

Real-time aerodynamic load estimation for hypersonics via strain-based inverse maps

Julie V. Pham^{*}, Omar Ghattas[†], Noel T. Clemens[‡], and Karen E. Willcox[§]
The University of Texas at Austin, Austin, TX 78712, USA

This work develops an efficient real-time inverse formulation for inferring the aerodynamic surface pressures on a hypersonic vehicle from sparse measurements of the structural strain. The approach aims to provide real-time estimates of the aerodynamic loads acting on the vehicle for ground and flight testing, as well as guidance, navigation, and control applications. Specifically, the approach targets hypersonic flight conditions where direct measurement of the surface pressures is challenging due to the harsh aerothermal environment. For problems employing a linear elastic structural model, we show that the inference problem can be posed as a least-squares problem with a linear constraint arising from a finite element discretization of the governing elasticity partial differential equation. Due to the linearity of the problem, an explicit solution is given by the normal equations. Pre-computation of the resulting inverse map enables rapid evaluation of the surface pressure and corresponding integrated quantities, such as the force and moment coefficients. The inverse approach additionally allows for uncertainty quantification, providing insights for theoretical recoverability and robustness to sensor noise. Numerical studies demonstrate the estimator performance for reconstructing the surface pressure field, as well as the force and moment coefficients, for the Initial Concept 3.X (IC3X) conceptual hypersonic vehicle.

I. Introduction

Modern aerospace vehicles demand robust flight control systems that can operate autonomously under dynamic, uncertain, and extreme environments. Central to this requirement is the ability to obtain accurate information about the aerodynamic state of the vehicle. This information enables the deployment of advanced predictive capabilities for autonomy, digital twins, and guidance, navigation, and control (GNC). These desired capabilities become more challenging under hypersonic flight conditions, where the harsh aerothermal environment is inhospitable to many sensing technologies. In particular, it is difficult to obtain information about aerodynamic loads, such as the surface pressure field on the vehicle. The high-temperature hypersonic environment often limits the use of external surface sensors. Other conventional sensing technologies include inertial measurement units, GPS, satellites, and optical lidar, but these may not provide direct or sufficiently accurate estimates of the aerodynamic quantities of interest, and often have external dependencies which may be denied in adversarial scenarios. In this work we develop a vehicle-as-a-sensor concept [1], where the deformation of the vehicle is used to infer the instantaneous aerodynamic surface pressures. Using this strategy to obtain accurate aerodynamic load estimates in real time can help inform GNC systems, enhance maneuverability, and improve reliability for hypersonics.

Aerodynamic loads are characterized by their integrated quantities, namely, the force and moment coefficients. Estimation of these aerodynamic parameters is commonly achieved using filtering methods such as the extended Kalman filter and unscented Kalman filter [2, 3]. The measurement models rely on the availability of measurements of the air flow angles, such as the angle of attack and sideslip angle, which are difficult to measure accurately. To mitigate this, Morelli et al. reconstruct the air flow angles from inertial data [4] for estimating the aerodynamic parameters. In this paper, we take the approach of using discrete strain measurements to infer the full surface pressure field, and consequently the aerodynamic force and moment coefficients. To do this, we must obtain a (inverse) mapping from strain to aerodynamic quantities of interest. The strain response induced by the aerodynamic loads is governed by the partial differential equations (PDE) of linear elasticity. This PDE underpins the strain measurement model for the inference task, leading to a PDE-constrained inverse problem.

^{*}Graduate Research Assistant, Department of Aerospace Engineering and Engineering Mechanics, AIAA Student Member.

[†]Professor, Walker Department of Mechanical Engineering, Oden Institute for Computational Engineering and Sciences.

[‡]Professor, Department of Aerospace Engineering and Engineering Mechanics, AIAA Fellow.

[§]Professor, Department of Aerospace Engineering and Engineering Mechanics, AIAA Fellow.

Solutions to PDE-constrained inverse problems have been formulated in both the deterministic [5] and statistical (Bayesian) [6] settings, and have great utility in "outer-loop" problems such as aerospace design optimization [7–9]. These solutions are challenging to compute for a number of reasons. First, the measurements from which we infer the solution are often noisy and sparse. Additionally, PDE solution operators for many engineering problems of interest are smoothing, meaning that the output of the physical process is smoother than the input. For example, the stress field in a problem governed by linear elasticity is smoother than the loading field that induced it. These characteristics result in an *ill-posed* inverse problem. A further challenge is that these inverse solutions often require many evaluations of the high-fidelity physics simulations, which is computationally intractable for real-time sensing problems. However, the linearity of the elasticity PDE in our problem allows for the formulation of a statistical least-squares problem with linear constraints, which can be solved explicitly via the normal equations. The resulting estimator is a strain-based inverse map that can be pre-computed and queried rapidly in real-time. Notably, using the statistical formulation of the least-squares problem, the known uncertainty due to sensor noise in the observed data also leads to a quantifiable uncertainty for the estimated solution. Consequently, the estimator fully captures the elastic behavior of the hypersonic vehicle described by the high-fidelity physics, while achieving the desired goal of rapid evaluation and uncertainty quantification.

The proposed inverse approach relies on the availability and amenability of discrete, sparse strain measurements to provide sufficient information about the structural state induced by the aerodynamic loads. Strain-based sensing, particularly fiber-optics, has been demonstrated to be capable of recovering aeroelastic shape under both static and dynamic loads [10], and has also been used for real-time control and monitoring applications [11, 12]. Our work extends the application of strain-based sensing to recover the full aerodynamic pressure field, as well as aerodynamic force and moment parameters. Some additional modeling aspects of this work rely on computational fluid dynamics (CFD) simulations for hypersonics to produce prior information for the surface pressure field. We note that the advancement of both simulation capabilities [13–17] and measurement technologies [18] have progressed greatly, and are key enabling factors for inverse problems which must leverage both resources.

An alternative approach to solving the inverse problem is to employ machine learning to train a rapid-to-evaluate inverse mapping from the measurements to the quantities of interest. The training data can come from high-fidelity physics simulations, which are queried over a large range of flight conditions, as well as experimental data. Our previous work [1] has employed physics simulations to train optimal classification trees (OCTs) [19] for the inverse mapping. The OCTs demonstrated strong prediction accuracy, in addition to being interpretable and rapid to evaluate. Neural networks are another possible option [20] which are fast to evaluate and have powerful representative capacity for nonlinear mappings, but they often lack interpretability and may perform poorly in extrapolation. While the accuracy performance may be reasonable, these machine-learned inverse maps lack a strong theoretical foundation for quantifying uncertainty and ensuring robustness in the predictions. In this paper, we seek to offer the same evaluation speed as the machine learning tools while remaining physics-informed and providing reliable quantification of uncertainty.

The remainder of this paper is outlined as follows: in Section II, we present the formulation of the proposed estimator and the corresponding analytical uncertainty quantification. Section III presents the testbed problem for demonstrating this work, along with numerical results. Finally, we discuss the conclusions and future work in Section IV.

II. Real-time surface pressure estimation: Problem statement and approaches

In this section, we present the inverse problem statement and approaches. Section II.A describes the problem definition and the PDE forward model. Section II.B formulates the solution to the inverse problem via a least-squares estimator, with data-driven modeling considerations outlined in Section II.C. Section II.D provides a detailed analysis of the expected performance and uncertainty of the estimator in two possible settings.

A. Problem Definition

We consider the structural and aerodynamic state of a hypersonic vehicle to be quasi-static at a particular instant in time. The structural state is observed through measurements of strain, with sensors placed on the interior surface of the vehicle. We denote the observed data by an n_d -dimensional vector $\mathbf{d} \in \mathcal{D}$, where $\mathcal{D} \subset \mathbb{R}^{n_d}$ is the space of possible strain measurements. The quantities of interest, $\mathbf{q} \in \mathcal{Q}$, are an n_q -dimensional discrete representation of the pressure field over the surface of the vehicle, where $\mathcal{Q} \subset \mathbb{R}^{n_q}$ is the space of possible pressure field representations. The goal is to infer \mathbf{q} from the observed data \mathbf{d} . To perform the inference task, we seek an estimator that represents the inverse mapping

$$\mathcal{T} : \mathcal{D} \rightarrow \mathcal{Q} \quad (1)$$

which can be evaluated rapidly in real-time. The relationship between the true aerodynamic surface pressure and the vehicle strain is described by a physics-based structural forward model, which we will later exploit in performing the inference task. Consider a discrete pressure field $\mathbf{p}(\boldsymbol{\mu})$ of dimension n_p which is dependent on the flight conditions at a particular instant in time, $\boldsymbol{\mu} = [M, \alpha, \beta, H]^\top$, where M is the Mach number, α the angle of attack, β the sideslip angle, and H the altitude. This pressure can be obtained, for example, by solving a computational fluid dynamics (CFD) model. Our inverse problem formulation works with \mathbf{q} , which are the pressure quantities of interest corresponding to a discrete pressure field $\mathbf{p}(\boldsymbol{\mu})$. For example, \mathbf{q} could be a vector containing $n_q < n_p$ nodal values of \mathbf{p} , or \mathbf{q} could be a vector containing n_q modal coefficients that define a representation of \mathbf{p} in some modal basis. The pressure quantities of interest \mathbf{q} are the inputs to the structural forward model, denoted by $F_{\text{struct}}(\mathbf{u}_{\text{struct}}; \mathbf{q}) = 0$, where $\mathbf{u}_{\text{struct}}$ is the displacement of the deformed vehicle. After solving the forward model for $\mathbf{u}_{\text{struct}}$, the model strain data at the sensor locations is computed through the observation equation $\mathbf{y} = G_{\text{struct}}(\mathbf{u}_{\text{struct}})$. Ultimately, the model strain response \mathbf{y} can be compared to the measured strain data \mathbf{d} , thus driving the formulation of the inverse problem. Figure 1 illustrates the overview of the relationship between \mathbf{p} , \mathbf{d} , and intermediate quantities.

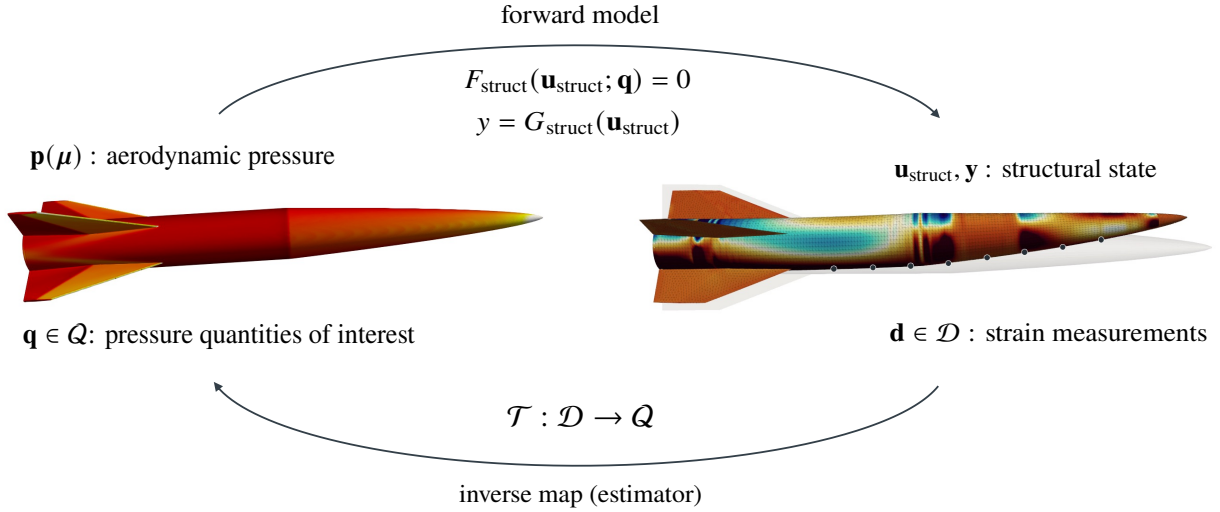


Fig. 1 Inverse problem overview: the forward model maps surface pressure quantities of interest to the structural state, and the inverse map relates strain measurements to pressure quantities of interest.

To explicitly define F_{struct} and G_{struct} , we consider the equations of linear elasticity on the domain $\Omega \subset \mathbb{R}^3$, with exterior and interior surfaces $\partial\Omega_{\text{ext}}$ and $\partial\Omega_{\text{int}}$, respectively, and aft end $\partial\Omega_{\text{aft}}$. The governing partial differential equation (PDE) and boundary conditions for the displacement $\mathbf{u}(\mathbf{x})$ for $\mathbf{x} \in \Omega$ are given by

$$\begin{cases} -\nabla \cdot \boldsymbol{\sigma}(\mathbf{u}) = \mathbf{0} & \text{in } \Omega \\ \boldsymbol{\sigma}(\mathbf{u}) \cdot \mathbf{n} = \mathbf{0} & \text{on } \partial\Omega_{\text{int}} \\ \boldsymbol{\sigma}(\mathbf{u}) \cdot \mathbf{n} = \mathbf{t} & \text{on } \partial\Omega_{\text{ext}} \\ \mathbf{u} = \mathbf{0} & \text{on } \partial\Omega_{\text{aft}} \end{cases} \quad (2)$$

where $\boldsymbol{\sigma}$ is the Cauchy stress tensor, \mathbf{n} is the outward-pointing surface normal, and \mathbf{t} is the traction defined on the external boundary $\partial\Omega_{\text{ext}}$. The traction boundary condition is derived from the aerodynamic surface pressure. The strain-displacement equation is given by $\boldsymbol{\varepsilon} = \frac{1}{2}(\nabla\mathbf{u} + \nabla\mathbf{u}^\top)$ where $\boldsymbol{\varepsilon}$ is the infinitesimal strain tensor, and the constitutive law is given by $\boldsymbol{\sigma} = 2S\boldsymbol{\varepsilon} + \frac{S(E-2S)}{3S-E}\text{tr}(\boldsymbol{\varepsilon})\mathbf{I}$, where S is the shear modulus and E is the Young's modulus of the material. We consider a finite element spatial discretization of the above elasticity equations. Upon discretization, the forward structural model is given by $F_{\text{struct}} = \mathbf{A}\mathbf{u}_{\text{struct}} - \mathbf{C}\mathbf{q} = \mathbf{0}$, where \mathbf{A} is the stiffness matrix, and \mathbf{C} relates the pressure quantities of interest \mathbf{q} to the nodal forces \mathbf{f} on $\partial\Omega_{\text{ext}}$, such that $\mathbf{f} = \mathbf{C}\mathbf{q}$. The result is a large, sparse linear system with dimension n_s , the number of degrees of freedom in the structural discretization. We also define $\mathbf{y} = G_{\text{struct}}(\mathbf{u}_{\text{struct}}) = \mathbf{B}\mathbf{u}_{\text{struct}}$, where \mathbf{B} is a spatial observation operator that (1) maps the discrete displacements to discrete strains using the strain-displacement relation and (2) extracts the particular strains corresponding to the direction and location of each sensor from the forward model.

B. Real-time Inverse Solutions via Statistical Least Squares

We formulate the inference problem of inferring pressure quantities of interest \mathbf{q} from measurements \mathbf{d} in general as a regularized, weighted least-squares problem. The main objective is to minimize the *data misfit* between a given strain measurement and the forward model strain predictions at the sensor locations. We add regularization and weights to provide statistical information about the measurements, via the sensor noise covariance $\mathbf{\Gamma}_n$, as well as the quantities of interest, in the form of a prior mean and covariance, $\bar{\mathbf{q}}$ and $\mathbf{\Gamma}_{pr}$, respectively. The selection of these properties are further detailed in Section II.C. Then, the least-squares objective is given by

$$\min_{\mathbf{q}} \underbrace{\frac{1}{2} \|\mathbf{B}\mathbf{u}_{\text{struct}} - \mathbf{d}\|_{\mathbf{\Gamma}_n^{-1}}^2}_{\text{data misfit}} + \underbrace{\frac{\gamma}{2} \|\mathbf{q} - \bar{\mathbf{q}}\|_{\mathbf{\Gamma}_{pr}^{-1}}^2}_{\text{regularization}} \quad \text{subject to } \mathbf{A}\mathbf{u}_{\text{struct}} - \mathbf{C}\mathbf{q} = \mathbf{0}, \quad (3)$$

where γ is a regularization parameter. The notable challenge is that the constraint is derived from a high-dimensional finite element model, and the solution of Equation 3 typically necessitates many evaluations of this expensive forward model. However, due to the linearity of the constraint, the least-squares problem offers an explicit solution via the normal equations, given by

$$\hat{\mathbf{q}} = (\mathbf{C}^T \mathbf{A}^{-T} \mathbf{B}^T \mathbf{\Gamma}_n^{-1} \mathbf{B} \mathbf{A}^{-1} \mathbf{C} + \gamma \mathbf{\Gamma}_{pr}^{-1})^{-1} (\mathbf{C}^T \mathbf{A}^{-T} \mathbf{B}^T \mathbf{\Gamma}_n^{-1} \mathbf{d} + \gamma \mathbf{\Gamma}_{pr}^{-1} \bar{\mathbf{q}}). \quad (4)$$

This is the general form of our estimator, or inverse map, obtained via a least-squares approach. We note that $\mathbf{u}_{\text{struct}} = \mathbf{A}^{-1} \mathbf{C} \mathbf{q}$; it follows that the parameter-to-observable map is given by $\mathbf{y} = \mathbf{B}\mathbf{u}_{\text{struct}} = \mathbf{B}\mathbf{A}^{-1} \mathbf{C} \mathbf{q}$. We let $\mathbf{Z} = \mathbf{B}\mathbf{A}^{-1} \mathbf{C}$ denote the Jacobian of the parameter-to-observable map. Substituting in Equation 4 and expanding, we obtain

$$\hat{\mathbf{q}} = (\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \gamma \mathbf{\Gamma}_{pr}^{-1})^{-1} (\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{d} + \gamma \mathbf{\Gamma}_{pr}^{-1} \bar{\mathbf{q}}) \quad (5)$$

$$= \underbrace{[(\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \gamma \mathbf{\Gamma}_{pr}^{-1})^{-1} \mathbf{Z}^T \mathbf{\Gamma}_n^{-1}]}_{:=\mathbf{T}} \mathbf{d} + \underbrace{(\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \gamma \mathbf{\Gamma}_{pr}^{-1})^{-1} \gamma \mathbf{\Gamma}_{pr}^{-1} \bar{\mathbf{q}}}_{:=\mathbf{k}} \quad (6)$$

We note that the matrix \mathbf{Z} contains the system matrices \mathbf{A} , \mathbf{B} and \mathbf{C} , which are measurement-independent, and can therefore be constructed in advance to pre-compute \mathbf{Z} . We additionally have a known, fixed $\mathbf{\Gamma}_n$, $\bar{\mathbf{q}}$, and $\mathbf{\Gamma}_{pr}$. Therefore, in Equation 6, the matrix $\mathbf{T} = (\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \gamma \mathbf{\Gamma}_{pr}^{-1})^{-1} \mathbf{Z}^T \mathbf{\Gamma}_n^{-1}$ can be entirely pre-computed, resulting in a matrix of size $n_q \times n_d$. Similarly, the term $\mathbf{k} = (\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \gamma \mathbf{\Gamma}_{pr}^{-1})^{-1} \gamma \mathbf{\Gamma}_{pr}^{-1} \bar{\mathbf{q}}$ can be pre-computed. Upon deployment of the estimator, for a new measurement \mathbf{d} , we can rapidly estimate the quantities of interest \mathbf{q} by simply performing the matrix-vector product $\mathbf{T}\mathbf{d}$, plus a shifting term \mathbf{k} due to the prior mean. Notably, this computation can realistically be performed in real-time, even for large n_q , since the sensor dimension n_d is typically moderate. We note that integrated quantities of interest, such as force and moment coefficients, can also easily be computed by reconstructing the pressure field from \mathbf{q} , and may be especially useful for real-time guidance, navigation, and control applications. Consequently, this estimator is a possible efficient real-time sensing strategy for hypersonic vehicles.

C. Data-Driven Modeling Considerations

1. Noise and prior distributions

We characterize two important features for statistical modeling: (1) the measurement noise model, and (2) the prior distribution. For the noise model, we consider additive Gaussian noise such that $\mathbf{d} = \mathbf{B}\mathbf{u}_{\text{struct}} + \boldsymbol{\eta} = \mathbf{y} + \boldsymbol{\eta}$, where the random variable $\boldsymbol{\eta} \sim (\mathbf{0}, \mathbf{\Gamma}_n)$ is the noise discrepancy between the model strain response and the measurement. Here, we assume independent, identically distributed sensor measurements, with zero mean and variance σ^2 , such that $\mathbf{\Gamma}_n = \sigma^2 \mathbf{I}$. This variance can be approximated through characterization of the physical sensor noise levels. Substituting $\mathbf{\Gamma}_n$ in Equation 4, the noise covariance-weighted least-squares solution simplifies to the ordinary least-squares solution if there is no regularization. In practice, we also perform a whitening transformation to normalize the variance $\boldsymbol{\eta}' = \mathbf{L}^{-1} \boldsymbol{\eta}$, where \mathbf{L} is computed through the Cholesky factorization $\mathbf{\Gamma}_n = \mathbf{L}\mathbf{L}^T$ for numerical stability.

The prior distribution of the quantities of interest is modeled by a Gaussian with mean $\bar{\mathbf{p}}$ and covariance $\mathbf{\Gamma}_{pr}$. We construct the mean and covariance in a data-driven fashion by using a CFD model, denoted by $F_{\text{aero}}(\mathbf{p}; \boldsymbol{\mu}) = 0$, to compute surface pressure snapshots over a range of flight conditions $\boldsymbol{\mu}_j$ for $j = 1, \dots, N$, where N is the total number of flight condition combinations of interest. The pressure snapshots are denoted by \mathbf{p}_j , each of which correspond to the

flight conditions μ_j . We define $\bar{\mathbf{p}}$ to be the mean over the snapshots. Then, we construct the covariance of the prior by computing the covariance over the snapshots,

$$\mathbf{\Gamma}_{\text{pr}} = \sum_{j=1}^N \frac{(\mathbf{p}_j - \bar{\mathbf{p}})(\mathbf{p}_j - \bar{\mathbf{p}})^\top}{N-1}. \quad (7)$$

This covariance captures spatial correlations in the pressure field from the data, including features such as smoothness or discontinuities. This is important for recovery of these features in the estimated pressure field if they are not strictly informed by the measurements, which we further describe in Section II.D.2.

2. Dimensionality reduction

Since the surface pressure field \mathbf{p} is high-dimensional, we may seek to represent the pressure field in a low-dimensional subspace using the proper orthogonal decomposition (POD). The pressure snapshots described above, \mathbf{p}_j , are collected in a snapshot matrix \mathbf{P} . We center the snapshot matrix by the column-wise mean, $\bar{\mathbf{p}}$, to obtain $\tilde{\mathbf{P}}$. We perform the singular value decomposition on the centered snapshot matrix, $\tilde{\mathbf{P}} = \mathbf{V}\mathbf{\Sigma}\mathbf{W}^\top$. We then choose to retain the first r left singular vectors. The reconstructed surface pressure field is then given by

$$\mathbf{p} \approx \bar{\mathbf{p}} + \sum_{i=1}^r c_i \mathbf{v}_i \quad (8)$$

where c_i and \mathbf{v}_i are the i -th POD coefficient and left singular vector, respectively. This allows us to represent the high-dimensional pressure field using r coefficients.

D. Analysis and Uncertainty Quantification

Up to this point, the specific representation of pressure quantities of interest \mathbf{q} have not been detailed. In the following sections, we outline two cases for \mathbf{q} : first, the case where we have fewer parameters than measurements, $n_q \leq n_d$, such that the least squares problem is over-determined, and second, the case where there are more parameters than measurements, $n_q > n_d$, such that the inverse problem is ill-posed. We provide a discussion of uncertainty quantification for each case, as well as an analysis of recoverable surface pressure information for a given sensor configuration.

1. Case 1: $n_q \leq n_d$

Consider the case where the number of parameters n_q is less than or equal to the number of measurements, n_d . In this setting, we let \mathbf{q} be a low-dimensional parameterization of the surface pressure field \mathbf{p} so that $n_q \leq n_d$. Here, we employ the POD decomposition given in Equation 8. First, let \mathbf{C}_{map} define the mapping from the pressure field to the resulting nodal forces, $\mathbf{f} = \mathbf{C}_{\text{map}}\mathbf{p}$. We then define $\mathbf{C} := \mathbf{C}_{\text{map}}\mathbf{V}_r$, where the columns of \mathbf{V}_r are the first r left singular vectors of the snapshot matrix, and $q_i := c_i$, the corresponding POD coefficients. We seek to estimate the coefficients \mathbf{q} from the data to reconstruct the pressure field. Here, we consider the weighted least squares problem without regularization ($\gamma = 0$), meaning we seek the unique solution to the over-determined system of equations $\mathbf{Z}\mathbf{q} = \mathbf{d}$, where \mathbf{Z} is full-rank. To analyze this problem, we consider an infinitesimal perturbation $\delta\mathbf{d}$ to the strain response $\mathbf{y} = \mathbf{Z}\mathbf{q}$. We seek to characterize the corresponding error in the least squares solution $\delta\mathbf{q} = \hat{\mathbf{q}} - \mathbf{q}$ with respect to the perturbation $\delta\mathbf{d}$. Standard sensitivity analysis [21] of the ordinary least-squares solution to small perturbations in the right-hand side yields a relative condition number given by

$$\mathcal{K} = \sup_{\delta\mathbf{d}} \left(\frac{\|\delta\mathbf{q}\|}{\|\mathbf{q}\|} \middle/ \frac{\|\delta\mathbf{d}\|}{\|\mathbf{d}\|} \right) = \frac{\kappa(\mathbf{Z})}{\nu \cos \theta} \quad (9)$$

where $\kappa(\mathbf{Z})$ is the condition number of the parameter to observable map, $\nu = \frac{\|\mathbf{Z}\|\|\mathbf{q}\|}{\|\mathbf{Z}\mathbf{q}\|}$, and $\cos \theta = \frac{\|\mathbf{Z}\mathbf{q}\|}{\|\mathbf{d}\|}$. For $\mathcal{K} \gg 1$, the problem is ill-conditioned and we can expect large relative errors in the solution, even though $\delta\mathbf{d}$ is infinitesimal; this error becomes even larger once we consider a larger perturbation $\boldsymbol{\eta}$. We can improve the conditioning of the problem through the number of sensors as well as sensor placement, since the condition number is directly affected by the information gained from the sensors in the problem. In many least squares problems, it is also useful to add regularization so that the errors are not magnified due to the matrix conditioning. For our problem in Case 1, when $n_q \leq n_d$, instead we drive $\kappa(\mathbf{Z})$ to be sufficiently small through the sensor configuration, so we do not require regularization ($\gamma = 0$). To

further quantify the uncertainty in the estimate, the covariance matrix of the estimator can be computed, given by the inverse of the Hessian of the data misfit term in our least squares objective function,

$$\mathbf{H}_{\text{misfit}}^{-1} = (\mathbf{C}^\top \mathbf{A}^{-\top} \mathbf{B}^\top \mathbf{\Gamma}_n^{-1} \mathbf{B} \mathbf{A}^{-1} \mathbf{C})^{-1} = (\mathbf{Z}^\top \mathbf{\Gamma}_n^{-1} \mathbf{Z})^{-1} \quad (10)$$

This is easy to compute for low-dimensional \mathbf{q} , and it provides an explicit quantification of the uncertainty on the estimated parameters.

2. Case 2: $n_q > n_d$

We now consider the case where the number of parameters n_q is larger than the number of measurements n_d . Particularly, we seek to estimate the distributed loads without restrictions to a particular low-dimensional parameterization using POD or otherwise. Instead, we seek to directly infer the surface load parameter field. In this setting, we define $\mathbf{q} := \mathbf{p}$, and $\mathbf{C} := \mathbf{C}_{\text{map}}$. This formulation results in an under-determined system of linear equations, for which there are many solutions. Of these solutions, it is sometimes sufficient to choose the minimum-norm solution. This solution lies within the row space of \mathbf{Z} , which is the data-informed subspace of the parameter field, and is the best we can achieve without prior information. However, we can significantly improve upon this result by constructing a prior, which enters through the regularization term in the least-squares problem.

We compute the covariance matrix as in Equation 7. We note that this data-driven covariance will be rank-deficient when the number of snapshots N used to construct the covariance is less than n_q (the rank of this matrix will be at most N). Thus, we do not have $\mathbf{\Gamma}_{\text{pr}}^{-1}$ in the explicit solution in Equation 4. However, this can be mitigated by adding a small perturbation $\epsilon \mathbf{I}$ to the prior covariance. Then, the Woodbury matrix identity* can be used for cheap computation of the inverse of the Hessian of the least-squares objective in Equation 3, given by

$$\mathbf{H}^{-1} = (\mathbf{Z}^\top \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \gamma \mathbf{\Gamma}_{\text{pr}}^{-1})^{-1} = \frac{1}{\gamma} (\mathbf{\Gamma}_{\text{pr}} - \mathbf{\Gamma}_{\text{pr}} \mathbf{Z}^\top (\gamma \mathbf{\Gamma}_n + \mathbf{Z} \mathbf{\Gamma}_{\text{pr}} \mathbf{Z}^\top)^{-1} \mathbf{Z} \mathbf{\Gamma}_{\text{pr}}). \quad (11)$$

Analysis of the Hessian provides insights for pressure field recoverability through the prior and the data. To compute the eigenvectors, we factor the Hessian as demonstrated in [22] and solve the eigenvalue problem for non-zero eigenvalues of the prior-preconditioned Hessian, formulated as

$$\mathbf{\Gamma}_{\text{pr}} \mathbf{H}_{\text{misfit}} \mathbf{b}_k = \lambda_k \mathbf{b}_k \quad (12)$$

where λ_k and \mathbf{b}_k are the k -th largest eigenvalue and corresponding eigenvector, respectively, and $\mathbf{H}_{\text{misfit}} = \mathbf{Z}^\top \mathbf{\Gamma}_n^{-1} \mathbf{Z}$. The eigenvectors (modes) of the resulting low-rank spectral decomposition of the prior-preconditioned Hessian define the possible solution space of the least-squares estimator. In a noise-free setting, these modes can be perfectly recovered through the prior and the data. The orthogonal complement, which corresponds to higher frequency (spatially oscillatory) modes, is eliminated from the possible solution space. Since we expect our parameter field to be mostly smooth, this elimination results in a physically tractable solution. Additionally, standard spectral analysis [5] shows that the errors due to noise in the Fourier coefficient corresponding to the k -th mode will scale with $\frac{1}{\lambda_k}$, which leads to amplified errors when λ_k is small. In turn, the regularization parameter γ must be chosen to shift the eigenvalues to reduce the noise amplification, which then also eliminates the contribution of the modes corresponding to eigenvalues $\lambda_k \ll \gamma$. We tune this parameter based on the expected sensor noise levels in order to balance estimator accuracy and noise mitigation.

III. Application: IC3X Testbed Problem

To demonstrate the inverse approaches outlined above, we consider the Initial Concept 3.X (IC3X) hypersonic vehicle. The IC3X was initially proposed by Pasilio et al. [23], and a detailed finite element model for the vehicle was developed by Witeof et al. [24]. The IC3X is a representative hypersonic vehicle with boost, cruise, and terminal phases in a nominal trajectory. The vehicle is propelled by a scramjet in the cruise phase. The geometry of the vehicle is shown in Figure 2. In this paper, we conduct an initial study using a modified version of the IC3X vehicle with removed fins (fuselage only), and a hollow internal structure.

*The Woodbury identity may suffer from ill-conditioning depending on the construction of the regularization. It is important to verify the numerical stability of the computation to avoid this potential pitfall.




Fig. 2 IC3X geometry.

A. IC3X modeling

The forward aerodynamics for the IC3X are solved using CART3D [25–27], which solves the compressible Euler equations for inviscid steady fluid dynamics. CART3D allows for rapid databasing of surface pressure solutions over a wide range of flight conditions. The structural model is constructed and solved using FEniCS[28, 29], which is a software project providing a collection of packages, including DOLFIN[30], that can be used for solving PDEs with the finite element method. Specifically, we employ FEniCS for the explicit construction of the required system matrices \mathbf{B} , \mathbf{A} , \mathbf{C} that appear in Equation 4. We note that the estimator can also be implemented with access to only the system matrix actions on a vector, which may be needed for large-scale problems where it is computationally infeasible to store the system matrices explicitly. We produce the surface mesh for the CFD solver and the volume mesh for the structural model using Gmsh [31]. The resulting surface mesh contains $n_p = 60,538$ degrees of freedom, and the resulting structural mesh contains $n_s = 461,664$ degrees of freedom. The boundary surfaces for the structural model, with boundary conditions given in Equation 2, are illustrated in Figure 3.




Fig. 3 Boundary surfaces of the modified IC3X. We seek to predict the pressure field on the exterior surface $\partial\Omega_{\text{ext}}$. Sensors are placed on the interior surface $\partial\Omega_{\text{int}}$, with fixed aft end $\partial\Omega_{\text{aft}}$.

1. Sensor configuration and noise statistics

In this work, the strain sensor configurations consist of strain gauges placed in rows on the inner surface of the hollow IC3X vehicle structure. Each row consists of nine evenly spaced strain gauges. The in-plane strain of the inner surface can be measured in a particular direction. Measuring the strain in the stream-wise direction (i.e., x-direction) captures the bending strain, which dominates the strain response when the total angle of attack is non-zero. The circumferential strain can also be measured, capturing the strain due to hoop stresses. Figure 4 depicts the sensor placement for two and four rows of sensors, which we will refer to as Configuration 1 and 2, respectively. Configuration 1 contains 27 sensors, of which 18 sensors measure the x-direction strain, and 9 sensors measure the circumferential strain. Configuration 2 contains 54 sensors, which are the same as Configuration 1 plus their mirror through the x-axis. The observation operator \mathbf{B} is constructed for each of these configurations to map the forward model displacement to the strain in the corresponding direction and location as the physical sensors. As described in Section II.C, we assume independent, identically distributed Gaussian noise for each sensor with zero mean and standard deviation σ . For this study, we choose σ to be equal to 1% of the median strain magnitude over all sensors and flight conditions, based on estimated noise characteristics for physical strain gauges.




Fig. 4 Strain sensor configurations: (top) Configuration 1 with 27 sensors, (bottom) Configuration 2 with 54 sensors. (Right) yz -plane view indicating sensor row measurement direction, (left) 3D view of sensor locations, with arrows indicating measurement direction.

2. Flight conditions and surface pressure

For a nominal trajectory of the IC3X, we consider the following flight condition ranges of interest: Mach number 5 – 7, angle of attack 0° – 10° , sideslip angle 0° – 10° , and altitude 20 km. For dimensionality reduction in Case 1, we solve the forward aerodynamics for the set P_1 of all combinations of the parameters $M = \{5.0, 5.5, 6.0, 6.5, 7.0\}$, $\alpha = \{0, 2, 4, 6, 8, 10\}$, $\beta = \{0, 5, 10\}$ at $H = 20\text{km}$ altitude, and collect the corresponding pressure snapshots. This results in a total of $N = 90$ snapshots. We compute the POD using the centered snapshots as in Equation 8. For the data-driven prior in Case 2, we use the set P_2 of all combinations of the parameters $M = \{5, 6, 7\}$, $\alpha = \{-8, -6, -4, -2, 0, 2, 4, 6, 8\}$ and $\beta = \{-8, -6, -4, -2, 0, 2, 4, 6, 8\}$, at $H = 20\text{km}$ altitude. The corresponding $N = 243$ snapshots are then used to compute the prior mean and covariance as in Equation 7. Figure 5 gives a visualization of the pressure at flight conditions $M = 5$, $\alpha = 8$, and $\beta = 0$. At the axial station $x = 1.75\text{m}$ where the angle of the fuselage changes, there is a clear discontinuity in the surface pressure, which is caused by the geometry and occurs at all flight conditions. Additionally, at larger angles of attack, there exists a discontinuity (shock) due to counter-rotating vortices that occur on the leeward side of the vehicle, leading to a region of higher surface pressure along the body. This is visible in the top right view of Figure 5. These types of features are common for hypersonic vehicles, and we seek to accurately estimate these surface pressure fields.




Fig. 5 Pressure visualization at conditions $M = 5$, $\alpha = 8$, $\beta = 0$. Left: Side view of non-dimensional pressure over 2D slice of fluid domain. Right: (top) leeward, (middle) side, and (bottom) windward views of the surface pressure field.[†]

Table 1 Relative condition number for sensor configurations.

	Number of Sensors	Relative Condition Number \mathcal{K}
Configuration 1	27	89.61
Configuration 2	54	47.96

B. Case 1 : Numerical Results

In this section, we provide numerical results for the scenario where $n_q \leq n_d$. The quantities of interest \mathbf{q} we seek to estimate are the POD coefficients parameterizing the surface pressure. For the pressure snapshots corresponding to P_1 , the cumulative energy retained by the first five modes is 99.1%, thus we choose to retain $r = 5$ singular vectors. Therefore, the parameter-to-observable map \mathbf{Z} has dimensions $n_d \times r$. First, we demonstrate the impact of the sensor configuration on the conditioning of the least squares problem. Table 1 shows the relative condition number \mathcal{K} using $\|\mathbf{d}\| = 100\sqrt{n_d}$ and $\|\mathbf{q}\| = 1$ for each sensor configuration. Configuration 2 has the smaller condition number, since it has a larger number of sensors, and will therefore have less sensitivity to noise perturbations for estimating \mathbf{q} .

To numerically demonstrate the performance of the estimator, we produce synthetic noisy measurements by sampling the noise model $\boldsymbol{\eta} \sim (\mathbf{0}, \boldsymbol{\Gamma}_n)$, and adding the realization of noise to the model strain response. The estimated POD coefficients $\hat{\mathbf{q}} = \mathbf{T}\mathbf{d}$ with $\gamma = 0$ are computed for each synthetic measurement. Each estimated coefficient \hat{q}_i is compared to the reference POD coefficient q_i at a particular flight condition using the normalized difference, computed by

$$e_i^{\text{POD}} = \frac{\hat{q}_i - q_i}{\text{range}(q_i)} \quad (13)$$

where $\text{range}(q_i)$ is the range of the i -th POD coefficient over all POD representations of the pressure snapshots in P_1 . Figure 6 shows the normalized difference metric in the predicted coefficients over 50 samples from each flight condition in P_1 , resulting in a total of 4500 samples. We also show the 3σ uncertainty bounds for each e_i^{POD} denoted by the triangle markers, where σ_i^{POD} is the normalized square root of the i -th diagonal of the covariance matrix from Equation 10. The results show that the median difference for both sensor configurations is zero, however, the distribution of differences for Configuration 2 has a smaller variance. This again arises from the improved conditioning in the least-squares problem for Configuration 2, resulting in lower uncertainty in the estimated coefficients. From these estimated POD coefficients, we reconstruct the surface pressure field $\hat{\mathbf{p}}$ as in Equation 8. Figure 7 shows the relative L2 errors in the reconstructed pressure field with respect to the reference pressure field, computed by

$$e^{\text{recon}} = \frac{\|\hat{\mathbf{p}} - \mathbf{p}\|}{\|\mathbf{p}\|}. \quad (14)$$

We observe that Configuration 1 has a median reconstruction error of 7.5% in the reconstruction error, but has a large distribution with some errors reaching over 30%. Configuration 2 has a median of 5.5% and significantly fewer outliers, with all estimated pressure reconstructions lower than 15% relative error.

To visualize the reconstructed pressure fields, Figure 8 and Figure 9 each show an example reconstructed surface pressure from the estimated POD coefficients with a high and low relative error, respectively. The flight conditions in Figure 8 are $M = 5$, $\alpha = 6$, $\beta = 5$, which is in the set P_1 . The reconstruction error was 15.9% using a synthetic noise realization for Configuration 1. The flight conditions in Figure 8 are $M = 7$, $\alpha = 2$, $\beta = 8$, which is a testing condition not in P_1 . The reconstruction error was 3.9% using a synthetic noise realization for Configuration 2. In both cases, the general surface pressure is generally well recovered. We observe that the errors primarily arise from the leeward shock—the discontinuity is not well captured, which is a result of both noise and the POD surface pressure parameterization, particularly in the testing condition, since the shock location is outside of the range space of the pressure snapshots.

[†]For visual clarity, the color range in this and subsequent pressure color bars is truncated at the indicated maximum. Higher pressures exist in a small region at the IC3X nose.

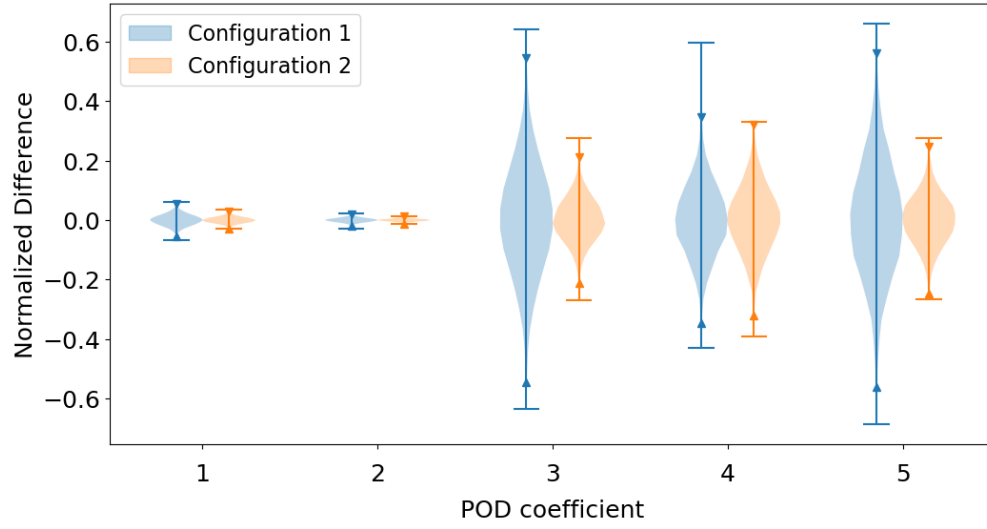


Fig. 6 Normalized difference in estimated POD coefficients for Configuration 1 and 2 with 3σ uncertainty bounds denoted by the triangle markers.




Fig. 7 Relative L2 error in reconstructed pressure field for sensor configuration 1 and 2.

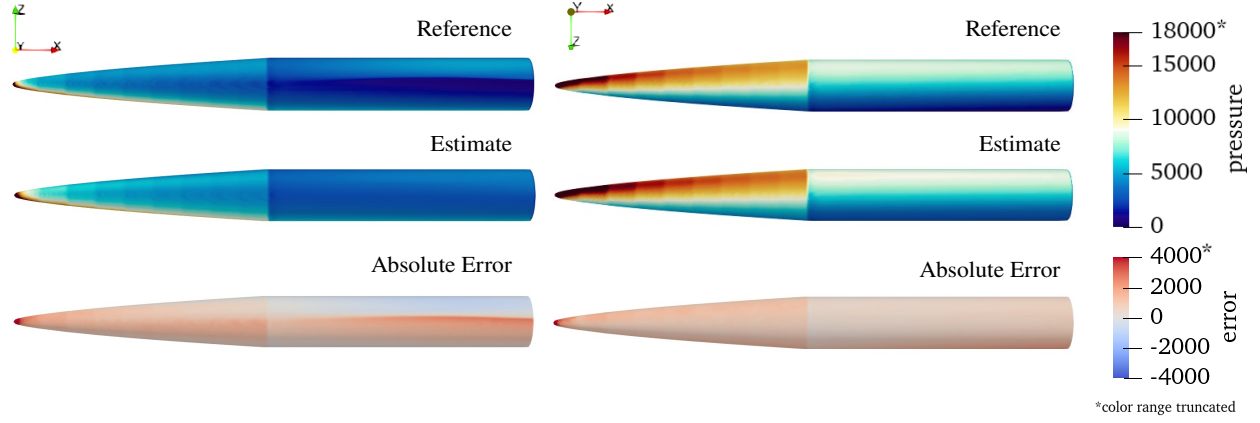


Fig. 8 Visualization of estimated surface pressure compared to the reference surface pressure at conditions $M = 5, \alpha = 6, \beta = 5$. (Left) Positive and (right) negative xz -plane view of reference, estimate, and absolute error in pressure, from top to bottom.

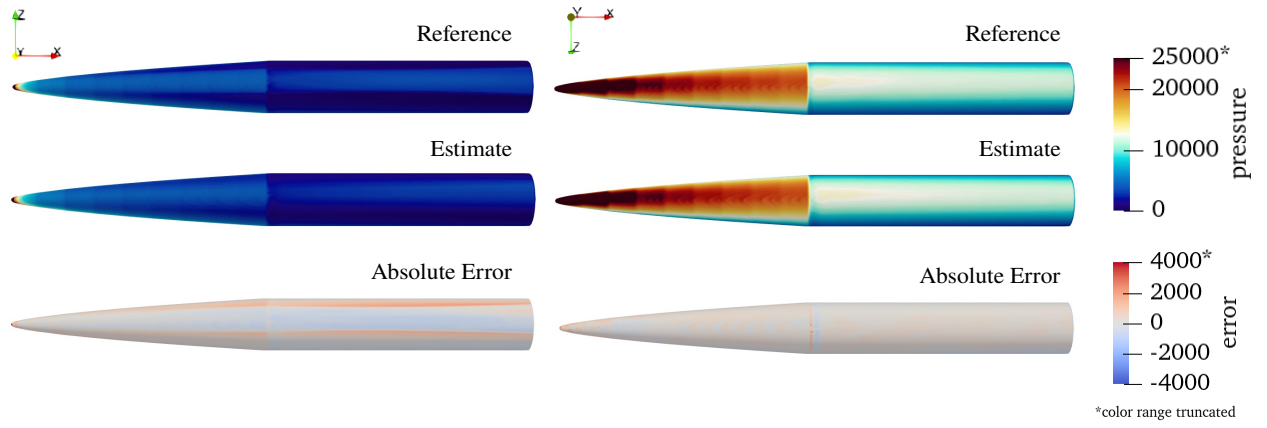


Fig. 9 Visualization of estimated surface pressure compared to the reference surface pressure at test conditions $M = 7, \alpha = 2, \beta = 8$. (Left) Positive and (right) negative xz -plane view of reference, estimate, and absolute error in pressure, from top to bottom.

Furthermore, we compute the aerodynamic force and moment coefficients from the estimated POD coefficients. The aerodynamic coefficients, which are the integrated field quantities of interest, can be computed for a pressure field \mathbf{p} . We denote the body frame force and moment coefficients by the vector $\mathbf{q}_{\text{coeff}} = [C_A, C_N, C_Y, M_P, M_Y]^T$, the elements of which are the coefficients of axial, normal, and side force coefficients, and the pitch and yaw moment coefficients, respectively. These coefficients are computed by

$$\mathbf{q}_{\text{coeff}} = \mathbf{G}\mathbf{f} = \mathbf{G}\mathbf{C}_{\text{map}}\mathbf{p} \quad (15)$$

where \mathbf{G} maps the nodal forces to the force and moment coefficients. For Case 1, since $\mathbf{p} \approx \mathbf{V}_r\mathbf{q} + \bar{\mathbf{p}}$, we can further expand $\mathbf{q}_{\text{coeff}} = \mathbf{G}\mathbf{C}_{\text{map}}\mathbf{V}_r\mathbf{q} + \mathbf{G}\mathbf{C}_{\text{map}}\bar{\mathbf{p}}$. Note that $\mathbf{G}\mathbf{C}_{\text{map}}\mathbf{V}_r$ and $\mathbf{G}\mathbf{C}_{\text{map}}\bar{\mathbf{p}}$ can be precomputed, allowing rapid computation of the force and moment coefficients in real-time, in the case of guidance and control applications. Figure 10 shows the normalized difference for the estimated force and moment coefficients over the same measurement sample set for Configuration 2, computed by

$$e_k^{\text{coeff}} = \frac{\hat{C}_k - C_k}{\max\{\epsilon_k, |C_k|\}} \quad (16)$$

where C_k denotes the k -th coefficient in $\mathbf{q}_{\text{coeff}}$, and ϵ_k is the tolerance where the k -th coefficient is close to zero. We observe that the median normalized difference is close to zero for all coefficients, and the distribution is again a result of the measurement noise.




Fig. 10 Normalized difference in estimated force and moment coefficients compared to reference values.

C. Case 2: Numerical Results

For the case where $n_q > n_d$, we seek to estimate the full surface pressure field. Here, we focus on sensor configuration 2 with $n_d = 54$ sensors. We compute the prior mean and covariance using Equation 7 for the set of flight conditions given by P_2 . To analyze the modes which are well-informed by both the prior and the data, we solve the eigenvalue problem with the prior-preconditioned data misfit Hessian, as described in Equation 12. The decay of the non-zero eigenvalues are shown in Figure 11. We observe that there is rapid decay in the eigenvalues, with a six order-of-magnitude difference between the first and the eighth eigenvalue. Figure 12 provides a visualization of the eigenvectors for $k = 1, 3, 6, 23, 53$. We observe that modes corresponding to larger eigenvalues are relatively smooth, and the modes corresponding to smaller eigenvalues contain higher spatial frequency features. These modes represent the features captured by the prior that are also informed by the measurements. However, due to the rapid decay of the eigenvalues, not all of the modes will be recoverable because these eigenvalues will amplify the noise in the measurements. To mitigate this noise amplification, we require tuning of the regularization parameter γ .




Fig. 11 Eigenvalue decay of prior-preconditioned data-misfit Hessian.




Fig. 12 (Left) Side view and (right) front view of the eigenvectors of the prior-preconditioned Hessian for $k = 1, 3, 6, 23, 53$. We observe increasing spatial frequency features with increasing k .

For a regularization parameter γ , the contribution of the eigenvectors which correspond to eigenvalues $\lambda_k \ll \gamma$ will be very small. Consequently, the projection of a centered pressure field \mathbf{p} onto the span of the first r eigenvectors, choosing r such that $O(\lambda_r) = O(\gamma)$ gives an approximate best-case estimate for a given γ . To demonstrate, we first consider $\gamma = 10^{-6}$ in a noise-free setting to estimate the pressure field at conditions $M = 5, \alpha = 6, \beta = 6$. This is equivalent to retaining all 54 modes in the solution, since all eigenvalues are larger than γ . Using this parameter, the estimated pressure field is computed for a zero-noise measurement, resulting in a relative error of 2.8%. Figure 13 shows the reference pressure, the projection of the reference onto $r = 54$ leading eigenvectors, the estimated pressure, and the absolute error between the reference and the estimate. Here, we observe that the leeward pressure discontinuity is captured by the modes, demonstrating that the strain measurements sufficiently inform this feature of the surface pressure.

With the introduction of noise, we must tune γ to prevent the noise amplification due to the smaller eigenvalues. For the assumed noise level, the regularization was selected using a Morozov-like criterion, where the parameter γ is selected such that norm of the data misfit between the model strain and the measurement, $\|\mathbf{Z}\mathbf{q} - \mathbf{d}\|$, is approximately equal to the noise level, defined as $\delta = \mathbb{E}[\|\boldsymbol{\eta}\|] = \sigma\sqrt{n_d}$. This resulted in a regularization parameter of $\gamma = 10^{1/2}$. We note that a single value for γ was selected for estimation over all flight conditions. Considering the same flight conditions as above, $M = 5, \alpha = 6, \beta = 6$, the estimated pressure field was computed for one hundred different noise realizations, resulting in a relative error range of 13.3% to 27.9%. Figure 14 shows the reference pressure field, the projection $r = 7$, the estimated pressure field for a particular noise realization, which has a relative error of 17.5%, and the absolute error between the reference and the estimate. We observe that solution no longer captures the leeward pressure feature because the regularization has eliminated (smoothed) the modes which inform the discontinuity. However, we have sufficiently prevented exploding errors produced by noise amplification. To further assess the estimator performance, we compute the force and moment coefficients for the estimated pressure field. Over the same one hundred synthetic noise realizations for the flight condition above, the relative errors in the coefficient of normal and side force were below 2%, and the relative errors in the pitch and yaw moment coefficients were below 3.4%. The largest axial force moment error was larger, reaching 16.8%, but with a mean of 5.1% over the one hundred noise realizations. This shows that




Fig. 13 (Left) Leeward view and (right) front view of (top to bottom) the reference pressure, projected reference, zero-noise estimated pressure, and absolute error for $\gamma = 10^{-6}$ at flight conditions $M = 5$, $\alpha = 6$, $\beta = 6$.

the integrated quantity relative errors can be quite low in comparison to the full-field relative error. To understand this behavior, returning to Figure 14, we examine the front view (right column) of the reference and estimated surface pressure for the same flight conditions, as well as the absolute difference. In this view, we can see symmetric structure in the absolute error across the plane with normal $[0, -\cos \frac{\pi}{4}, \cos \frac{\pi}{4}]$ (direction of the total angle of attack). This results in a similar bending moment in comparison to the reference pressure field, which dominates the strain response. Consequently, the coefficients related to normal and side forces have small relative errors. This behavior arises from the terms in the regularized least-squares problem: we seek a data misfit which is of the same order as the noise level, thus the resulting bending moment is accurate enough to invoke a similar strain response, however we are unable to correctly recover higher-frequency features in the pressure field due to the regularization term.

Additionally, we consider flight condition $M = 6$, $\alpha = 0$, $\beta = 5$, which was not in the set P_2 used to construct the prior covariance. Figure 15 shows the reference pressure field, the projection $r = 7$, the estimated pressure field for a particular noise realization, which has a relative error of 7.6%, and the absolute error between the reference and the estimate. For this flight condition, the leeward discontinuity is less prominent in the reference solution because the total angle of attack is smaller. Again, we observe that the leeward feature is not captured by the estimated pressure field, due to the regularization. Over one hundred noise realizations, the relative error range was 4.0% up to 12.0%. This shows consistent estimator performance for flight conditions outside of P_2 . It also shows that there is smaller uncertainty in the estimated pressure field for flight conditions at lower angles of attack due to a smoother leeward pressure feature. This highlights the challenge with correctly identifying the discontinuities in a surface pressure field using strain-based inverse maps.

In summary, the strain-based inverse map using our example sensor configuration captures the structural behavior sufficiently to accurately recover surface pressures over a wide range of flight conditions in a noise-free setting. Therefore, the limitation in pressure recoverability is due to the inherently noisy nature of measurements. We observe that the flexibility of the high-dimensional parameter field in Case 2 still allows the solution to adapt to best match the strain response, resulting in small errors in the performance quantities of interest like the force and moment coefficients. In contrast, the POD parameterization in Case 1 does not enjoy the same advantage, resulting in larger force and moment errors in the presence of sensor noise. In both cases, high spatial frequency features, such as a shock, remain difficult to recover due to sensor noise.

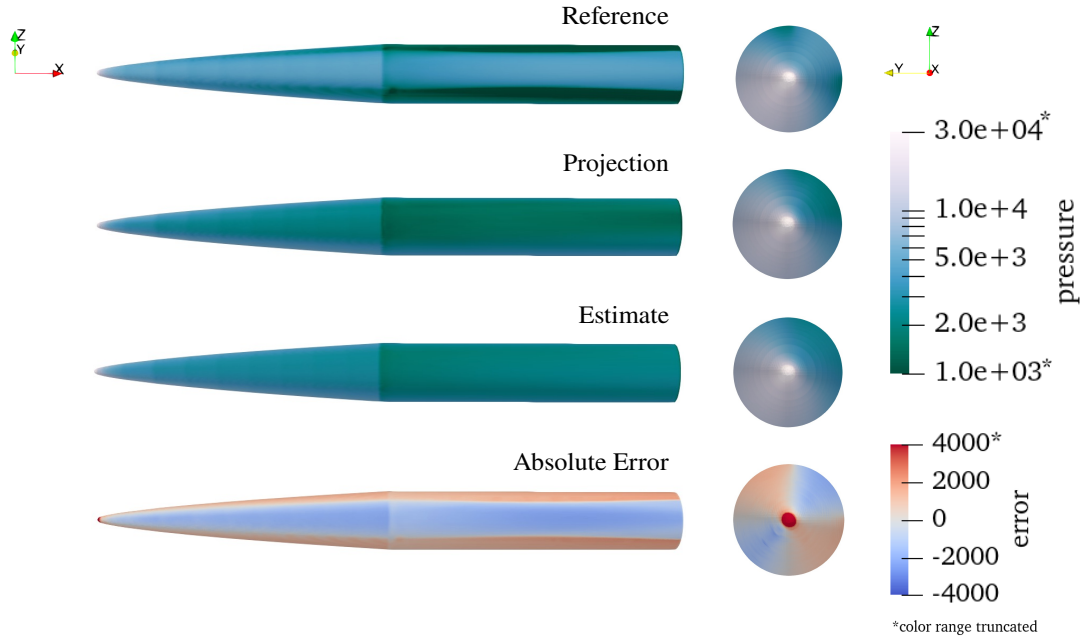


Fig. 14 (Left) Leeward view and (right) front view of (top to bottom) the reference pressure, projected reference, estimated pressure, and absolute error for $\gamma = 10^{1/2}$ at flight conditions $M = 5$, $\alpha = 6$, $\beta = 6$. The relative error of the estimated pressure compared to the reference pressure is 17.5% using a synthetic noisy measurement.




Fig. 15 (Left) Side view and (right) front view of (top to bottom) the reference pressure, projected reference, estimated pressure, and absolute error for $\gamma = 10^{1/2}$ at testing flight conditions $M = 6$, $\alpha = 0$, $\beta = 5$. The relative error of the estimated pressure compared to the reference pressure is 7.6% using a synthetic noisy measurement.

IV. Conclusion

In this work, we have formulated a real-time estimator for aerodynamic loads via a strain-based strategy. Pre-computation of the high-dimensional system matrix multiplications in the resultant estimator, which embed the physics in the solution, enables real-time tractability of the approach. Given sensor noise statistics, the approach also enables uncertainty quantification for the estimated quantities of interest. In the case where the number of unknown pressure quantities is smaller than the number of measurements (i.e., $n_q \leq n_d$), we directly compute the variance of the estimated parameters, demonstrating the impacts of the conditioning of the parameter-to-observable map. In the case where the number of pressure quantities exceeds the number of measurements (i.e., $n_q > n_d$), we analyze the prior-preconditioned Hessian to expose the theoretical recoverability of pressure fields under measurement uncertainty through the prior and the data. Although two different sensor configurations were considered in this paper, the question of optimal sensor placement for strain-based sensing strategies will be addressed in future work. Future work will also consider aerothermal heating effects, which impact the strain response of the vehicle. The developments in this paper, as well as the planned studies, are critical steps towards a deployable strain-based aerodynamic measurement strategy for controls, ground and flight testing in hypersonics.

Acknowledgments

This work was supported by AFOSR grant FA9550-21-1-0089 under the NASA University Leadership Initiative (ULI), the DOE ASCR grant DE-SC0021239, and the DOD MURI grant FA9550-21-1-0084. The authors would like to thank the members of the FAST ULI team for the engaging discussions and support in this work.

References

- [1] Pham, J., Morreale, B. J., Clemens, N., and Willcox, K. E., “Aerodynamic sensing for hypersonics via scientific machine learning,” *AIAA Aviation Forum*, 2022, paper 2022-3717.
- [2] Chowdhary, G., and Jategaonkar, R., “Aerodynamic parameter estimation from flight data applying extended and unscented Kalman filter,” *Aerospace Science and Technology*, Vol. 14, No. 2, 2010, pp. 106–117.
- [3] Garcia-Velo, J., and Walker, B. K., “Aerodynamic Parameter Estimation for High-Performance Aircraft Using Extended Kalman Filtering,” *Journal of Guidance, Control, and Dynamics*, Vol. 20, No. 6, 1997, pp. 1257–1260.
- [4] Morelli, E. A., “Real-Time Aerodynamic Parameter Estimation Without Air Flow Angle Measurements,” *Journal of Aircraft*, Vol. 49, No. 4, 2012, pp. 1064–1074.
- [5] Ghattas, O., and Willcox, K., “Learning physics-based models from data: perspectives from inverse problems and model reduction,” *Acta Numerica*, Vol. 30, 2021, p. 445–554.
- [6] Stuart, A. M., “Inverse problems: a Bayesian perspective,” *Acta Numerica*, Vol. 19, 2010, pp. 451–559.
- [7] Martins, J., Alonso, J., and Reuther, J., “A Coupled-Adjoint Sensitivity Analysis Method for High-Fidelity Aero-Structural Design,” *Optimization and Engineering*, Vol. 6, 2005, p. 33–62.
- [8] Hicken, J. E., and Zingg, D. W., “Aerodynamic Optimization Algorithm with Integrated Geometry Parameterization and Mesh Movement,” *AIAA Journal*, Vol. 48, No. 2, 2010, pp. 400–413.
- [9] Martins, J. R. R. A., and Lambe, A. B., “Multidisciplinary Design Optimization: A Survey of Architectures,” *AIAA Journal*, Vol. 51, No. 9, 2013, pp. 2049–2075.
- [10] Freydin, M., Rattner, M. K., Raveh, D. E., Kressel, I., Davidi, R., and Tur, M., “Fiber-Optics-Based Aeroelastic Shape Sensing,” *AIAA Journal*, Vol. 57, No. 12, 2019, pp. 5094–5103.
- [11] Derkevorkian, A., Masri, S. F., Alvarenga, J., Boussalis, H., Bakalyar, J., and Richards, W. L., “Strain-Based Deformation Shape-Estimation Algorithm for Control and Monitoring Applications,” *AIAA Journal*, Vol. 51, No. 9, 2013, pp. 2231–2240.
- [12] Pena, F., Martins, B., and Richards, W. L., “Active In-flight Load Redistribution Utilizing Fiber-Optic Shape Sensing and Multiple Control Surfaces,” *NASA Technical Memorandum*, 2018, NASA-TM 20180002147.
- [13] Knutson, A. L., Thome, J. S., and Candler, G. V., “Numerical Simulation of Instabilities in the Boundary-Layer Transition Experiment Flowfield,” *Journal of Spacecraft and Rockets*, Vol. 58, No. 1, 2021, pp. 90–99.
- [14] McNamara, J. J., and Friedmann, P. P., “Aeroelastic and Aerothermoelastic Analysis in Hypersonic Flow: Past, Present, and Future,” *AIAA Journal*, Vol. 49, No. 6, 2011, pp. 1089–1122.
- [15] San, O., and Iliescu, T., “Proper Orthogonal Decomposition Closure Models for Fluid Flows: Burgers Equation,” *International Journal of Numerical Analysis Modeling Series B*, Vol. 5, No. 3, 2014, pp. 217–237.
- [16] Falkiewicz, N. J., Cesnik, C. E. S., Crowell, A. R., and McNamara, J. J., “Reduced-Order Aerothermoelastic Framework for Hypersonic Vehicle Control Simulation,” *AIAA Journal*, Vol. 49, No. 8, 2011, pp. 1625–1646.
- [17] Klock, R. J., and Cesnik, C. E. S., “Nonlinear Thermal Reduced-Order Modeling for Hypersonic Vehicles,” *AIAA Journal*, Vol. 55, No. 7, 2017, pp. 2358–2368.
- [18] Musta, M. N., Vanstone, L., Ahn, Y.-J., Eitner, M., Sirohi, J., and Clemens, N., “Investigation of flow-structure coupling for a compliant panel under a shock/boundary-layer interaction using fast-response PSP,” *AIAA Aviation Forum*, 2021, paper 2021-2809.
- [19] Bertsimas, D., and Dunn, J., “Optimal classification trees,” *Machine Learning*, Vol. 106, No. 7, 2017, pp. 1039–1082.
- [20] Meinicke, A. C., and Cesnik, C. E., “On the Efficient Generation of Training Data for High-Speed Aerodynamic State Recovery Through Strain-Load Neural Network Model,” *AIAA SciTech Forum*, 2024, paper 2024-1047.
- [21] Trefethen, L. N., and Bau, D., *Numerical Linear Algebra*, Society for Industrial and Applied Mathematics, 1997.
- [22] Isaac, T., Petra, N., Stadler, G., and Ghattas, O., “Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet,” *Journal of Computational Physics*, Vol. 296, 2015, pp. 348–368.

- [23] Pasiliao, C. L., Sytsma, M. J., Neergaard, L., Witeof, Z., and Troler, J. W., "Preliminary Aero-thermal Structural Simulation," *14th AIAA Aviation Technology, Integration, and Operations Conference*, 2014, AIAA 2014-2292.
- [24] Witeof, Z., and Neergaard, L., "Initial concept 3.0 finite element model definition," AFRL-RWWV-TN-2014-0013, *Eglin Air Force Base, Air Force Research Laboratory*, 2014.
- [25] NASA Ames Research Center, "Automated Triangle Geometry Processing for Surface Modeling and Cartesian Grid Generation (Cart3D)," <https://software.nasa.gov/software/ARC-14275-1>, Accessed 2022.
- [26] Aftosmis, M., Berger, M., and Adomavicius, G., "A parallel multilevel method for adaptively refined Cartesian grids with embedded boundaries," *38th Aerospace Sciences Meeting and Exhibit*, 2000, paper 2000-0808.
- [27] Aftosmis, M. J., Berger, M. J., and Melton, J. E., "Robust and Efficient Cartesian Mesh Generation for Component-Based Geometry," *AIAA Journal*, Vol. 36, 1998, pp. 952–960.
- [28] Alnaes, M. S., Blechta, J., Hake, J., Johansson, A., Kehlet, B., Logg, A., Richardson, C., Ring, J., Rognes, M. E., and Wells, G. N., "The FEniCS Project Version 1.5," *Archive of Numerical Software*, Vol. 3, 2015.
- [29] Logg, A., Mardal, K., Wells, G. N., et al., *Automated Solution of Differential Equations by the Finite Element Method*, Springer, 2012.
- [30] Logg, A., and Wells, G. N., "DOLFIN: Automated Finite Element Computing," *ACM Transactions on Mathematical Software*, Vol. 37, 2010.
- [31] Geuzaine, C., and Remacle, J.-F., "Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities," *International Journal for Numerical Methods in Engineering*, Vol. 79, No. 11, 2009, pp. 1309–1331.