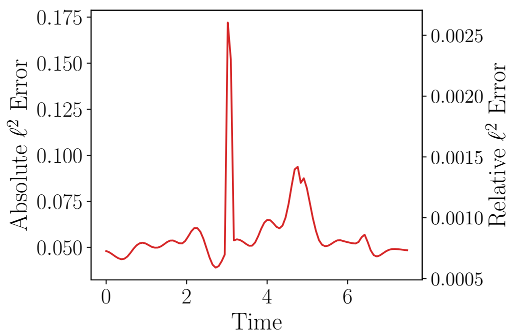

(a) Pressure

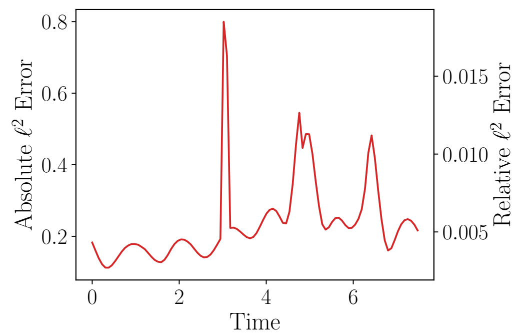

(b) x-Velocity

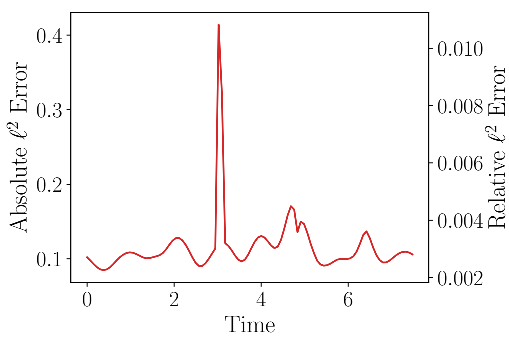

(c) z-Velocity

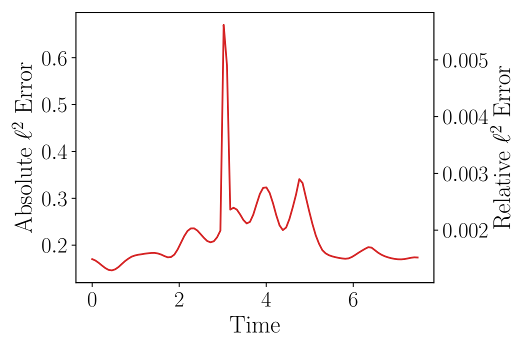

(d) Specific Volume

Figure 2.7: Absolute and relative ℓ^2 errors for the state variables.

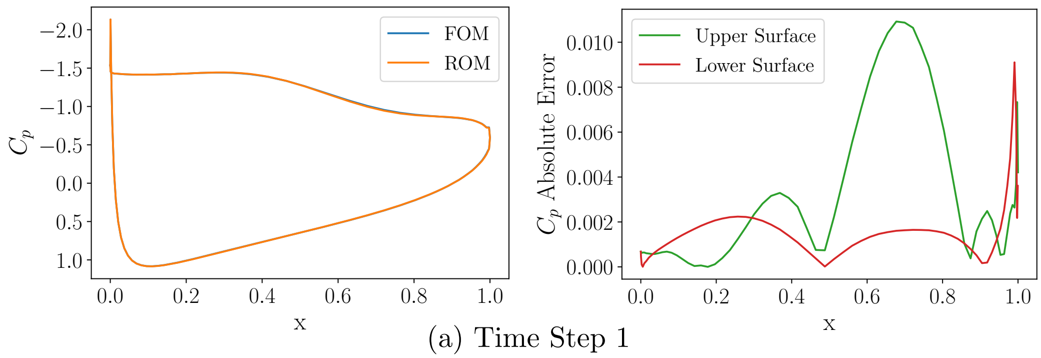

(a) Time Step 1

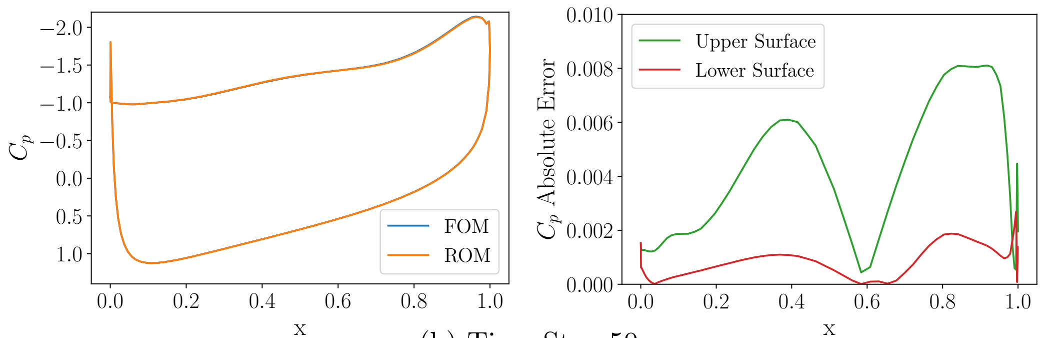

(b) Time Step 50

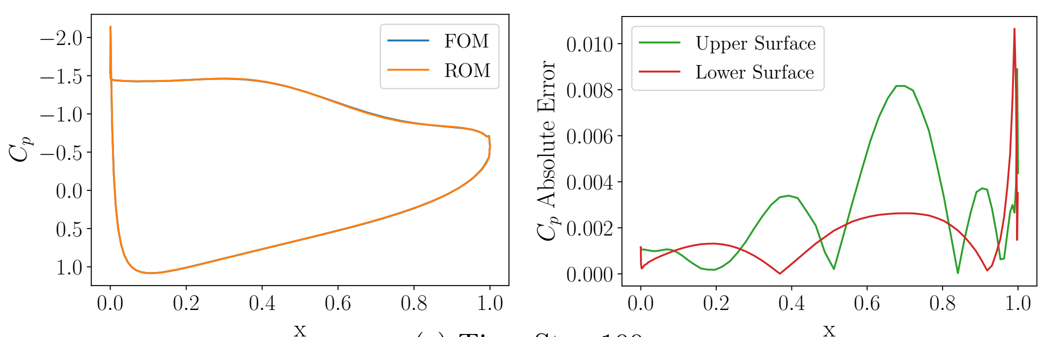

(c) Time Step 100

Figure 2.8: Predicted and actual pressure distribution with errors at snapshots.

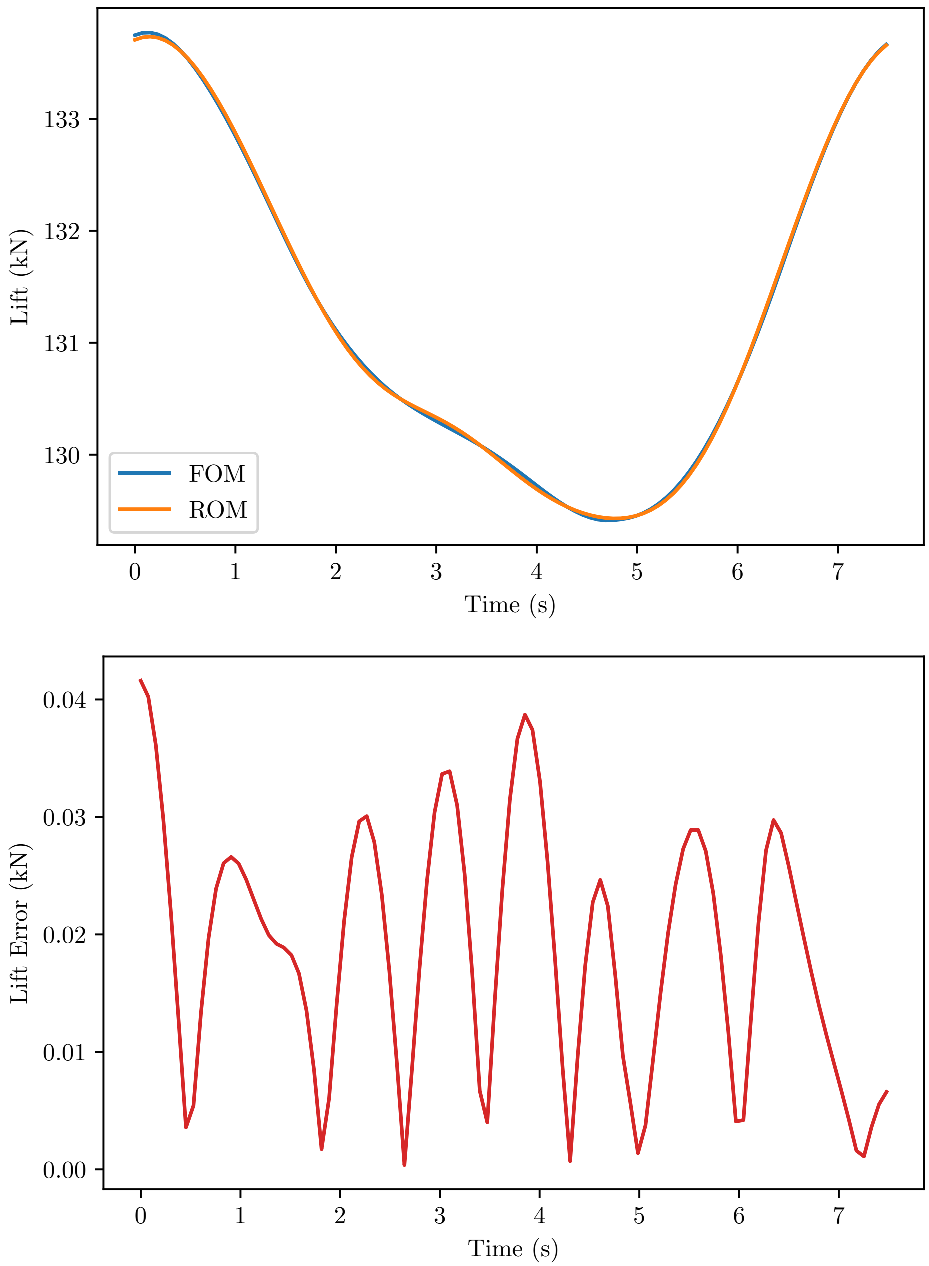

Figure 2.9: Predicted and actual lift (same as induced drag) over time.

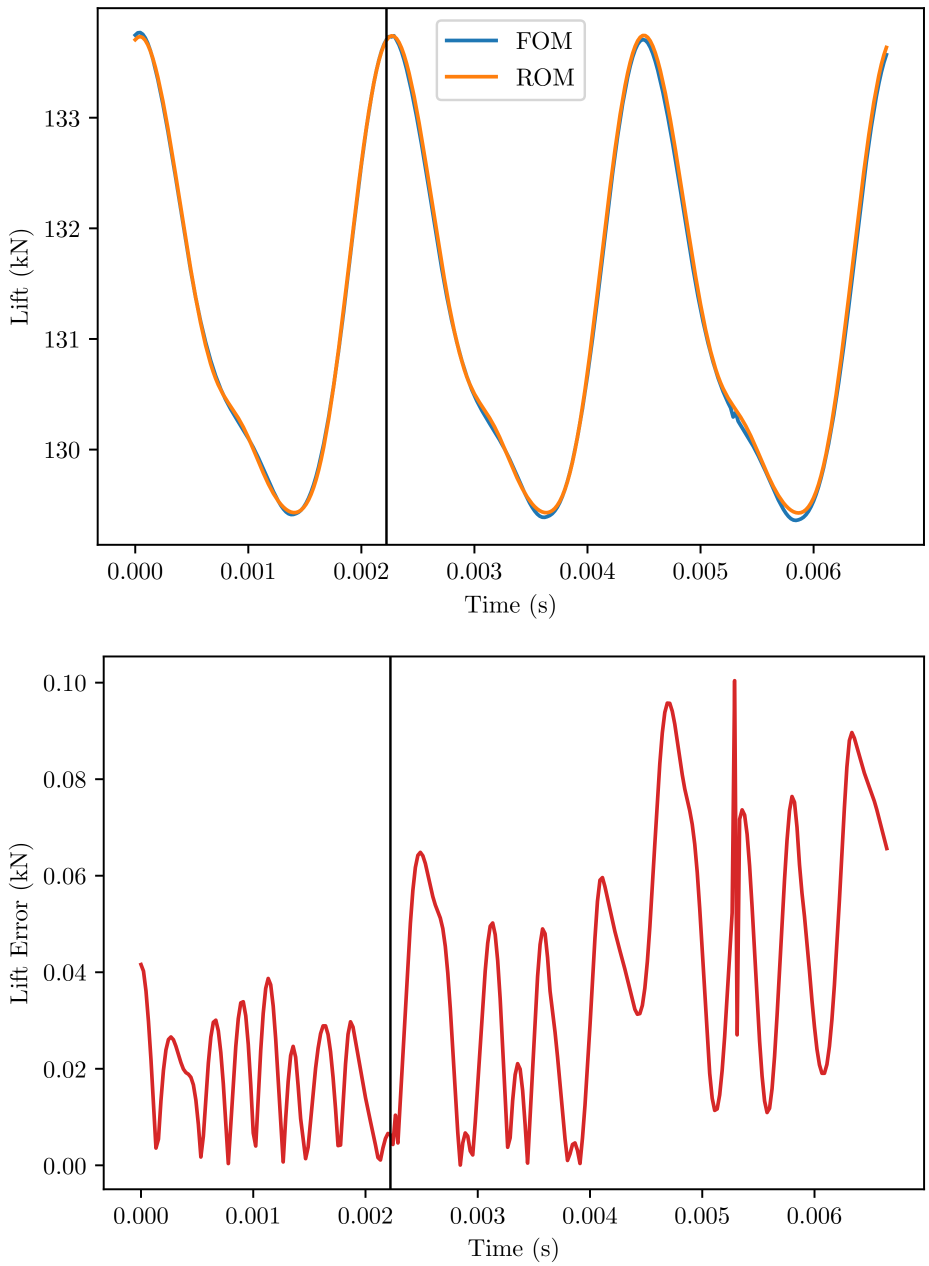

Figure 2.10: Predicted and actual lift (same as induced drag) over time; solid black vertical line indicates boundary between training and testing regimes.

Chapter 3

Experimental Data Collection and Analysis

The maiden flight of the testbed aircraft revealed that a more advanced sensing architecture is required to enable the self-aware UAV. This chapter [1] presents a recent effort to develop such an architecture. First, the bench-top experimental setup used for this work as well as the sensing technology adopted is described. Finally, results are presented that serve as a proof-of-concept for the application of this sensing architecture to the self-aware UAV.

3.1 Experimental Setup and Sensor Technology

A bench-top experimental setup was developed with the hardware testbed that enables controlled experiments and collection of realistic sensor data for the aircraft. In the bench-top setup, the wings are mounted upside-down to a wooden mount that mimics the fuselage. The opportunity also exists to mount the electric motor from the testbed onto this fuselage mount in order to excite vibrations in the wings that are characteristic of those expected in-flight. The experimental setup for these tests is shown in Figure 3.1.

[1] Work from this chapter previously published in [31]. Setup, experiments, and analysis performed by author in collaboration with M. Kapteyn and J. Pretorius.

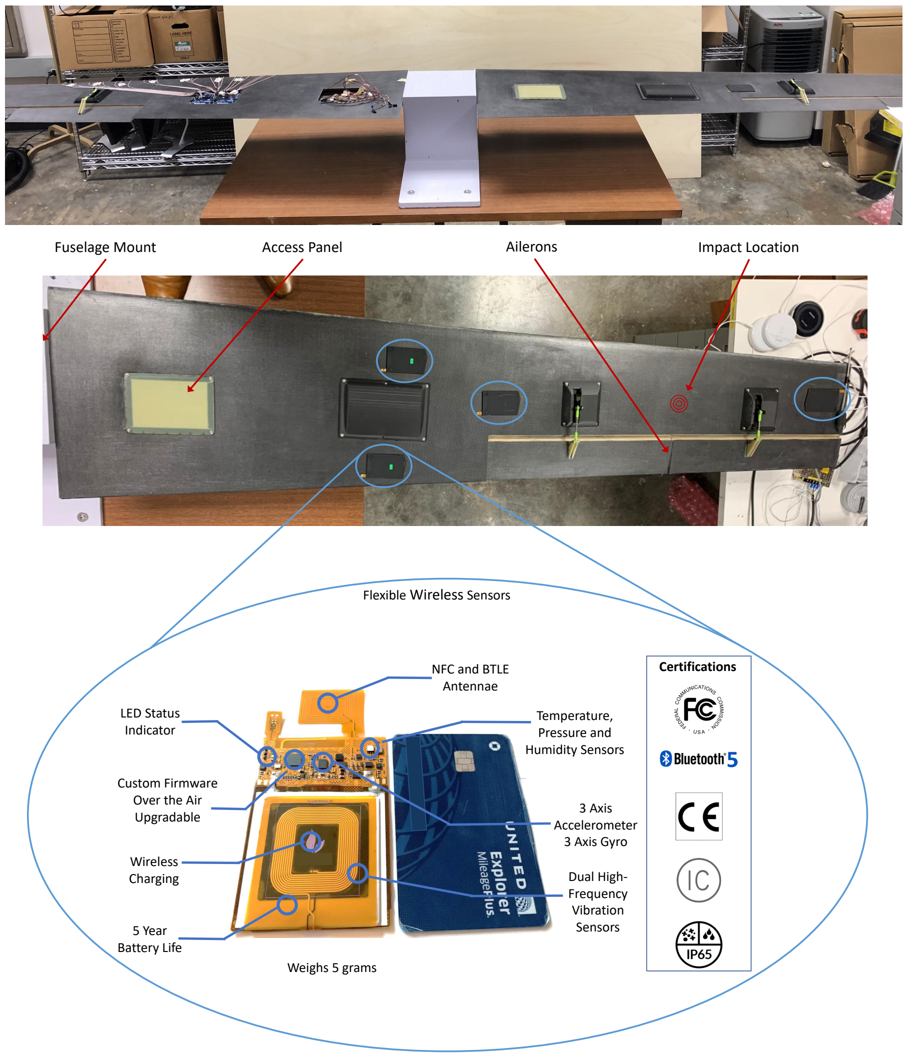

Figure 3.1: Experimental setup and wireless sensor used for data collection.

Based on the data collected during the validation test-flight, the decision was made to switch from the traditional uniaxial strain gauges mounted on the top surface of the right wing, to a set of dual high frequency dynamic strain sensors mounted on the bottom surface of the left wing. The primary motivation for this change was the improved signal-to-noise ratio. The wings on the testbed vehicle are relatively flexible, so the strains observed in-flight are typically dynamic with significant high frequency content. In this setting the dynamic strain gauges provide increased sensitivity, as well as reduced susceptibility to electromagnetic noise, and thus a higher signal-to-noise ratio.

The dynamic strain sensors used in this work are embedded in a set of wireless, self-adhering sensor suites, one of which is shown in detail in Figure 3.1. In addition to the dynamic strain sensors, each wireless sensor includes temperature, pressure, and humidity sensors, as well as a 3-axis accelerometer and gyroscope. In addition, the sensors have a built-in analog-to-digital converter, Bluetooth transmitter, and long-life battery. The wireless nature of the sensors provides additional benefits such as reduced weight, system complexity, and aerodynamic drag due to the absence of wires and other sensor hardware. Preliminary data for this work was collected using one of these wireless sensors, but the form factor and ease of installation would allow for many of these sensors to be used.

3.2 Proof-of-Concept Results

As shown in Figure 3.1, there is a removable access panel, originally intended for modifications to sensing hardware. However, this component also allows the testbed to represent a scenario in which the access panel is unintentionally left open or entirely detached. Customizing this panel also allows for emulation of different structural states. For example, a flexible

Table 3.1: Material properties of different access panel cases.

Material	Thickness	Elastic Modulus
Thick Carbon Fiber	1/16”	2400 ksi
Thin Carbon Fiber	1/32”	2400 ksi
PVC	1/16”	450 ksi
Nylon	1/16”	400 ksi

panel emulates a reduction in stiffness in the wing skin caused by damage or degradation. In preliminary data collection, both the thickness and elastic modulus of the panel were varied. In particular, the cases tested were carbon fiber panels of two different thicknesses, a PVC panel, a nylon panel, and a reference case with no access panel attached. The material properties of the varying panels are provided in Table 3.1. In the preliminary data collection, a small hammer is tapped at the impact location site (indicated in Figure 3.1) to induce high frequency vibrations in the wing, and one of the wireless sensors collects data through the vibration sensor at a sampling frequency of 5000 Hz. This hammer impact test is repeated using the different access panels described in Table 3.1.

The goal of these experiments is to process the sensor data from the hammer impact tests in order to extract features containing information about the structural response of the wing, and demonstrate how these features can be used to estimate the structural state of the wing, in this case represented by the access panel properties. Figure 3.2 (top and middle respectively), show the raw vibration sensor output for each of the panel cases and the vibration sensor output for each case after filtering with a 250 Hz high-pass filter. The high-

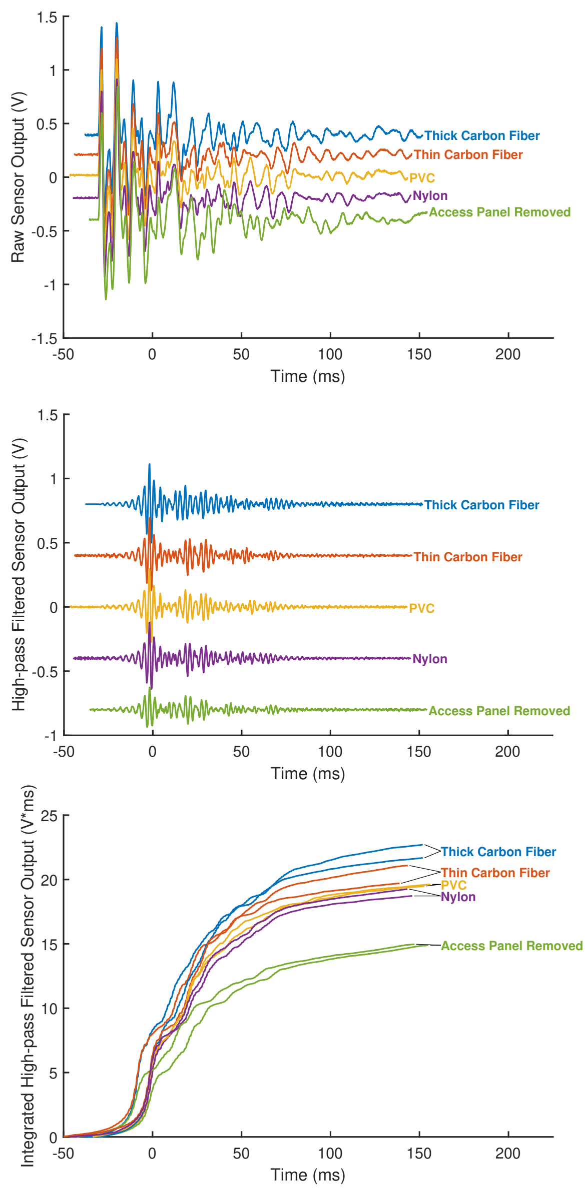

Figure 3.2: Preliminary experimental results. Top: Raw vibration sensor output for each access panel case. Middle: High-pass filtered vibration sensor output. Vertical offsets are added to better show the difference between cases. Bottom: Integrated high-pass filtered sensor data (two repeated trials for each case).

pass filter cutoff frequency of 250 Hz is chosen as a hyperparameter where meaningful discrimination between panel cases can be seen. The high-pass filtered data shows that there is a variation of high frequency content between each access panel case. This can be more clearly seen in Figure 3.2 (bottom), which shows the integrated high-pass filtered sensor output for two trials of each of the panel cases.

The integrated filter output shows a clear trend based on the access panel stiffness. As the stiffness of the panel is reduced, the integrated filter output is decreased with respect to time. The two trials show a fair degree of consistency, however in future work more data will be collected to ensure consistency and provide a more complete dataset. These preliminary experiments demonstrate that the sensor architecture is capable of detecting differences in the structural response of the wing, even when the difference in the underlying structural state is small; in this case only the properties of the (relatively small) access panel are varied.

In the future, the data collected using these bench-top experiments will be integrated into the self-aware UAV system framework, which is summarized by the information model shown in Figure 3.3. In this framework, the experimental data will be used in conjunction with structural models to train a classifier [2] in which features extracted from in-flight sensor data (in this case the amount of high-frequency content) can be used to estimate the structural state of the wing (in this case which access panel is attached to the wing). Online, this classifier can be used as part of a digital twin that enables

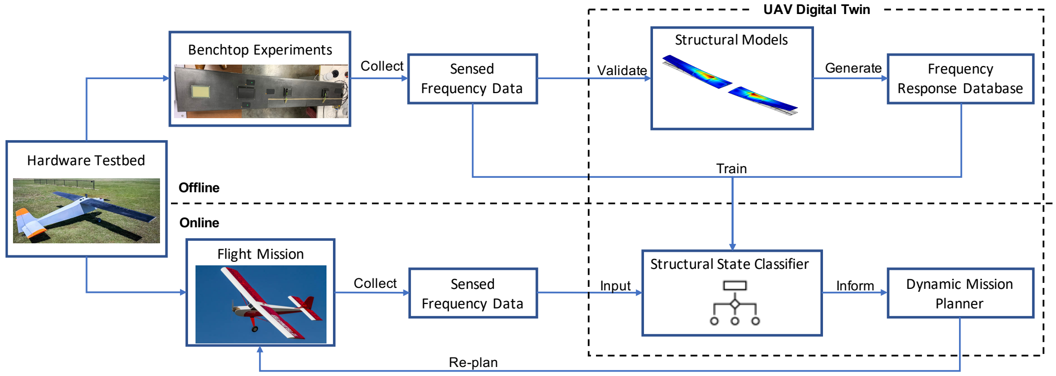

Figure 3.3: Information model for the proposed architecture.

condition-aware sensing and dynamic mission replanning, thus enabling the UAV to become self-aware.

Chapter 4

Conclusions

4.1 Research Summary

This research work advanced the framework for a predictive digital twin for unmanned aerial vehicles. First, non-intrusive ROM methods that work efficiently with physics simulations where internal access to the simulation code is unavailable were developed in application to unsteady flow over an airfoil. The maximum ROM size and the need for regularization in the operator inference approach were investigated to achieve a ROM with satisfactory predictive performance. The results comparing the reduced-order model prediction to the full-order model data show that the ROM was capable of reconstructing the time domain of the training data with very little error, and was able to accurately predict out two additional lifting cycles. Thus, these aerodynamic ROMs have the ability to improve the predictive capabilities of the digital twin by enabling a rapid, accurate, physics-based prediction of the in-flight aerodynamic loads.

In addition, this work presented an experimental data collection and analysis method to support the development of a self-aware UAV. A sensor architecture was developed that leverages wireless self-adhering dynamic strain sensors capable of measuring high frequency vibrations in the structural response of the wing. Bench-top experiments were conducted and the resulting data suggests that the degree of high-frequency content provides a useful feature for classifying the structural state of the wing. The development of a fully functional aircraft system capable of generating high quality experimental data serves as a key enabler towards validating and verifying a fully self-aware aircraft system.

4.2 Future Research Directions

The reduced-order modeling work validated the operator inference approach as applied to aerodynamic ROMs. The next direction would be to apply the developed framework to the larger scale problem of a 3D wing. Once the ROM is validated against a single flight condition, the goal is to extend the methods to accept parametric changes in the flight conditions or wing orientation. The goal for the final aerodynamic ROM is to take as input time-dependent measurements of the wing displacements and the flight conditions, and then compute time-dependent aerodynamic loading as output. Finally, the coupling between aerodynamic ROMs and structural ROMs can be investigated, seeking to create fluid-structure interaction ROMs, which would establish a robust digital twin in which multidisciplinary modeling enables advanced, intelligent decision making.

Future work on the experimental side will focus on gathering additional data sets, and further analysis of the data in order to extract additional features that can reliably inform classification of the wing structural state. Other tests include ring-down response experiments to capture natural frequency information, forced vibration tests to examine the structural response of the wing to vibrations induced by the aircraft's electric motor, and load-displacement tests to examine stiffness characteristics. While the preliminary experiments utilized only one sensor, multiple sensors can be used, and features extracted from different sensors can be added to the classifier in order to more reliably estimate the severity of structural defects, and also estimate the location of these defects.

These future research directions combine modeling and simulations with hardware experimentation that work to enable the successful creation of the robust, digital-twin-enabled self-aware UAV.

Appendices

Appendix A

Additional Snapshot Results

A.1 Pressure Contours

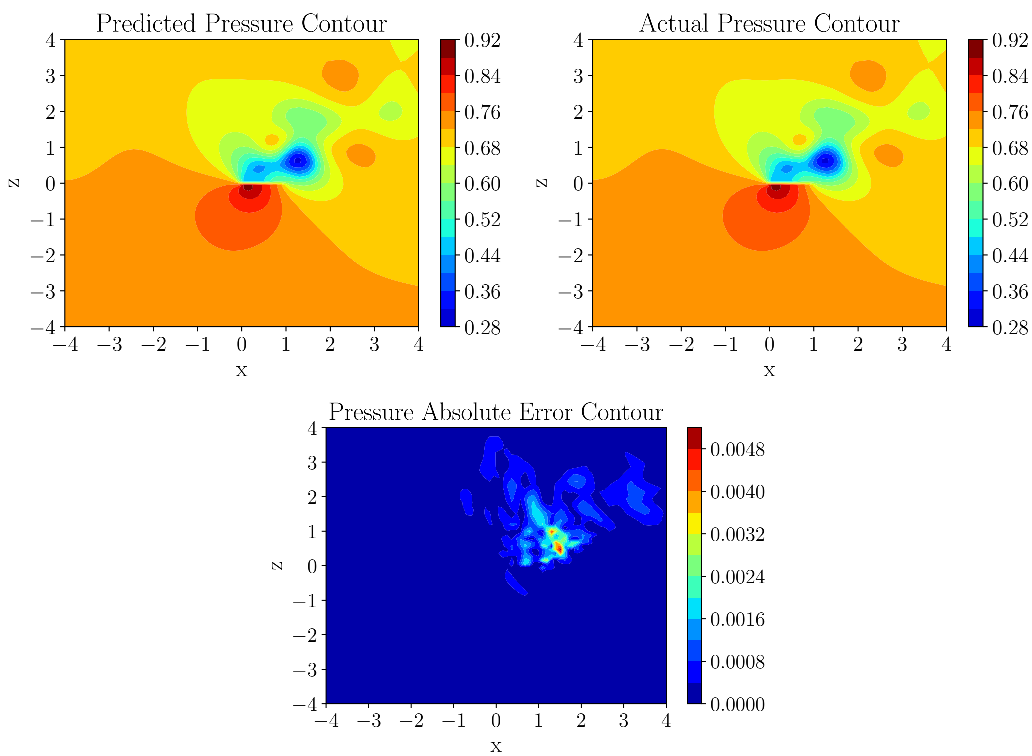

(a) Time Step 1

Figure A.1: Continued below.

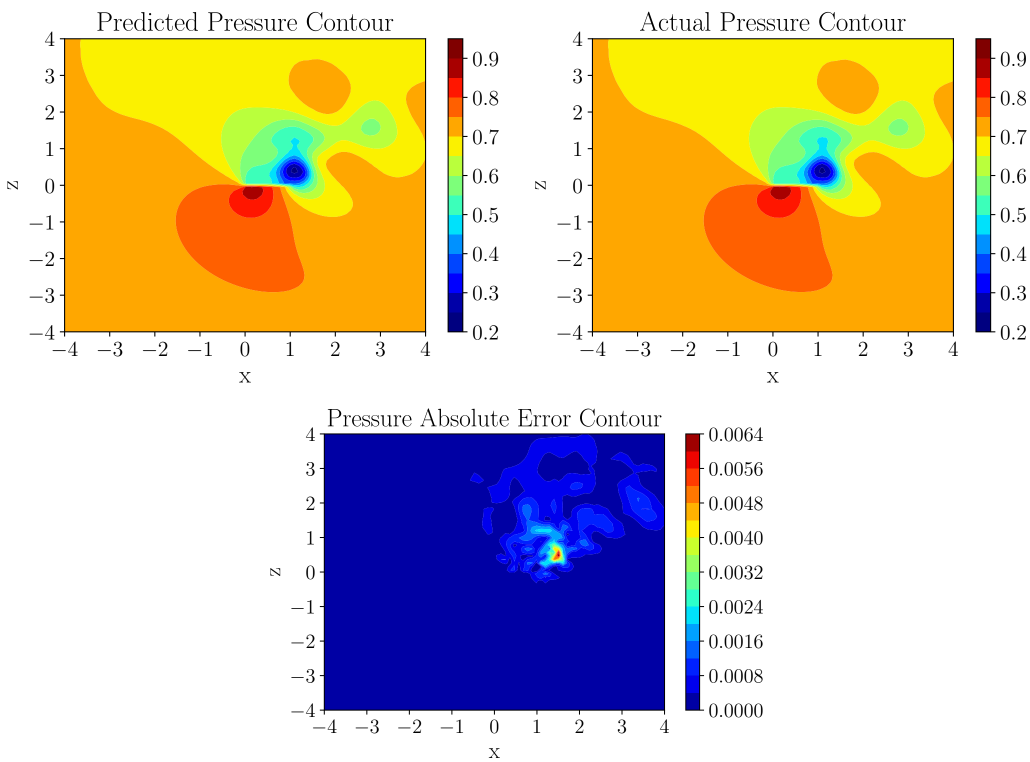

(b) Time Step 25

Figure A.1: Continued below.

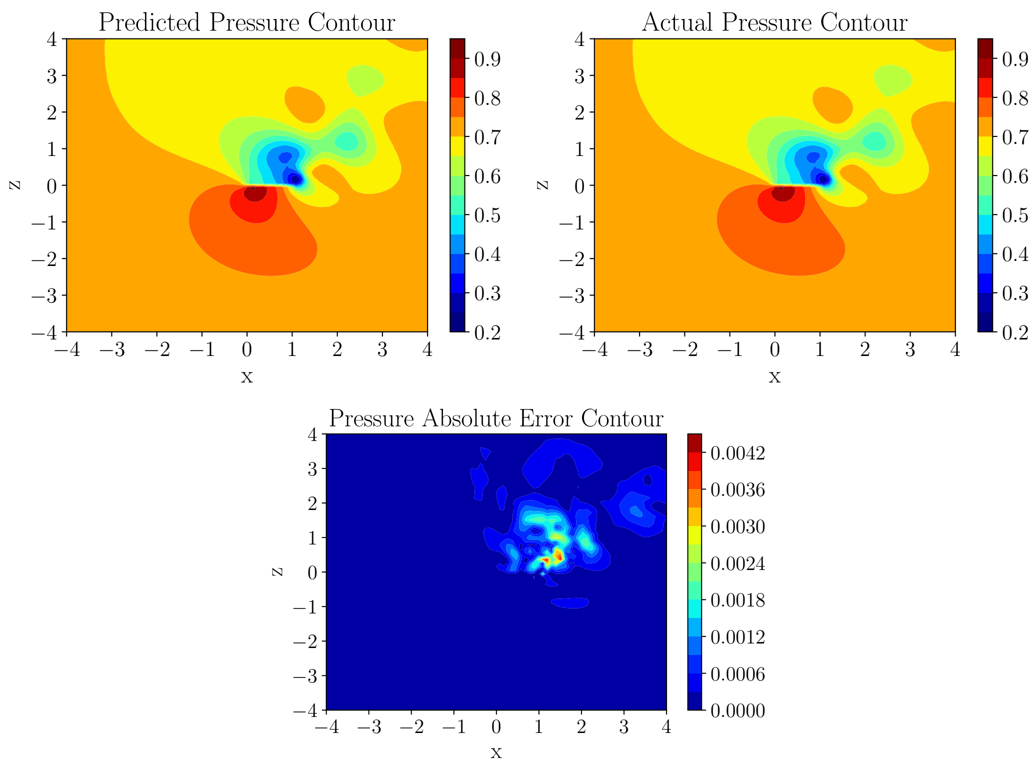

(c) Time Step 50

Figure A.1: Continued below.

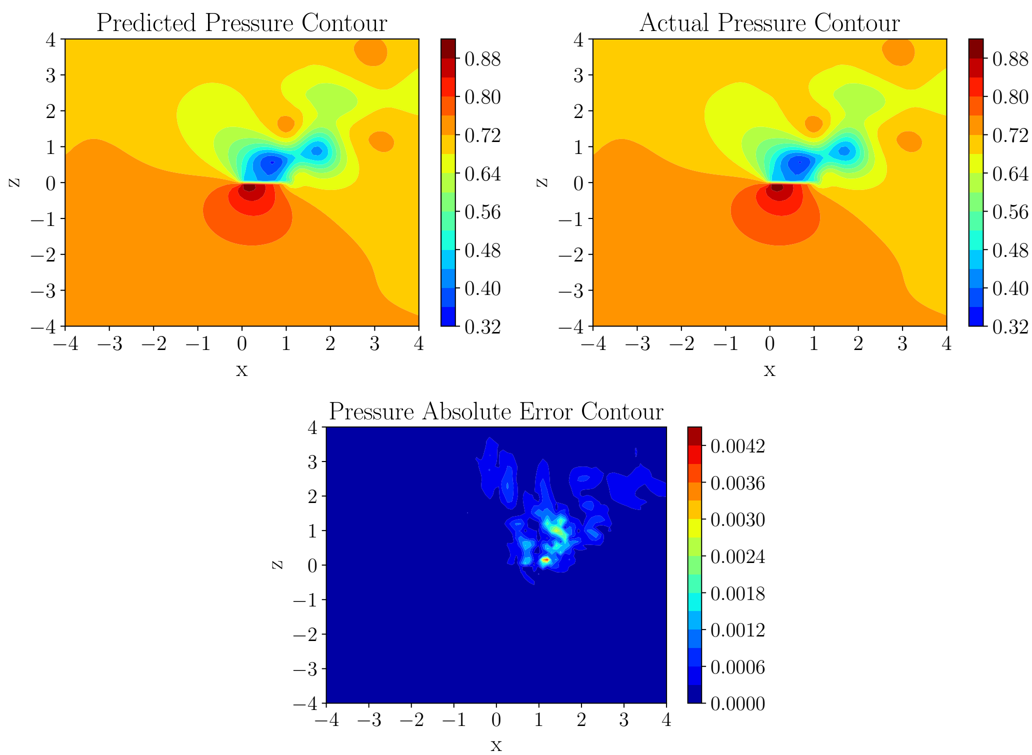

(d) Time Step 75

Figure A.1: Continued below.

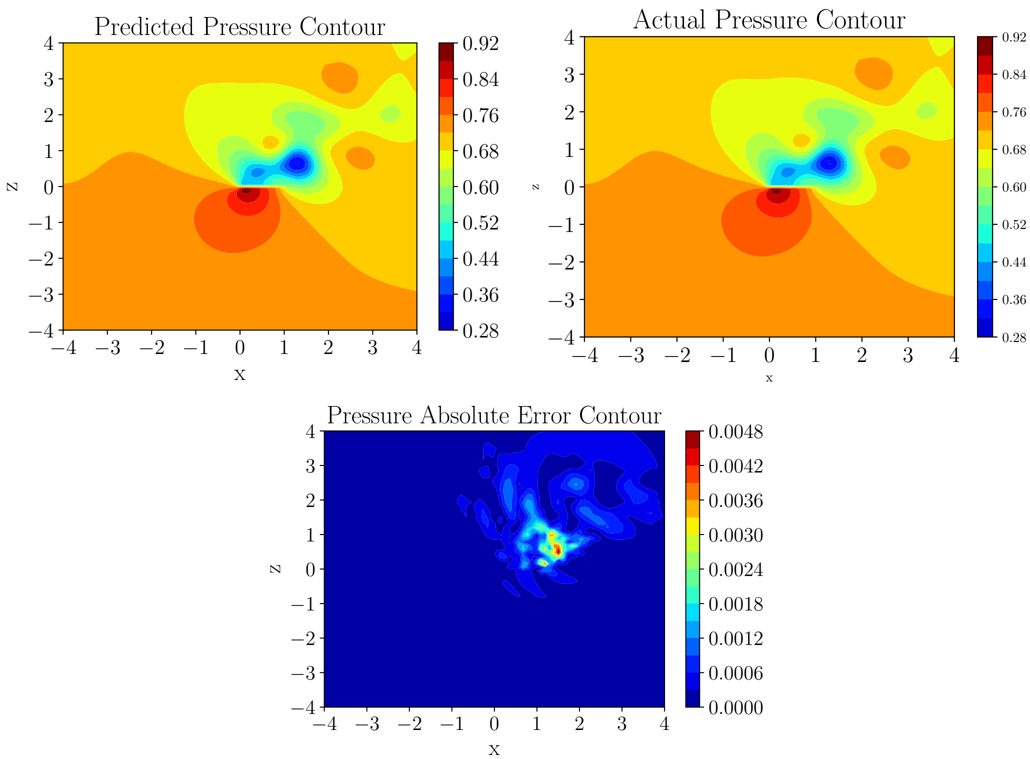

(e) Time Step 100

Figure A.1: Predicted and actual pressure contours at additional snapshots

A.2 Pressure Distributions

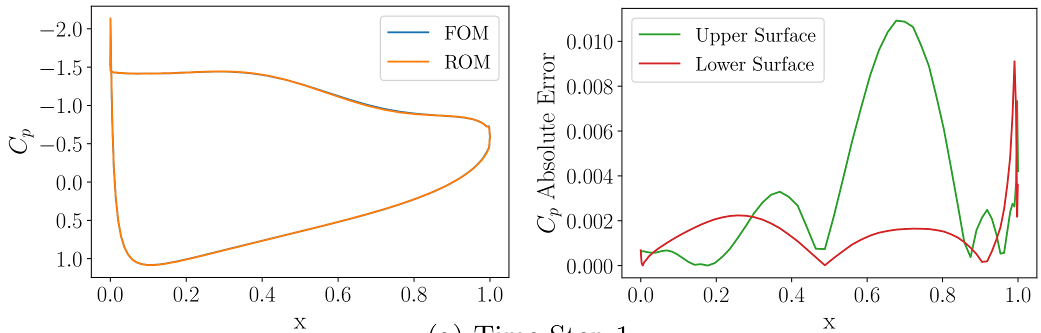

(a) Time Step 1

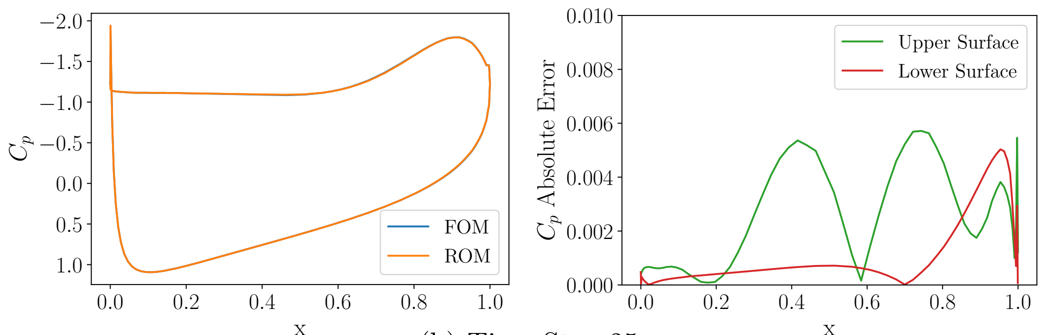

(b) Time Step 25

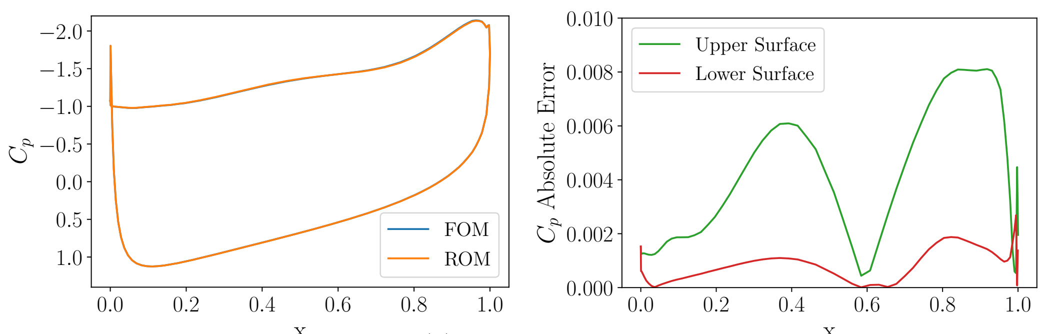

(c) Time Step 50

Figure A.2: Continued below.

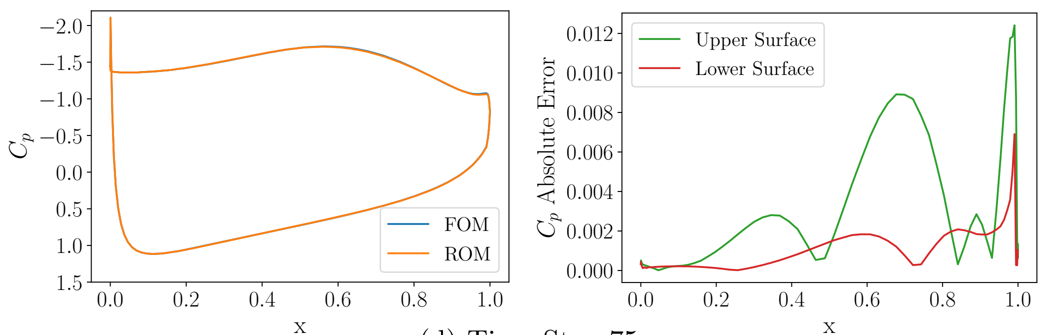

(d) Time Step 75

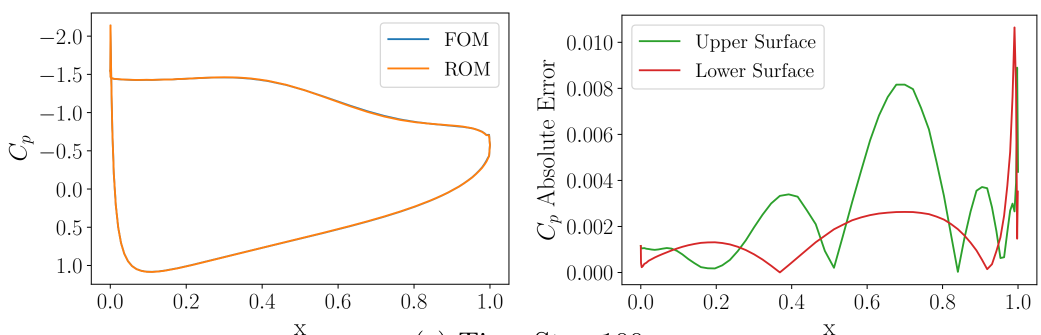

(e) Time Step 100

Figure A.2: Predicted and actual pressure distribution with errors at additional snapshots.

References

[1] AIAA Digital Engineering Integration Committee, *Digital Twin: Definition and Value*, American Institute of Aeronautics and Astronautics (AIAA) and Aerospace Industries Association (AIA) Position Paper, 2020.

[2] Kapteyn, M. G., Knezevic, D. J., Huynh, D., Tran, M., and Willcox, K., "Data-driven physics-based digital twins via a library of component-based reduced-order models," *International Journal for Numerical Methods in Engineering*, 2020.

[3] Tuegel, E. J., Ingraffea, A. R., Eason, T. G., and Spottswood, S. M., "Reengineering Aircraft Structural Life Prediction using a Digital Twin," *International Journal of Aerospace Engineering*, Vol. 2011, 2011, Article ID 154798.

[4] Glaessgen, E. and Stargel, D., "The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles," *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, AIAA 2012–1818, 2012.

[5] Li, C., Mahadevan, S., Ling, Y., Choze, S., and Wang, L., "Dynamic Bayesian Network for Aircraft Wing Health Monitoring Digital Twin," *AIAA Journal*, Vol. 55, No. 3, 2017, pp. 930–941.

[6] Reifsnider, K. and Majumdar, P., "Multiphysics stimulated simulation digital twin methods for fleet management," *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA 2013–1578, 2013.

[7] Bruynseels, K., Santoni de Sio, F., and van den Hoven, J., "Digital Twins in Health Care: Ethical Implications of an Emerging Engineering Paradigm," *Frontiers in Genetics*, Vol. 9, 2018, pp. 31.

[8] Tao, F., Zhang, M., and Nee, A., *Digital twin driven smart manufacturing*, Academic Press, 2019.

[9] Mohammadi, N. and Taylor, J. E., "Smart city digital twins," *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2017, pp. 1–5.

[10] Yu, H., Miao, C., Leung, C., and White, T. J., "Towards AI-powered personalization in MOOC learning," *npj Science of Learning*, Vol. 2, No. 1, 2017, pp. 1–5.

[11] Allaire, D., Biros, G., Chambers, J., Ghattas, O., Kordonowy, D., and Willcox, K., "Dynamic Data Driven Methods for Self-aware Aerospace Vehicles," *Procedia Computer Science*, Vol. 9, 2012, pp. 1206–1210.

[12] Singh, V. and Willcox, K. E., "Methodology for path planning with dynamic data-driven flight capability estimation," *AIAA Journal*, Vol. 55, No. 8, 2017, pp. 2727–2738.

[13] Antoulas, A. C., *Approximation of Large-Scale Dynamical Systems*, Society for Industrial and Applied Mathematics, 2005.

[14] Benner, P., Gugercin, S., and Willcox, K., "A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems," *SIAM Review*, Vol. 57, No. 4, 2015, pp. 483–531.

[15] Chinesta, F., Ladeveze, P., and Cueto, E., "A Short Review on Model Order Reduction Based on Proper Generalized Decomposition," *Archives of Computational Methods in Engineering*, Vol. 18, 11 2011, pp. 395–404.

[16] Quarteroni, A. and Rozza, G., *Reduced Order Methods for Modeling and Computational Reduction*, Springer, 2014.

[17] Peherstorfer, B. and Willcox, K. E., "Data-driven operator inference for non-intrusive projection-based model reduction," *Computer Methods in Applied Mechanics and Engineering*, Vol. 306, 2016, pp. 196–215.

[18] 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, Jun 2020, pp. 2658–2672.

[19] McQuarrie, S. A., Huang, C., and Willcox, K., "Data-driven reduced-order models via regularised operator inference for a single-injector combustion process," *Journal of the Royal Society of New Zealand*, 2021, pp. 1–18.

[20] Kapteyn, M. G. and Willcox, K., "From Physics-based models to Predictive Digital Twins via Interpretable Machine Learning," *arXiv preprint arXiv:2004.11356*, 2020.

[21] 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, May 2020, pp. 132401.

[22] Harvazinski, M. E., Huang, C., Sankaran, V., Feldman, T. W., Anderson, W. E., Merkle, C. L., and Talley, D. G., "Coupling between hydrodynamics, acoustics, and heat release in a self-excited unstable combustor," *Physics of Fluids*, Vol. 27, No. 4, 2015, pp. 045102.

[23] Martins, B. L. and Kosmatka, J. B., "Detecting damage in a UAV composite wing spar using distributed fiber optic strain sensors," *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA 2015–0447, 2015.

[24] Martins, B. L. and Kosmatka, J. B., "In-Situ Health Monitoring of Aerospace Structures via Dynamic Strain Measurements," *AIAA Scitech 2019 Forum*, AIAA 2019–1758, 2019.

[25] Martins, B. L. and Kosmatka, J. B., "Health Monitoring of Aerospace Structures via Dynamic Strain Measurements: An Experimental Demonstration," *AIAA Scitech 2020 Forum*, AIAA 2020–0701, 2020.

[26] Biedron, R. T., Carlson, J.-R., Derlaga, J. M., Gnoffo, P. A., Hammond, D. P., Jacobson, K. E., Jones, W. T., Kleb, B., Lee-Rausch, E. M., Nielsen, E. J., Park, M. A., Rumsey, C. L., Thomas, J. L., Thompson, K. B., Walden, A. C., Wang, L., and Wood, W. A., *FUN3D Manual 13.7*, NASA TM 20205010139, Langley Research Center, Hampton, Virginia, 2020.

[27] Wood, S., *FUN3D v13.4 Training Session 13: Time-Dependent Simulations*, NASA FUN3D Training Workshop, 2018.

[28] Spalart, P. and Allmaras, S., "A one-equation turbulence model for aerodynamic flows," *30th Aerospace Sciences Meeting and Exhibit*, AIAA 1992-439, January 1992.

[29] Tikhonov, A. N. and Arsenin, V. Y., *Solution of Ill-posed Problems*, Winston & Sons, 1977.

[30] Tikhonov, A. N., Goncharsky, A., Stepanov, V. V., and Yagola, A. G., *Numerical Methods for the Solution of Ill-Posed Problems*, Springer Netherlands, 1995.

[31] Salinger, S., Kapteyn, M., Kays, C., Pretorius, J., and Willcox, K., "A Hardware Testbed for Dynamic Data-Driven Aerospace Digital Twins," *Dynamic Data Driven Applications System*, edited by F. Darema, E. Blasch, S. Ravela, and A. Aved, Springer International Publishing, Cham, 2020, pp. 37–45.