

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

Julie Pham<sup>\*</sup>, Omar Ghattas<sup>†</sup>, and Karen Willcox<sup>‡</sup>  
*The University of Texas at Austin, Austin, TX 78712, USA*

**This work develops an efficient inverse formulation for inferring the aerodynamic surface pressures on a hypersonic vehicle from sparse measurements of the structural strain. The approach aims to provide accurate estimates of the aerodynamic loads acting on the vehicle for real-time guidance, navigation, and control with quantifiable uncertainty. 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 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 optimal sensor placement 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 hypersonic environment limits the availability of measurements such as external pressure sensors, since they would not withstand the thermal heating effects. 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.

---

<sup>\*</sup>Graduate Research Assistant, Department of Aerospace Engineering and Engineering Mechanics, AIAA Student Member.

<sup>†</sup>Professor, Oden Institute for Computational Science and Engineering.

<sup>‡</sup>Director, Oden Institute for Computational Science and Engineering, AIAA Fellow.

Solutions to PDE-constrained inverse problems have been formulated in both the deterministic [5] and statistical (Bayesian) [6] settings. The challenges in this field come from the nature of measurements, which are often noisy, sparse observations, as well as the smoothing effects of many PDE solution operators, both of which can result in ill-posedness. Often, 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 this 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 [7], and has also been used for real-time control and monitoring applications [8, 9]. This 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 [10–13] and measurement technologies [14] 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) [15] 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 which are fast to evaluate, but they often lack interpretability and 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 development of the strain-based inverse map is part of a larger ongoing effort to advance the vehicle-as-a-sensor concept, including experimental validation. Ultimately, we hope to improve vehicle design, testing, and control for hypersonics. The remainder of this paper is outlined as follows: in Section II.A, 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. Surface pressure estimation: problem statement and approaches

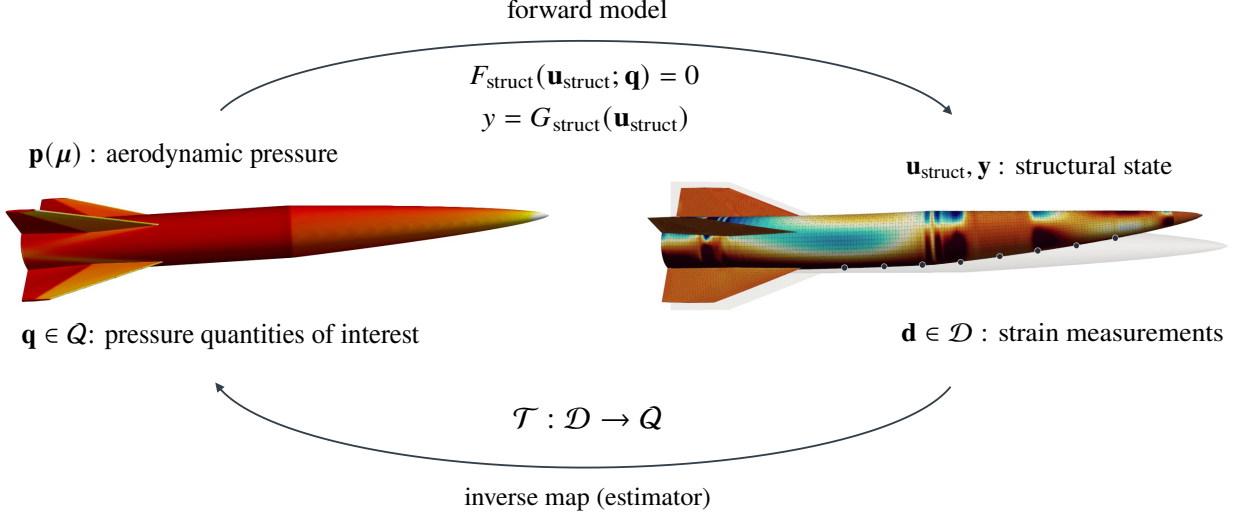
In this section, we present the inverse problem statement and approaches. Section II.A describes the problem definition and the PDE forward model. Section II.B formulates the solution to the inverse problem via a least-squares estimator, with data-driven modeling considerations outlined in Section II.C. 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  $\mathcal{D}$  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. 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

$$\mathcal{T} : \mathcal{D} \rightarrow \mathcal{Q} \tag{1}$$

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,  $\alpha$  the angle of attack,  $\beta$  the sideslip angle, and  $H$  the altitude. This pressure can be obtained, for example, by solving a computational fluid dynamics (CFD) model. The pressure  $\mathbf{p}(\boldsymbol{\mu})$  is represented by  $\mathbf{q}$ , which 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.**

To explicitly define  $F_{\text{struct}}$  and  $G_{\text{struct}}$ , 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 surface pressure  $\mathbf{p}$ . 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  $\mathbf{C}$  is the mapping from the pressure quantities of interest to the nodal forces on  $\partial\Omega_{\text{ext}}$  due to the traction boundary condition  $\mathbf{t}$ . 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. Efficient 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: we seek to minimize the difference between a given strain measurement and the forward model strain predictions at the sensor locations. We choose the 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 objective is given by

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

where  $\lambda$  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

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

This is the general form of the regularized least squares estimator with *a priori* information. 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, we obtain

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

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, the matrix  $\mathbf{T} = (\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \lambda \mathbf{\Gamma}_{pr}^{-1})^{-1} \mathbf{Z}^T \mathbf{\Gamma}_n^{-1}$  can be entirely pre-computed, resulting in a matrix of size  $n_q \times n_d$ . Similarly, the term  $\mathbf{k} = (\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \lambda \mathbf{\Gamma}_{pr}^{-1})^{-1} \lambda \mathbf{\Gamma}_{pr}^{-1} \bar{\mathbf{q}}$  can be pre-computed. Upon deployment of this 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. This computation can be realistically performed in real-time, even for large  $n_q$ , since the sensor dimension is typically moderate. Consequently, this estimator may be used as an efficient real-time sensing strategy for hypersonic vehicles.

## C. Data-Driven Modeling Considerations

### 1. Covariance of the noise and the prior

We must consider two important features for statistical modeling: (1) the measurement noise model, and (2) the prior covariance. 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  $\boldsymbol{\eta}$  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  $\sigma^2$ , 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 can 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 ease of implementation.

For the prior of the quantities of interest, we choose to model the distribution as 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  $\boldsymbol{\mu}_j$ . We define  $\bar{\mathbf{p}}$  to be the mean over these snapshots. The covariance matrix over the snapshots can be computed by

$$\mathbf{\Gamma}_{pr} = \sum_{j=1}^N \frac{(\mathbf{p}_j - \bar{\mathbf{p}})(\mathbf{p}_j - \bar{\mathbf{p}})^T}{N - 1}. \quad (6)$$

This provides prior information about what possible surface pressures may look like over the flight condition range captured by the snapshots, including features such as smoothness or discontinuities. This is important for recovery of these features in the estimated pressure field if they are not informed by the observed data.

## 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,  $\tilde{\mathbf{P}} = \mathbf{V}\Sigma\mathbf{W}^\top$ . We then choose to retain the first  $r$  left singular vectors. The reconstructed surface pressure field is then given by

$$\mathbf{p} \approx \bar{\mathbf{p}} + \sum_{i=1}^r c_i \mathbf{v}_i \quad (7)$$

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 the surface loads  $\mathbf{q}$  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 the information that can possibly be recovered 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 7 and define  $\mathbf{C} := \mathbf{V}_r$ , the 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. Additionally, the aerodynamic parameters we seek may be the integrated quantities from the surface loads, such as force and moment coefficients. These coefficients can easily be computed from the pressure field reconstruction, and may be especially useful for real-time guidance, navigation, and control applications.

To begin the analysis, we consider the weighted least squares problem without regularization ( $\lambda = 0$ ). In this case, we seek the unique solution to the over-determined system of equations  $\mathbf{Z}\mathbf{q} = \mathbf{d}$ . We consider the Gaussian random variable  $\boldsymbol{\eta}$  from the noise model as a small perturbation in  $\mathbf{d}$ , with the measurement model given by  $\mathbf{Z}\mathbf{q} = \mathbf{d} - \boldsymbol{\eta}$ . We seek to quantify the error in the least squares solution  $\Delta\mathbf{q} = \mathbf{q} - \hat{\mathbf{q}}$  with respect to the perturbation  $\boldsymbol{\eta}$ . Standard error analysis [16] of the ordinary least-squares solution to small perturbations in the right-hand side shows that the relative error in the solution, given the size of the perturbation  $\boldsymbol{\eta}$ , is bounded by

$$\frac{\|\Delta\mathbf{q}\|}{\|\mathbf{q}\|} \leq \frac{\kappa(\mathbf{Z})}{\alpha \cos \theta} \frac{\|\boldsymbol{\eta}\|}{\|\mathbf{d}\|} \quad (8)$$

where  $\kappa(\mathbf{Z})$  is the condition number of the parameter to observable map,  $\alpha = \frac{\|\mathbf{Z}\| \|\mathbf{q}\|}{\|\mathbf{Z}\mathbf{q}\|}$ , and  $\cos \theta = \frac{\|\mathbf{Z}\mathbf{q}\|}{\|\mathbf{d}\|}$ . We observe that the upper bound for the relative errors in the solution will be scaled by  $\mathcal{K} = \frac{\kappa(\mathbf{Z})}{\alpha \cos \theta}$ , so for  $\mathcal{K} \gg 1$ , we can expect large relative errors in the solution, even if  $\boldsymbol{\eta}$  is small. We can improve the conditioning of the problem through the number of sensors as well as optimal sensor placement, since the condition number is directly affected by the information gained from the sensors in the problem. In many least squares problems, when  $\mathcal{K} \gg 1$ , it is also useful to add regularization so that the errors due to conditioning are not magnified by the lower modes. For this problem, when  $n_q \leq n_d$ , instead we drive  $\kappa(\mathbf{Z})$  to be sufficiently small through the sensor configuration, so we eliminate the need for regularization ( $\lambda = 0$ ), since all parameters contribute independently to the observables. To further quantify the uncertainty in the estimate, the error covariance of the estimator can be computed by  $\text{var}(\hat{\mathbf{q}}) = \mathbb{E}[\hat{\mathbf{q}}\hat{\mathbf{q}}^\top]$ , since our estimator is unbiased. This is given by

$$\Gamma_{\hat{\mathbf{q}}\hat{\mathbf{q}}} = (\mathbf{C}^\top \mathbf{A}^{-\top} \mathbf{B}^\top \Gamma_n^{-1} \mathbf{B} \mathbf{A}^{-1} \mathbf{C})^{-1} \quad (9)$$

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 would like to directly infer the surface load parameter field. This can be used for estimating the distributed aerodynamic loads during wind tunnel or flight testing. In this setting, we let  $\mathbf{q} = \mathbf{p}$ , and let  $\mathbf{C}$  be the mapping from surface pressure to the nodal forces in the full structural forward model. 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 solution of *minimum-norm*. This solution, which lies in the row space of  $\mathbf{Z}$ , represents the data-informed subspace of the parameter field, and is the best we can hope to achieve with no prior information. However, we can improve upon this result using regularization: here, it is highly important to choose the prior and regularization parameter.

We compute the covariance matrix as in Equation 6. 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, it is not feasible to compute  $\mathbf{\Gamma}_{\text{pr}}^{-1}$  in the explicit solution in Equation 4. However, this can be mitigated by adding a very 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 regularized Hessian, under the assumption that  $\mathbf{\Gamma}_{\text{pr}}$  and  $\mathbf{\Gamma}_n$  are invertible, given by

$$\mathbf{H}^{-1} = (\mathbf{Z}^T \mathbf{\Gamma}_n^{-1} \mathbf{Z} + \lambda \mathbf{\Gamma}_{\text{pr}}^{-1})^{-1} = \frac{1}{\lambda} (\mathbf{\Gamma}_{\text{pr}} + \mathbf{\Gamma}_{\text{pr}} \mathbf{Z}^T (\lambda \mathbf{\Gamma}_n + \mathbf{Z} \mathbf{\Gamma}_{\text{pr}} \mathbf{Z}^T)^{-1} \mathbf{Z} \mathbf{\Gamma}_{\text{pr}}) \quad (10)$$

Substitution of this matrix inversion identity into Equation 5 shows that the solution will lie within the range of  $\mathbf{\Gamma}_{\text{pr}}$ . This can be viewed as constraining the solution to the  $N$ -dimensional subspace spanned by the snapshots used to construct the prior covariance, in consequence eliminating the higher frequency modes that are not informed by the prior or the data. These modes span the null space of the regularized Hessian, which yields the subspace of the parameter field that is unrecoverable through neither the data nor the prior. Since we expect our parameter field to be mostly smooth, this elimination results in a more physically tractable solution. The regularization parameter  $\lambda$  also plays an important role in determining the influence of the prior in relation to the data misfit in the estimator. Tuning  $\lambda$  is critical to obtaining accurate estimates of the parameter field.

## 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 Pasiliao et al. [17], and a detailed finite element model for the vehicle was developed by Witeof et al. [18]. 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.





Fig. 2 IC3X geometry.

### A. IC3X modeling

The forward aerodynamics for the IC3X are solved using CART3D [19], 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[20], which is a software


project providing a collection of packages, including DOLFIN[21], that can be used for solving PDEs with the finite element method. Specifically, FEniCs enables 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 2D surface mesh for the CFD solver and the 3D volume mesh for the structural model using Gmsh [22]. 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  $\partial\Omega_{ext}$ . Sensors are placed on the interior surface  $\partial\Omega_{int}$ , with fixed aft end  $\partial\Omega_{aft}$ .

### 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 mirrored across 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  $\sigma = 2.5e-6$ . This choice of  $\sigma$  is approximately 2% of the expected strain magnitude.



**Fig. 4** Strain sensor configurations, with arrows denoting the direction of the strain measurement. (Top) Configuration 1 with 27 sensors, (bottom) Configuration 2 with 54 sensors.

## 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 7. For the data-driven prior in Case 2, we use the set  $P_2$  of all combinations of the parameters  $\alpha = \{-8, -6, -4, -2, 0, 2, 4, 6, 8\}$  and  $\beta = \{-8, -6, -4, -2, 0, 2, 4, 6, 8\}$ , at Mach number  $M = 5$  and  $H = 20\text{km}$  altitude. The corresponding  $N = 81$  snapshots are then used to compute the prior mean and covariance as in Equation 6. 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 shock due to counter-rotating vortices that occurs on the leeward side of the vehicle, visible in the top right view of the figure. 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 fluid domain slice. Right: (top) leeward, (middle) side, and (bottom) windward views of the surface pressure field.**

## 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 99.4%, thus we choose to retain  $r = 5$  singular vectors. First, we demonstrate the impact of the sensor configuration on the upper bound for the relative errors in  $\mathbf{q}$  as in Equation 8, assuming  $\eta_i = \sigma$  and  $\|\mathbf{q}\| = 100$ . The table below shows the resulting upper bound for each configuration. Configuration 2 has a smaller upper bound on the error in the estimated coefficients since the parameter-to-observable map results in a lower conditioning  $\mathcal{K}$  for noise perturbations in the least-squares problem.

	Number of Sensors	$\frac{\ \Delta\mathbf{q}\ }{\ \mathbf{q}\ }$ upper bound
Configuration 1	27	0.3307
Configuration 2	54	0.1908

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  $\lambda = 0$  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 as

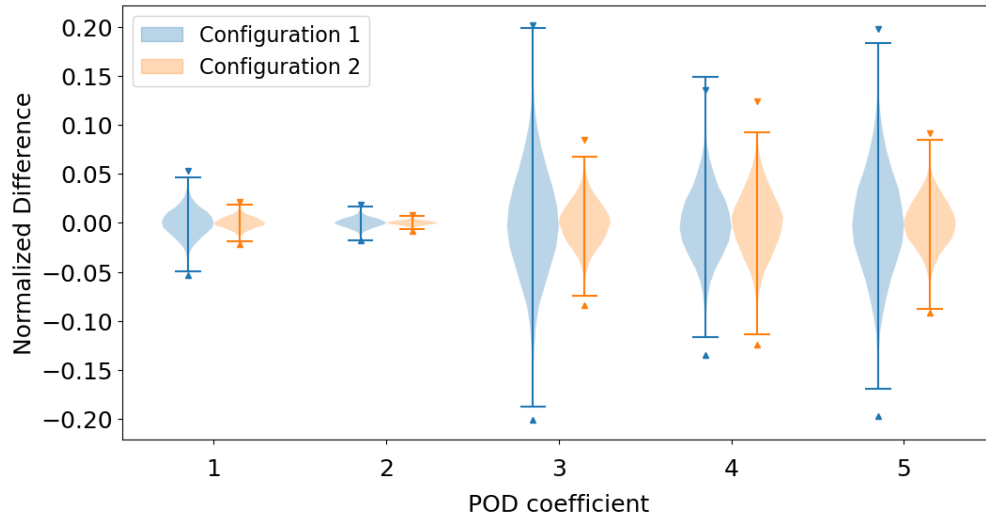
$$e_i^{\text{POD}} = \frac{q_i - \hat{q}_i}{\text{range}(q_i)} \quad (11)$$




where  $\text{range}(q_i)$  is the range of the  $i$ -th POD coefficient over the POD representations of the pressure snapshots. 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  $4\sigma$  uncertainty bounds for each POD coefficient, where  $\sigma_i$  is the square root of the  $i$ -th diagonal of the covariance matrix from Equation 9. The results show that the median difference for both sensor configurations is zero, however, the distribution of 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 7. Figure 7 shows the relative L2 errors in the reconstructed pressure field with respect to the reference pressure field, computed by

$$e^{\text{recon}} = \frac{\|\mathbf{p} - \hat{\mathbf{p}}\|}{\|\mathbf{p}\|} \quad (12)$$

We observe that Configuration 1 has a number of outliers in the reconstruction error, while Configuration 2 has few outliers with all estimated pressure reconstructions lower than 15% relative error.

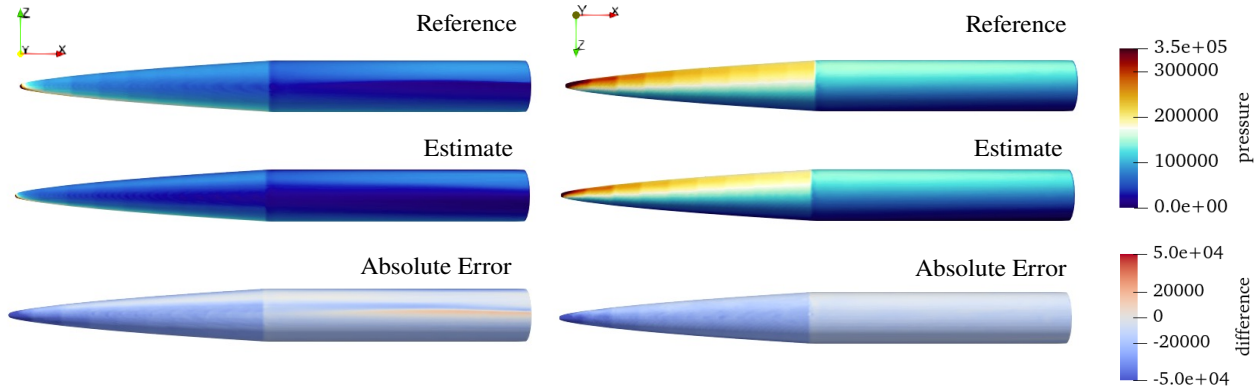


**Fig. 6** Normalized difference in estimated POD coefficients for Configuration 1 and 2 with  $4\sigma$  uncertainty bounds denoted by the triangle markers.

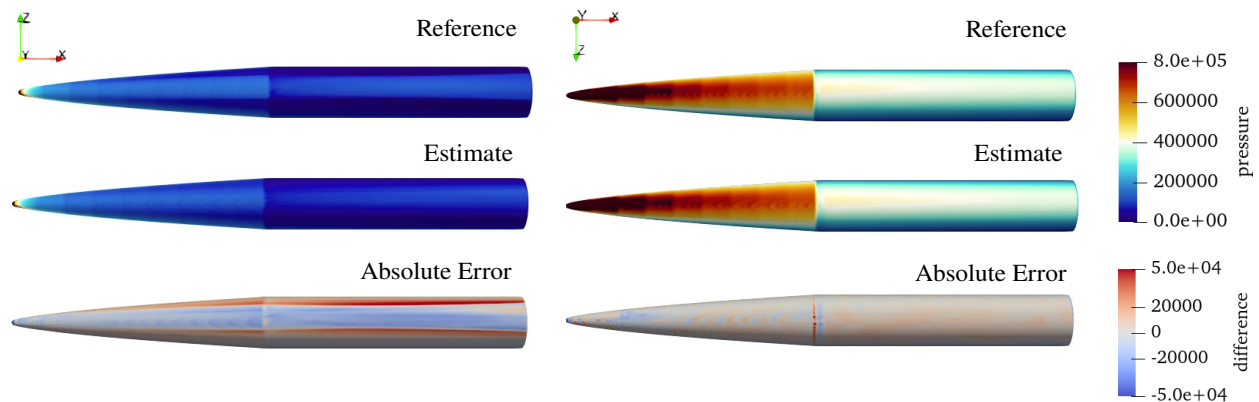


**Fig. 7** Relative L2 error in reconstructed pressure field for sensor configuration 1 and 2.

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 19.5% using a synthetic noise realization for Configuration 1. The flight conditions in Figure 9 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. We see that in both cases, the general surface pressure pattern compared to the reference pressure is recovered well. In the testing flight condition, we note that the errors primarily arise from the leeward shock, which is a result of the POD surface pressure parameterization, since the shock is outside of the range space of the pressure snapshots.



**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 (parameters) from the estimated quantities of interest. We first map  $\mathbf{q}$  to the nodal forces on the surface, and consequently use the distributed force information to compute the aerodynamic parameters. The force and moment coefficient parameters are denoted by  $\mathbf{C}_F = [C_D, C_L, C_S, M_P, M_Y]^T$ , the elements of which are the coefficients of drag, lift, and side force coefficients, and the pitch and yaw moment coefficients, respectively. The conversion to nodal forces from quantities of interest is given by  $\hat{\mathbf{f}} = \mathbf{C}\mathbf{q} + \bar{\mathbf{f}}$ , where  $\bar{\mathbf{f}}$  are the nodal forces corresponding to  $\bar{\mathbf{p}}$ . Then, the aerodynamic parameters are computed by

$$\mathbf{C}_F = \mathbf{G}\hat{\mathbf{f}} = \mathbf{G}\mathbf{C}\mathbf{q} + \mathbf{G}\bar{\mathbf{f}} \quad (13)$$

where  $\mathbf{G}$  is the mapping between the nodal forces and the force and moment coefficients. Note that  $\mathbf{G}\mathbf{C}$  can be precomputed, and will only be of size  $5 \times r$ . Figure 10 shows the normalized difference for the estimated force and moment coefficients over the same measurement sample set for Configuration 2, computed by

$$e_k^{\text{coeff}} = \frac{C_k - \hat{C}_k}{\max\{\epsilon_k, C_k\}} \quad (14)$$

where  $C_k$  denotes the  $k$ -th parameter in  $\mathbf{C}_F$ , and  $\epsilon_k$  is the tolerance for which the  $k$ -th parameter is considered very close to zero. We see that the median of the estimates for all parameters have very small differences to the reference values. We also note that the parameters corresponding to lift forces,  $C_L$  and  $M_P$ , have a smaller variance in comparison to the side or drag forces. This is due to the strain measurement configuration being more informative towards these parameters.




**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. First, we illustrate the ill-posedness of the problem by observing the minimum-norm solution, where  $\lambda \rightarrow 0$ . We consider the noise-free measurement  $\mathbf{d}$  corresponding to  $\mathbf{p}$  at flight conditions  $M = 5$ ,  $\alpha = 6$ , and  $\beta = 6$ . The minimum-norm surface pressure estimate is shown in Figure 11. We observe the estimate contains oscillatory features, due to the high-frequency modes of the parameter-to-observable map. Additionally, we obtain non-negligible negative pressure estimates, which are unphysical. We seek to use regularization via a data-driven prior to obtain better estimates of the surface pressure field.

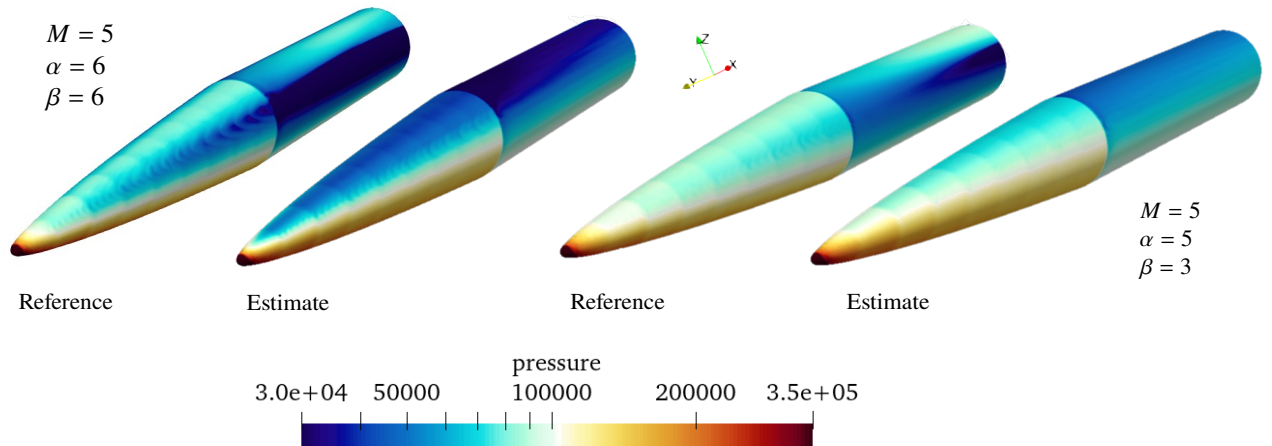
We compute the prior mean and covariance using Equation 6 for the set of flight conditions given by  $P_2$ . This results in a rank-81 matrix. The regularization is tuned for the assumed noise level,  $\sigma = 2.5\text{e-}6$ . The regularization was selected using a Morozov-like criterion, where the parameter  $\lambda$  is tuned such that data misfit between the model strain and the measurement,  $\|\mathbf{Z}\mathbf{q} - \mathbf{d}\|$ , is as close to the noise level as possible. Here, we define the noise level



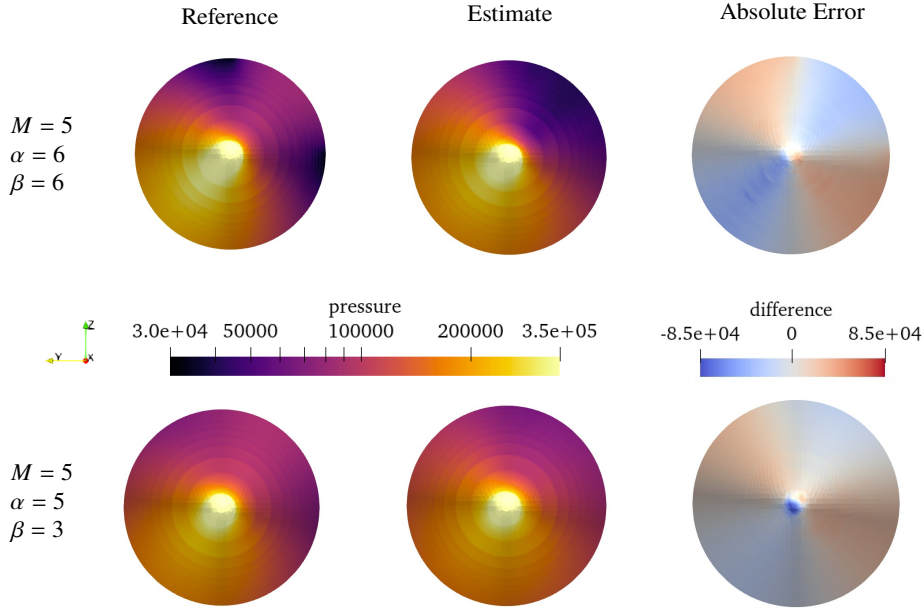
**Fig. 11** Minimum-norm surface pressure estimate for flight conditions  $M = 5, \alpha = 6, \beta = 6$ . (Top) Leeward view, and (bottom) windward view.

$\delta = \mathbb{E}[\|\eta\|] = \sigma\sqrt{n_d}$ . We observe that the data misfit is closest to the noise level when  $\lambda \approx 1.25 \times 10^6$ , which we use as the regularization parameter.

To illustrate the prediction errors, we consider two flight conditions:  $M = 5, \alpha = 6, \beta = 6$ , which is within the set  $P_1$ , as well as a testing flight condition where  $\mu_2 = M = 5, \alpha = 5, \beta = 3$ , which was not used as a snapshot for constructing the covariance. Figure 12 shows the reference surface pressure and estimated surface pressure for a synthetic measurement for each flight condition. The relative L2 errors for each flight condition were 22.3% and 16.4%, respectively. We observe that the predicted surface pressures are smooth except at the geometric discontinuity, since the prior covariance restricts the estimated solution to lie within the range space of the pressure snapshots, thus eliminating the high-frequency modes in the parameter-to-observable map. However, the leeward shock feature is not captured by the inverse solution because the data do not inform this feature. To further illustrate, Figure 13 shows a front view 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 perpendicular to the total angle of attack. This results in a very similar bending moment in comparison to the reference pressure field, which dominates the strain response. Specifically, the relative error in the pitch and yaw moment coefficients for noisy synthetic measurements were all within 4% and 12% for the flight condition in  $P_1$ , respectively. For the testing flight condition, the pitch and yaw moment coefficients were within 11% and 2%, respectively. In summary, the estimates of the pressure field remain as consistent with the data as possible under regularization; though the data do not inform certain features such as the leeward shock. This differs from the pressure field estimates in Case 1, where the shock was recovered, but slightly offset from the true location due to sensor noise.



**Fig. 12** Visualization of estimated surface pressure compared to the reference pressure. Left: Reference and estimate for  $M = 5, \alpha = 6, \beta = 6$ . Right: Reference and estimate for  $M = 5, \alpha = 5, \beta = 3$  (testing flight condition).



**Fig. 13 Front view of surface pressures. Top: Reference, estimate and absolute error for  $M = 5$ ,  $\alpha = 6$ ,  $\beta = 6$ . Bottom: Reference, estimate and absolute error for  $M = 5$ ,  $\alpha = 5$ ,  $\beta = 3$  (testing flight condition).**

#### IV. Conclusion

We have formulated a strain-based estimator for aerodynamic loads that is informed by a high-fidelity structural model and can be evaluated rapidly for real-time sensing applications. Given sensor noise statistics, we also provide uncertainty quantification and analysis. In the case where the number of unknown pressure quantities is smaller than the number of measurements (i.e.,  $n_q \leq n_d$ ), the variance of the estimated parameters is explicitly computed, and is dependent on 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 show that a data-driven prior can inform the estimator by restricting the solution to lie within the range space of the prior covariance. Consequently, the contributions of the higher frequency modes in the parameter-to-observable map are eliminated, and the data misfit corresponding to the regularization is minimized.

In this paper, we demonstrated the impact of sensor number and configuration on the performance and uncertainty of the estimator. Future work will consider the optimal sensor placement problem, where the goal is to minimize the number of sensors and determine their placement so that the uncertainty is reduced to within the desired tolerance. The uncertainty quantification in this paper can be used as a metric for solving this type of problem. Additionally, we seek to consider aerothermal heating effects, since the hypersonic environment causes additional thermal loads that will impact the strain response of the vehicle. Finally, experimental verification and validation studies for this work are planned, including both benchtop and wind tunnel testbeds.

#### Acknowledgments

This work was supported by AFOSR grant FA9550-21-1-0089 under the NASA University Leadership Initiative (ULI). 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 paper 2022-3717, AIAA Aviation 2022 Forum*, 2022. <https://doi.org/10.2514/6.2022-3717>.
- [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] 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.
- [8] 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.
- [9] Pena, F., Martins, B., and Richards, W. L., "Active In-flight Load Redistribution Utilizing Fiber-Optic Shape Sensing and Multiple Control Surfaces," NASA-TM 20180002147, *NASA Technical Memorandum*, 2018.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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 2021 Forum*, 2021, AIAA 2021-2809.
- [15] Bertsimas, D., and Dunn, J., "Optimal classification trees," *Machine Learning*, Vol. 106, No. 7, 2017, pp. 1039–1082.
- [16] Trefethen, L. N., and Bau, D., *Numerical Linear Algebra*, Society for Industrial and Applied Mathematics, 1997.
- [17] Pasilio, C. L., Sytsma, M. J., Neergaard, L., Witeof, Z., and Trolier, J. W., "Preliminary Aero-thermal Structural Simulation," *14th AIAA Aviation Technology, Integration, and Operations Conference*, 2014, AIAA 2014-2292.
- [18] 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.
- [19] 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.
- [20] 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. <https://doi.org/10.11588/ans.2015.100.20553>.
- [21] Logg, A., and Wells, G. N., "DOLFIN: Automated Finite Element Computing," *ACM Transactions on Mathematical Software*, Vol. 37, 2010. <https://doi.org/10.1145/1731022.1731030>.
- [22] 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.