

Neural operator-enabled aerodynamic load estimation for hypersonics

Julie Pham*

The University of Texas at Austin, Austin, TX, 78712

Patrick Blonigan†

Sandia National Laboratories, Albuquerque, NM, 87123

Thomas O'Leary-Roseberry‡

The Ohio State University, Columbus, OH, 43210

Omar Ghattas§ and Karen Willcox¶

The University of Texas at Austin, Austin, TX, 78712

This work expands on a strain-based aerodynamic sensing strategy for hypersonics to account for nonlinear temperature effects in real time. The sensing strategy uses sparse strain observations to infer the aerodynamic pressure loads, which is mathematically posed as a partial differential equation (PDE)-constrained inverse problem. In previous work, this inverse problem was shown to have a closed-form solution, where offline computation of the operations requiring the PDE solution was exploited to enable real-time evaluation speeds. In this work, the temperature effects preclude the offline pre-computation acceleration because the PDE operator is nonlinearly temperature dependent. To address this challenge, the recently developed neural matrix operator (NEMO) approach is employed to account for the temperature dependence. The NEMO method explicitly incorporates the physics structure of the governing equations, and thus preserves the availability of a closed-form inverse solution that can be computed rapidly onboard the vehicle. This work further considers the tasks of estimating the temperature field from sparse temperature measurements, and compensation for the thermal strain. The overall performance is demonstrated on the Initial Concept 3.X hypersonic vehicle. The results show strong approximation performance of NEMO and corresponding inverse solution accuracy, but further work is necessary to reduce errors in the thermal strain compensation.

I. Introduction

Aerodynamic sensing is critical for guidance and control of hypersonic vehicles. Particularly, the aerodynamic pressure loads and the corresponding force and moment coefficients are key information for flight control systems. Measurement of these aerodynamic quantities, however, poses a large challenge in the extreme aerothermal conditions of hypersonic flight. We seek to address this challenge through a strain-based sensing approach [1–3], which can enable rapid estimation of aerodynamic quantities of interest from internal strain sensors. This vehicle-as-a-sensor concept requires the rapid solution of a partial differential equation (PDE)-constrained inverse problem, where the PDE describes the governing physics relating the inversion parameters (pressure) and observables (strain). In Pham et al. [1], the solution to the inverse problem is shown to have a closed-form solution, and all operations involving the computationally expensive PDE solution can be precomputed offline, enabling real-time performance for online deployment.

In this paper, we extend the prior work to include temperature effects, which requires addressing two key additional considerations: (1) changes in the temperature-dependent, spatially-varying material properties in the PDE operator, and (2) compensation for the temperature-induced strain in the strain measurements. With these considerations, the inverse problem solution theoretically has an explicit solution of the same form as in [1], but the nonlinear temperature

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

†Principal Member of the Technical Staff, AIAA Senior Member.

‡Assistant Professor, Department of Mathematics.

§Professor, Walker Department of Mechanical Engineering.

¶Professor, Department of Aerospace Engineering and Engineering Mechanics, AIAA Fellow.

dependence no longer allows for offline pre-computations for rapid online evaluation. To achieve faster evaluation compared to traditional PDE solvers in nonlinear settings, approaches such as model reduction [4–6] and machine learning [7] are often used for design and optimization tasks. Here, we seek to address the nonlinear temperature dependence using machine learning in the form of neural operators. Neural operators approximate maps between function spaces for spatio-temporal fields, and have been shown to be scalable and accurate surrogates for forward problems [8–10]. These methods have also demonstrated success in solving inverse problems [11–14]. To address (1), due to the physics structure in this problem, we use the recently developed neural matrix operator (NEMO) [15] method that maps the temperature field T to the PDE parameter-to-observable (p2o) operator, denoted by $\mathcal{F}(T)$. Crucially for real-time performance, the NEMO formulation retains the availability of the closed-form solution to the inverse problem. NEMO considers reduced basis representations of the pressure and temperature fields to enable scalable training and fast online evaluation, with the additional benefit of providing regularization for the inverse problem. This approach contrasts with other machine learning approaches which directly learn the inverse map from measurement to parameter, as previously explored in [3, 16, 17]. While these approaches provide real-time evaluation speeds, they do not expose the underlying structure of the physics to allow for uncertainty quantification, interpretability, or recoverability analysis under ill-posed inverse problem conditions.

To address (2), we introduce an additional neural operator that estimates the thermal strain $\mathbf{z}(T)$ for a given temperature field. This allows us to compensate for the thermal strain in the observations, so that the inverse problem is performed on only the pressure-induced strain. Both (1) and (2) require knowledge of the temperature field; we also consider the task of reconstructing the temperature field from sparse temperature measurements. Figure 1 shows a block diagram of the PDE forward problem that represents the governing physics mapping pressure and temperature to strain, and the neural operator surrogates and dimension reduction strategies implemented in order to enable real-time inverse solutions. In the remainder of the paper, Section II describes the problem setup and methodology. Section III demonstrates the numerical results for a canonical hypersonic vehicle geometry. We summarize the discussion and future work in Section IV.

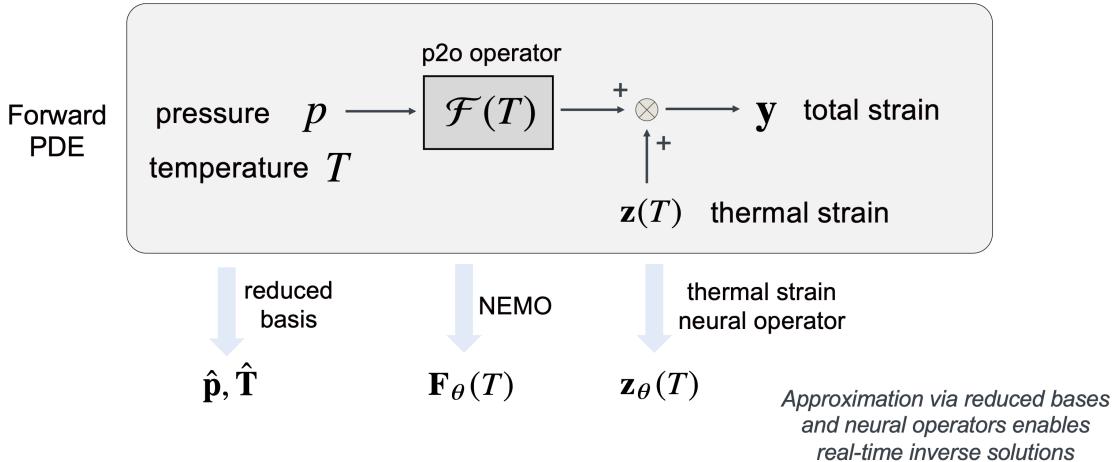


Fig. 1 Block diagram of forward problem components and corresponding approximations to enable real-time hypersonic aerodynamic sensing.

II. Methodology

In this section, we outline the methodology for the strain-based sensing approach with thermal effects. First, we define the forward PDE model describing the governing physics and the formulation of the neural operator learning problem using reduced basis representations of the pressure and temperature fields. We then illustrate how the neural operators are used to enable real-time inverse problem solutions on a reduced subspace.

A. Forward model definition

The relationship between the pressure, temperature, and strain is governed by the partial differential equation of elasticity. Consider the elasticity equation for the displacement $u(x)$ on domain $\Omega \subset \mathbb{R}^3$, with governing equations and boundary conditions given by

$$\left\{ \begin{array}{ll} -\nabla \cdot \sigma(u) = 0 & \text{in } \Omega \\ \sigma(u) \cdot n = 0 & \text{on } \partial\Omega_{\text{int}} \\ \sigma(u) \cdot n = t & \text{on } \partial\Omega_{\text{ext}} \\ u = 0 & \text{on } \partial\Omega_{\text{aft}} \end{array} \right. \quad (1)$$

where $\partial\Omega_{\text{int}}$, $\partial\Omega_{\text{ext}}$, $\partial\Omega_{\text{aft}}$ are the internal, external, and aft boundary surfaces, respectively, and $t = -p \cdot n$ where p is the aerodynamic surface pressure field and n is the outward-pointing unit normal. The strain-displacement relation is given by $\varepsilon = \frac{1}{2}(\nabla u + \nabla u^\top)$, where ε is the strain tensor. For spatially varying temperature field T , the constitutive equation for the stress tensor σ is given by

$$\sigma = \lambda \text{tr}(\varepsilon)I + 2\mu\varepsilon - \underbrace{\alpha(3\lambda + 2\mu)(T - T_0)I}_{\text{thermal load}} \quad (2)$$

where $\lambda(T)$, $\mu(T)$, $\alpha(T)$ are temperature-dependent material properties (Lamé parameters, thermal expansion coefficient), and T_0 is a reference temperature. From the governing equations, and defining $A(T)$ as the temperature-dependent PDE operator, the forward problem can be written as

$$A(T)u = Cp + f(T) \quad (3a)$$

$$\mathbf{y} = Bu \quad (3b)$$

where \mathbf{y} are the discrete strain observations at the sensor locations, B is the observation operator that extracts the observable strains from the state u , C maps the pressure to forces, and $f(T)$ is the thermal load. We can write the full parameter-to-observable (p2o) map as

$$\mathbf{y} = BA(T)^{-1}(Cp + f(T)). \quad (4)$$

We define $\mathcal{F}(T) = BA(T)^{-1}C$ as the p2o operator, which maps the pressure field to the strains observed at the sensor locations. The p2o operator is expensive to apply, since it requires the solution of the expensive PDE (represented mathematically by $A(T)^{-1}$). The p2o operator is also nonlinearly temperature-dependent, and thus cannot be pre-computed offline as in our prior work [1]. Consequently, we seek surrogate models in the form of neural operators to be used by the inverse problem solution for rapid, real-time tractability. In the next section, we outline the neural operator learning problem, which utilizes reduced basis representations of the functions p and T with the NEMO architecture to explicitly incorporate the physics structure of the problem. As we will show in Section II.D, the physics-structure preservation enables a closed-form solution that can be evaluated rapidly in real-time.

B. Reduced basis neural operators using governing physics structure

We consider orthonormal reduced basis representations of the functions $p \in \mathcal{P}$ and $T \in \mathcal{T}$, where \mathcal{P}, \mathcal{T} are real-valued Hilbert spaces. The approximations are given by

$$p \approx \bar{p} + \sum_{i=1}^{r_p} \hat{p}_i \phi_i, \quad T \approx \bar{T} + \sum_{i=1}^{r_T} \hat{T}_i \psi_i \quad (5)$$

where \hat{p}_i, \hat{T}_i are the i -th modal coefficients, ϕ_i, ψ_i are the i -th basis functions, \bar{p}, \bar{T} are the affine shifts, and r_p, r_T are the number of retained modes, for the pressure and temperature fields respectively. We denote the vector of reduced coordinates for each function as $\hat{\mathbf{p}}, \hat{\mathbf{T}}$. These are computed from the original functions as $\hat{\mathbf{p}} = \Phi_r^*(p - \bar{p})$ and $\hat{\mathbf{T}} = \Psi_r^*(T - \bar{T})$, where Φ_r is the vector of pressure basis functions ϕ_i for $i = 1, \dots, r_p$, Ψ_r is the vector of temperature basis functions ψ_i for $i = 1, \dots, r_T$, and $*$ denotes the adjoint. Substituting the reduced coordinate representation of the

pressure field, the expanded equation (4) becomes

$$\mathbf{y} = BA(T)^{-1} \left(C \left(\bar{p} + \sum_i^r \hat{p}_i \phi_i \right) + f(T) \right) \quad (6a)$$

$$\begin{aligned} &= BA(T)^{-1} C \Phi_r \hat{\mathbf{p}} \\ &\quad + BA(T)^{-1} f(T) \\ &\quad + BA(T)^{-1} C \bar{p} \end{aligned} \quad (6b)$$

Observing the first term in Equation (6b), we define the reduced p2o (rp2o) operator $BA(T)^{-1} C \Phi_r = \mathcal{F}(T) \Phi_r$, which maps the reduced coordinates of pressure to the pressure-induced (mechanical) strain. Here, we employ a NEMO approximation of the rp2o operator. Due to the reduced basis representation of the pressure field, the rp2o operator is a matrix of size $n_y \times r_p$, where n_y is the number of strain sensors. Since these dimensions are moderate, the NEMO method is able to directly output the rp2o matrix. In other words, NEMO represents the mapping $T \rightarrow \mathcal{F}(T) \Phi_r$. Direct output of the rp2o operator preserves the linear relationship between the strain and the reduced pressure coordinates. We will see that this is instrumental for real-time inverse problem solutions, which is further detailed in Section II.D. For scalability of method, the temperature is also represented in its reduced coordinate form as the input to the NEMO network. We will refer to the NEMO approximation of the rp2o operator as $\mathbf{F}_\theta(T)$, where θ indicates the neural network weight parameters. We train NEMO with objective function given by

$$\min_{\theta} \frac{1}{N} \sum_{j=1}^N \|\mathcal{F}(T_j) \Phi_r - \mathbf{F}_\theta(T_j)\|_F^2 \quad (7)$$

where N is the total number of training samples, and F denotes the Frobenius norm.

We must also consider the remaining strain component terms in (6b). The thermal strain component, which we denote as $\mathbf{z} = BA(T)^{-1} f(T)$, is also temperature dependent. To evaluate this term rapidly, we seek to learn an additional neural operator which approximates the map $T \rightarrow \mathbf{z}(T)$. The thermal strain neural operator (NO) output is similarly denoted by $\mathbf{z}_\theta(T)$. To train the thermal strain model, we utilize the following objective:

$$\min_{\theta} \frac{1}{N} \sum_{j=1}^N \|\mathbf{z}(T_j) - \mathbf{z}_\theta(T_j)\|_2^2. \quad (8)$$

The magnitude of the remaining component $BA(T)^{-1} C \bar{p}$ is small, and the temperature-dependent variation of this component is smaller than the considered noise levels. Therefore, we assume it to be a constant, $\bar{\mathbf{y}} = BA(\bar{T})^{-1} C \bar{p}$.

C. Constructing reduced bases

Dimension reduction is the foundational methodology that enables scalability of the neural operator learning problems. Defining appropriate reduced bases is also critical to the performance of the inverse problem. In this work, we consider linear subspaces constructed from data by utilizing additional physics simulation to obtain representative fields of the pressure and temperature. This approach allows us to construct reduced bases that are informed by physics models and characteristic of the flight condition ranges of interest, as opposed to using more general, smoothing eigenfunctions. Specifically, we use computational fluid dynamics (CFD) and transient thermal simulations to produce snapshots of the pressure and temperature fields over a set of flight conditions of interest. These data are used to construct appropriate subspaces via the proper orthogonal decomposition (POD). We compute samples for a function p from a distribution ρ that is characterized by the selected flight condition range. Then, we compute the reduced basis modes (eigenfunctions) by solving the eigenvalue problem given by

$$\mathbb{E}_{p \sim \rho} [\langle p - \bar{p}, \phi_j \rangle_{\mathcal{P}} (p - \bar{p})] = \lambda_j \phi_j \quad \langle \phi_i, \phi_j \rangle_{\mathcal{P}} = \delta_{ij}. \quad (9)$$

where λ_j denotes the eigenvalue corresponding to the j -th eigenfunction ϕ_j , δ_{ij} is the Kronecker delta, and $\langle \cdot, \cdot \rangle_{\mathcal{P}}$ denotes the inner product on \mathcal{P} . We can choose the number of basis functions r for truncation based on the eigenvalue decay and reconstruction accuracy requirements.

D. Real-time inverse problem solution with neural operators

The inverse problem is posed in the reduced subspace as a least-squares problem with regularization. We define the measured strain as $\tilde{\mathbf{y}} = \mathbf{y} + \boldsymbol{\eta}$, where $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma}_n)$ is the Gaussian noise model with zero mean and covariance $\boldsymbol{\Gamma}_n$. The inverse problem in the reduced subspace with the neural operator approximations is given by

$$\begin{aligned} \min_{\hat{\mathbf{p}}} \|\tilde{\mathbf{y}} - \mathbf{y}\|_{\boldsymbol{\Gamma}_n^{-1}}^2 + \gamma \|\hat{\mathbf{p}}\|_2^2 \\ \text{s.t. } \mathbf{y} = \mathbf{F}_\theta(T)\hat{\mathbf{p}} + \bar{\mathbf{y}} + \mathbf{z}_\theta(T). \end{aligned} \quad (10)$$

The solution to the subspace inverse problem is given by

$$\mathbf{p}_* = (\mathbf{F}_\theta^\top(T)\boldsymbol{\Gamma}_n^{-1}\mathbf{F}_\theta(T) + \gamma\mathbf{I})^{-1}\mathbf{F}_\theta^\top(T)\boldsymbol{\Gamma}_n^{-1}(\tilde{\mathbf{y}} - \bar{\mathbf{y}} - \mathbf{z}_\theta(T)). \quad (11)$$

By using \mathbf{F}_θ in place of the true rp2o operator, we are able to directly evaluate the closed-form solution without solving the expensive PDE online. The main contributors to the computational cost are the evaluation of NEMO and the thermal strain neural operator, and the matrix inverse in (11). Using NEMO instead of computing the true rp2o using the PDE results in over three orders of magnitude speedup, as shown in [15]. To be considered real-time, the solution must be able to be evaluated at the typical onboard sensor measurement frequency, which is approximately 100 Hz. Since we solve the problem on a reduced subspace, the matrix to be inverted is of size $r_p \times r_p$. Since r_p can typically be chosen to be small, this enables the solution to be evaluated rapidly. Further, posing the problem in the reduced subspace is advantageous because it provides regularization for the solution, which is critical since the inverse problem is ill-posed. The POD subspace also provides characteristics that are informed by the CFD models. Similar to other regularization approaches, such as truncated SVD, the subspace approach eliminates the higher spatial frequency modes which are not recoverable due to ill-posedness [18].

In the solution in Eq. (11), we compensate for the thermal strain by subtracting \mathbf{z}_θ from the measured strain $\tilde{\mathbf{y}}$. We note that the errors in estimating the thermal strain corrupts the remaining signal (the pressure-induced mechanical strain) from which the aerodynamic pressure must be inferred. Controlling this error is therefore critical to the inverse problem performance.

III. Numerical Results

In this section, we present the model vehicle and numerical results. First, we describe the numerical testbed setup and modeling details in Section III.A and III.B. In Sections III.C and III.D, we show the errors associated with estimating the reduced temperature coordinates from measurements, and assess neural operator errors with the exact and estimated reduced temperature inputs. Then, we assess the impacts of these errors on inverse problem solution quality combined with NEMO in Section III.E, and the thermal strain neural operator in Section III.F.

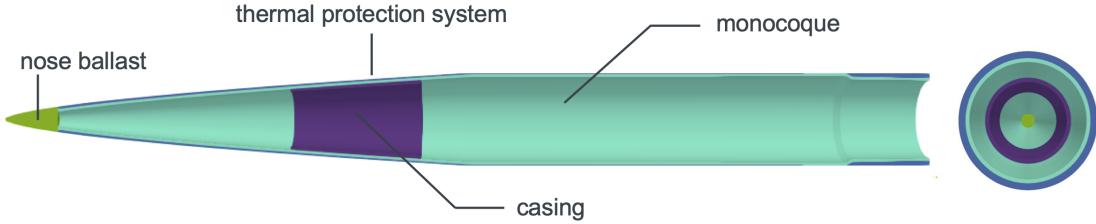


Fig. 2 IC3X components.

A. IC3X modeling

We demonstrate the methodology on a conceptual hypersonic vehicle, the Initial Concept 3.X (IC3X) [19]. We consider a modified version of the IC3X, which has the same outer mold line but adapted structural components as shown in Figure 2. The values for the material properties, including the thermal conductivity, specific heat capacity, Young's modulus and thermal expansion coefficients, are taken from Section 5.2.2 in [20]. For the structural model, the

thermal expansion coefficient for all components is assumed constant. The Young's modulus of the nose ballast, casing, and thermal protection system (TPS) is constant with temperature, while the Young's modulus of the monocoque is temperature dependent for the structural model.

We solve the PDE in Eq. (1) with a structural finite element model using FEniCS [21] to obtain the displacement u for the pressure and temperature at a given flight condition. The surface pressure fields p are obtained by solving the steady, inviscid Euler equations using CART3D [22] at a given flight condition. The temperature fields T (over the full domain Ω) are obtained by solving a transient thermal problem (heat equation) using SIERRA/Aria [23] with an aerothermal heating boundary condition based on thin film theory. For each temperature field, the reference temperature T_0 is taken to be the temperature at a fixed point on the internal structure of the vehicle. The aeroheating fields needed by the thermal solver are obtained using the Sandia Parallel Aerodynamics and Reentry Code (SPARC) [24] at a steady flight condition. The flow over the vehicle is assumed to be fully turbulent and modeled using a Reynolds-Averaged Navier-Stokes solver with the Spalart-Allmaras turbulence model. Since the aeroheating boundary conditions are not expected to be sensitive to wall temperature for flight conditions considered, a constant wall temperature of 350K is used for the SPARC simulation.

We consider a set of strain and temperature sensors at axial locations (x -positions) with eight circumferential locations at each x -position, as shown in Figure 3. For the strain sensors, we consider measurement of strain in both the axial and circumferential directions on the interior surface of the monocoque at each location (two sensors). The total number of strain sensors is 128. For the temperature sensors, we consider interior surface measurements at all axial locations, and outer monocoque temperature measurements (below the thermal protection layer) at two of the axial locations. Two additional temperature sensors are placed at the nose of the vehicle. The total number of temperature sensors is 50.

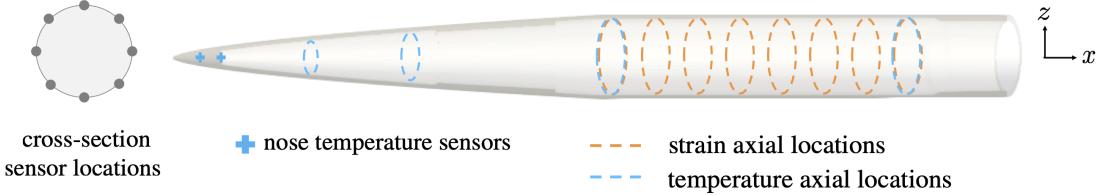


Fig. 3 Pressure and temperature sensor locations.

B. Reduced bases

We compute reduced bases for the pressure and temperature fields by collecting representative samples over a range of flight conditions and computing the POD basis. For the pressure fields, we consider a range of $[-10, 10]$ degrees for the angle of attack α and sideslip angle β , and a Mach number range of $[5, 7]$ at an altitude of 20 kilometers. The flight conditions are grid sampled with angle spacing of two degrees and Mach number spacing of one, resulting in 363 pressure samples. The temperature fields are taken from a time span of $[20, 80]$ seconds in the thermal simulations which each use a constant flight condition. For the thermal simulations, we consider angle of attack and sideslip ranges $[-6, 6]$ degrees over the same Mach number and altitude ranges as the pressure fields. Using these data, we compute the POD basis for both the pressure and temperature fields as in Equation (9) using hippyflow[25]. Figure 4 shows the eigenvalue decay for the temperature and pressure modes. Both functions show rapid decay, indicating they can be well-represented in a low-dimensional basis. We choose $r_p = 20$ for pressure field and $r_t = 7$ for the temperature field, resulting in a decrease of more than two orders of magnitude from the first to the r -th eigenvalue.

C. Temperature reconstruction

The thermal state of the vehicle is inferred from point measurements of the temperature, denoted by $\tilde{\mathbf{T}}$. We employ a gappy POD approach as a least-squares problem with regularization, given by

$$\mathbf{T}_* = \min_{\tilde{\mathbf{T}}} \|\tilde{\mathbf{T}} - S(\Psi_r \hat{\mathbf{T}} + \bar{\mathbf{T}})\|_2^2 + \zeta \|\hat{\mathbf{T}}\|_2^2 \quad (12)$$

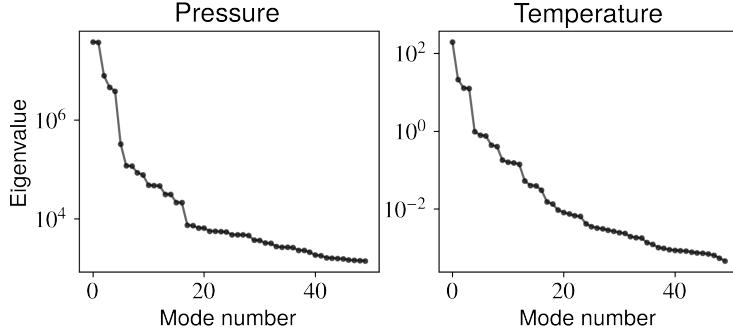


Fig. 4 Eigenvalue decay for the pressure and temperature modes.

where S is the selection (observation) operator for the temperature sensors, and ζ is a regularization parameter. This minimizes the error between the reconstructed temperature field at the sensor locations and the measurements. The regularization is added for robustness to noise in the temperature measurements. To produce synthetic temperature measurements, we use an additive Gaussian noise model defined by $\eta_i \sim \mathcal{N}(0, \sigma^2)$, assuming independent noise for all temperature sensors, and we define 3σ to be 1.5% of the magnitude of the temperature for each measurement, where σ is the standard deviation. We then solve Eq. (12) via the normal equations to obtain estimate \mathbf{T}_\star , which we refer to as the estimated input to NEMO and the thermal strain neural operator models.

D. Surrogate modeling with neural operators

For both NEMO and the thermal strain neural operator, the models consist of fully-connected feed-forward neural networks with ReLU activation functions, and are trained using PyTorch using the Adam optimizer with an initial learning rate of 5e-4 and batch size 64. The temperature input dimension is $r_T = 7$, the dimension of the reduced temperature basis. For the NEMO operator, the network output size is $n_y \times r_p = 128 \times 20$. The model is trained for 2000 epochs over 500 training samples, and 100 validation samples. For the thermal strain neural operator, the output size is the number of strain sensors, $n_y = 128$. The model is trained for 15000 epochs with 650 training samples, and 100 validation samples.

Here, we demonstrate the test accuracy of the surrogate models. The results are shown on a test set consisting of 100 unseen temperature field samples. To provide a reference for the performance benefits of the NEMO, we consider a reference constant (temperature-independent) rp2o operator. We choose this reference operator to be the element-wise mean over the true rp2o operators from the test set, which we refer to as the mean rp2o operator. Figure 5 shows the relative errors of the models, using inputs of both the exact and estimated (from solving (12)) reduced temperature coordinates. For the reduced p2o operators, we show the relative Frobenius errors of NEMO and the mean rp2o compared to the true rp2o operator. We observe that the NEMO relative errors are small, indicating a good approximation of the true rp2o operator. The mean rp2o operator errors are small when the mean is a good representation for the true temperature field. However, when the mean is not close to the truth, the errors are larger compared to NEMO. For the thermal strains, we show the relative L2 errors of the neural operator outputs. We observe low relative errors in the thermal strain neural operator for exact inputs, with most test points below 2% error. However, the thermal strain neural operator is sensitive to errors due to the estimated inputs, with errors reaching above 10% in the worst cases.

E. Inverse problem with NEMO

In this section, we assess the quality of inverse problem solutions using NEMO. Here, we assume exact knowledge of the thermal strain (i.e., no errors in $\mathbf{z}(T)$). This is equivalent to measuring the pressure-induced mechanical strain in isolation from the thermal strain. We produce synthetic noisy observations with an additive Gaussian noise model defined by $\eta_i \sim \mathcal{N}(0, \sigma^2)$, assuming independent, identically distributed noise for all measurements, where σ is the standard deviation. The regularization parameter γ is chosen using the L-curve criterion. The inverse problem test set consists of all combinations of the 100 test temperature samples and 30 pressure samples, resulting in 3000 total test samples for the inverse problems.

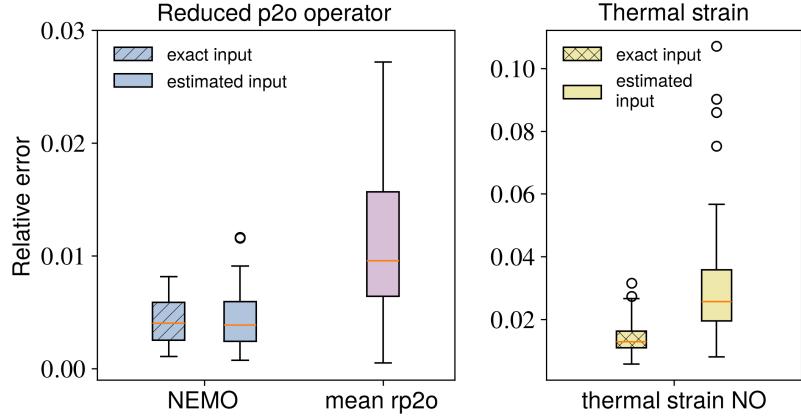


Fig. 5 (Left) Relative Frobenius error of NEMO and mean rp2o operators. (Right) Relative L2 error of thermal strain neural operator outputs.

We compute the subspace inverse problem solution \mathbf{p}_* using Equation (11). We reconstruct the surface pressure as $p = \bar{p} + \Phi_r \mathbf{p}_*$. Figure 6 shows the relative error in the pressure field reconstruction from the inverse problem solution when using the true rp2o, NEMO with estimated inputs, and the mean rp2o operators. We define two different noise levels, choosing three standard deviations 3σ of the noise distribution to be 1.5% and 5% of the mean strain response over all sensors and samples in the test set. In the low noise case (left), we use different regularization parameters γ for the three cases, since NEMO and the mean rp2o introduce approximation errors that require additional regularization. We observe smaller pressure errors with NEMO compared to the mean rp2o. The mean rp2o reconstruction errors again are large when the temperature field is not well-represented by the mean. The NEMO reconstruction errors result from both the estimated input errors and the NEMO approximation errors. In the high noise case (right), the same γ is used for all three cases. Here, all three reduced operators achieve similar performance. This indicates that when the noise level is high, the reduced operator temperature-dependent effects on the strain response are within the noise levels. This mitigates the benefits of accounting for this temperature dependence using NEMO.

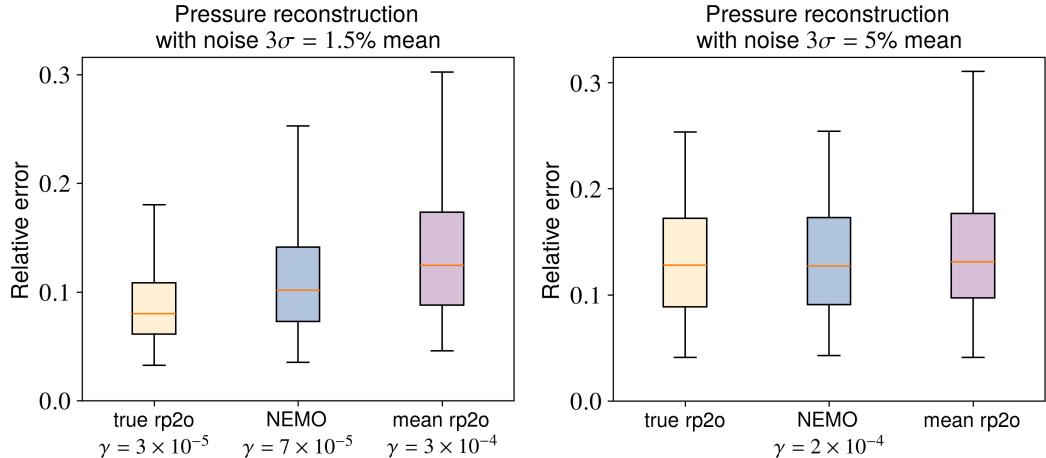


Fig. 6 Relative errors in the pressure inverse problem solution for (left) low noise level and (right) high noise level using the true rp2o, NEMO, and mean rp2o operators on the test set.

From the reconstructed pressure fields, we also compute the estimated aerodynamic force and moment coefficients in the body frame, which are computed by integrating the forces due to the reconstructed pressure over the vehicle. Figure 7 shows the coefficient normalized errors on the test set for the low noise case, where the normalized error is

computed by

$$\epsilon_C = \frac{C_* - C}{\max\{0.2, |C|\}} \quad \epsilon_M = \frac{M_* - M}{\max\{0.4, |M|\}}. \quad (13)$$

where C denotes a force coefficient and M denotes a moment coefficient. This metric is a relative error that avoids division by small true values of the aerodynamic coefficients, with the sign indicating over- or under- prediction of the coefficient. We observe that the errors in the aerodynamic coefficients are smaller than their full-field pressure reconstruction counterparts from Figure 6, since they are integrated (smoothed) quantities. We also observe that the NEMO coefficient errors are smaller compared to the mean rp2o.

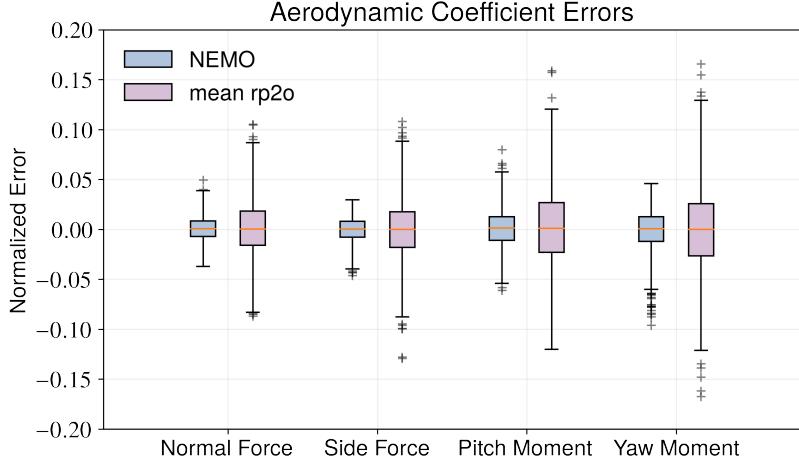


Fig. 7 Normalized errors in the estimated aerodynamic coefficients on the test set with known thermal strain.

We visualize an example pressure reconstruction using NEMO in the low noise case in Figure 8. The flight conditions are Mach 6, with angle of attack 2 degrees and sideslip angle 10 degrees. This reconstruction has a relative error of 5.6%. We observe that the NEMO inverse problem smooths the discontinuities in the true pressure, since they cannot be resolved through the inverse problem. This is a result of the problem ill-posedness, and the corresponding required regularization to account for the noise and sparsity of the measurements.

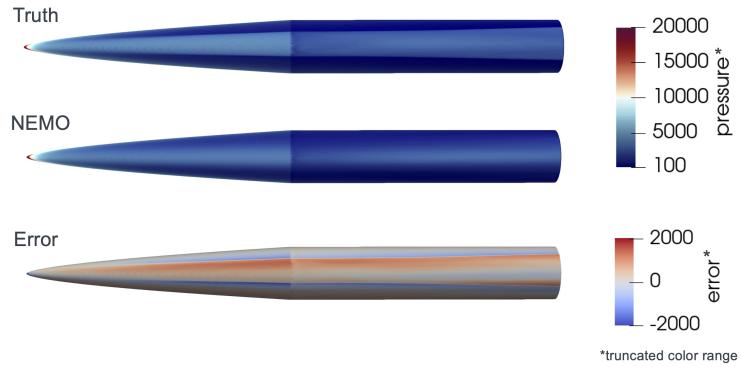


Fig. 8 Example inverse problem aerodynamic surface pressure reconstruction with NEMO.

F. Inverse problem with NEMO and thermal strain neural operator

We now consider the inverse problem solution quality using the thermal strain neural operator for thermal compensation along with NEMO. Figure 9 shows the pressure reconstruction errors for the inverse problem using thermal strain neural operator with the exact and estimated input for both NEMO and the mean rp2o. First, the errors in the

pressure reconstruction increase when considering the estimated thermal strain from the neural operator, compared to the errors in Figure 6, where the thermal strain was known. This arises from the fact that the thermal strain magnitude is often larger than the pressure-induced mechanical strain. Consequently, with any errors in estimating the thermal strain, the remaining signal (the mechanical strain) computed by $\hat{\mathbf{y}} - \mathbf{z}_\theta$ in Equation (11) has a low signal-to-noise ratio. Therefore, even in the exact input case, where the thermal strain errors are around 2% (as shown in Figure 5), we observe that more regularization (larger γ) is needed compared to either noise scenario in Figure 6. The reconstruction errors increase further in the estimated input case, emphasizing the effects of the thermal strain neural operator sensitivity to input errors. Finally, there is not a discernible benefit of using NEMO compared to the mean $rp2o$, since the errors introduced by the thermal strain estimate are dominant over the errors due to neglecting the material temperature-dependence.

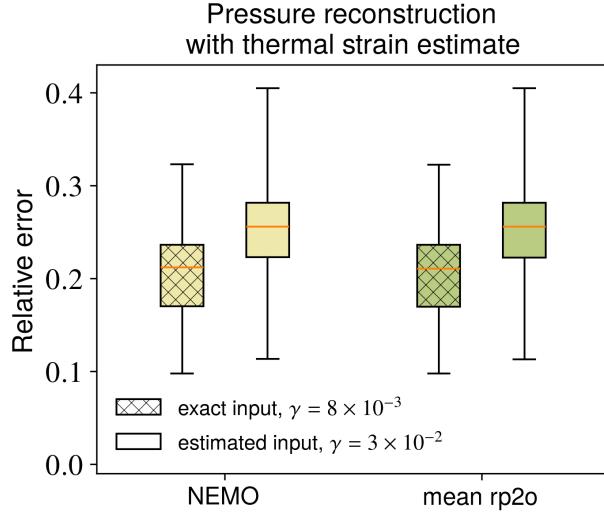


Fig. 9 Relative errors in the pressure inverse problem solution on the test set with NEMO and the mean $rp2o$, and the thermal strain neural operator with exact and estimated inputs.

Figure 10 shows the aerodynamic coefficient normalized errors using NEMO with the thermal strain neural operator for compensation, alongside NEMO with the exact thermal strain as a reference. As expected, the errors are increased when using the thermal strain neural operator compared to the exact known thermal strain.

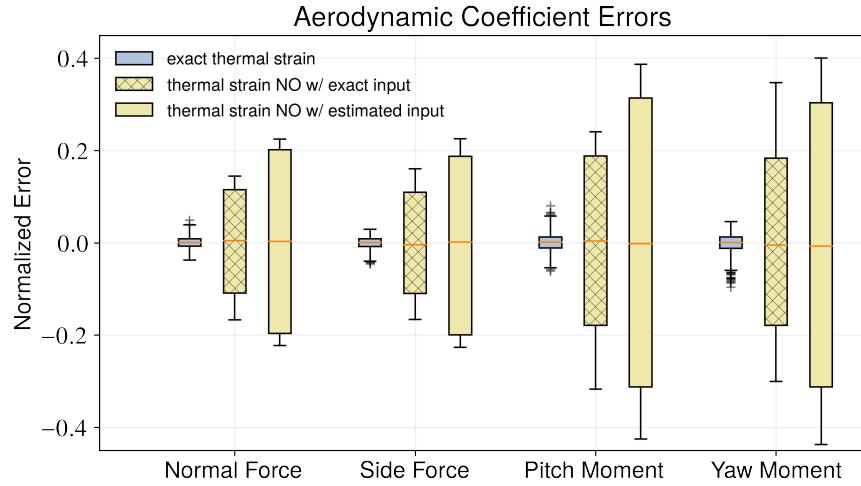


Fig. 10 Normalized errors in the estimated aerodynamic coefficients on the test set with NEMO and thermal strain compensation.

These results show that more accurate compensation for the thermal strain is necessary for the success of a strain-based sensing strategy for hypersonic vehicles. Similar conclusions were drawn in the experimental setting on a benchtop testbed setup, where the inverse problem performance in estimating point loads was impacted by the effects of thermal compensation [26]. Future work will benefit from improved modeling techniques to estimate the thermal strain from temperature measurements, as well as the consideration of filtering techniques in the time or frequency domain [27]. Another possible direction is sensor technology with built-in temperature compensation to enable accurate isolation of the mechanical strain in these settings. Further, the vehicle structure itself may be specifically designed to increase the mechanical strain signal and limit the impacts of thermal effects.

IV. Conclusion

In this work, a real-time strain-based aerodynamic sensing approach for hypersonics is extended to account for nonlinear temperature effects via the introduction of neural operators. The two main challenges due to temperature effects addressed in this work are the temperature-dependent material properties that appear in the PDE parameter-to-observable operator, and the thermally-induced strain that must be compensated for in the inverse problem. The recently developed neural matrix operator (NEMO) is used to account for the temperature-dependence of the parameter-to-observable operator. This method preserves the physics structure of the problem, allowing for a rapid closed-form inverse solution. An additional neural operator is introduced to compensate for the thermal strain in the inverse problem. The results show the inverse problem quality for full surface pressure reconstructions, as well as the aerodynamic coefficients. The effects of different sources of errors are demonstrated, including the thermal state estimation from temperature measurements, the resultant NEMO and thermal strain neural operator accuracy, and the corresponding impact on the inverse problem performance. NEMO demonstrates strong performance in both approximation accuracy and the inverse problem. However, the thermal strain neural operator is susceptible to input noise, and the thermal strain compensation methods result in lower quality inverse problem solutions due to the magnitude of the thermal strain.

Future work should focus particular attention to the thermal strain compensation. Isolating the mechanical strain from the thermal strain while preserving a reasonable signal-to-noise ratio for the inverse problem is a significant challenge. Improved modeling approaches or sensor technologies may improve the performance observed in this work. Beyond the quasi-static setting considered in this work, signal processing and filtering techniques may be able to improve the thermal strain characterization by exploiting the timescale difference between thermal and aerodynamic state changes of the vehicle. Further, we note that there are still many open modeling questions to address in future work. Specifically, the monocoque thickness, control surfaces, volumetric finite element modeling, temperature strain locations and fixed aft boundary conditions may be chosen to better reflect a specific real-world hypersonic system. These modeling considerations should be addressed along with extensive model calibration.

Acknowledgements This work was supported by AFOSR grant FA9550-21-1-0089 under the NASA University Leadership Initiative (ULI) and the Multifaceted Mathematics for Predictive Digital Twins (M2DT) grant DE-SC002317 under the U.S. Department of Energy's Office of Advanced Scientific Computing Research. This material is also based on work supported by the National Science Foundation (NSF) Graduate Research Fellowship under grant DGE 2137420, and the NSF grant 2313033 under the Office of Advanced Cyberinfrastructure. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing computational resources that have contributed to the results in this paper. The authors would like to thank Sean Babiniec for his assistance setting up the thermal model of the IC3X flight vehicle in Sierra/Aria, and the members of the FAST ULI team for helpful discussions.

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC (NTESS), a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration (DOE/NNSA) under contract DE-NA0003525. This written work is authored by an employee of NTESS. The employee, not NTESS, owns the right, title and interest in and to the written work and is responsible for its contents. Any subjective views or opinions that might be expressed in the written work do not necessarily represent the views of the U.S. Government. The publisher acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this written work or allow others to do so, for U.S. Government purposes. The DOE will provide public access to results of federally sponsored research in accordance with the DOE Public Access Plan. SAND2025-15036C

References

- [1] Pham, J., Ghattas, O., Clemens, N., and Willcox, K., “Real-Time Aerodynamic Load Estimation for Hypersonics via Strain-Based Inverse Maps,” *AIAA Journal*, Vol. 63, No. 1, 2025, pp. 91–101. <https://doi.org/10.2514/1.J064375>.
- [2] Panigrahi, A., Villa, B. E. D., Blocher, B. L., Sirohi, J., and Clemens, N. T., “Data-Driven Reconstruction of Distributed Force on a Hypersonic Wind Tunnel Model Using Measured Strain,” *AIAA SciTech Forum*, 2025, paper AIAA 2025-1209. <https://doi.org/10.2514/6.2025-1209>.
- [3] 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. <https://doi.org/10.2514/6.2024-1047>.
- [4] 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. <https://doi.org/10.2514/6.2010-7928>.
- [5] Crowell, A. R., and McNamara, J. J., “Model Reduction of Computational Aerothermodynamics for Hypersonic Aerothermoelasticity,” *AIAA Journal*, Vol. 50, No. 1, 2012, pp. 74–84.
- [6] 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.
- [7] Li, J., Du, X., and Martins, J. R., “Machine Learning in Aerodynamic Shape Optimization,” *Progress in Aerospace Sciences*, Vol. 134, 2022, p. 100849.
- [8] Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A., “Fourier Neural Operator for Parameteric Partial Differential Equations,” *International Conference on Learning Representations*, 2021.
- [9] Lu, L., Jin, P., Pang, G., Zhang, Z., and Karniadakis, G. E., “Learning Nonlinear Operators via DeepONet Based on the Universal Approximation Theorem of Operators,” *Nature Machine Intelligence*, Vol. 3, No. 3, 2021, p. 218–229. <https://doi.org/10.1038/s42256-021-00302-5>.
- [10] O’Leary-Roseberry, T., Chen, P., Villa, U., and Ghattas, O., “Derivative-Informed Neural Operator: An Efficient Framework for High-Dimensional Parametric Derivative Learning,” , 2023.
- [11] Kaltenbach, S., Perdikaris, P., and Koutsourelakis, P.-S., “Semi-Supervised Invertible Neural Operators for Bayesian Inverse Problems,” *Computational Mechanics*, Vol. 72, 2023, p. 451–470.
- [12] Molinaro, R., Yang, Y., Engquist, B., and Mishra, S., “Neural Inverse Operators for Solving PDE Inverse Problems,” *The 40th International Conference on Machine Learning*, 2023, pp. 25105 – 25139.
- [13] Wang, T., and Wang, C., “Latent Neural Operator for Solving Forward and Inverse PDE Problems,” *The 38th Annual Conference on Neural Information Processing Systems*, 2024, pp. 33085 – 33107.
- [14] Cao, L., O’Leary-Roseberry, T., and Ghattas, O., “Derivative-Informed Neural Operator Acceleration of Geometric MCMC for Infinite-Dimensional Bayesian Inverse Problems,” *Journal of Machine Learning Research*, Vol. 26, No. 78, 2025, pp. 1–68.
- [15] Pham, J., O’Leary-Roseberry, T., Ghattas, O., and Willcox, K., “Real-Time Inverse Solutions for Digital Twins via Neural Matrix Operators,” *Oden Institute Report, The University of Texas at Austin*, 2025. URL <https://oden.utexas.edu/research/publications?searchNeedle=Real-time%20inverse%20solutions%20for%20digital%20twins%20via%20neural%20matrix>.
- [16] 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. <https://doi.org/10.2514/6.2022-3717>.
- [17] Meinicke, A. C., Cesnik, C. E., Blocher, B., Panigrahi, A., Sirohi, J., and Clemens, N., “Experimental Validation of Strain-Load Neural Network Model on a Slender Hypersonic Vehicle,” *International Forum on Aeroelasticity and Structural Dynamics*, 2024, paper IFASD-2024-185.
- [18] 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. <https://doi.org/10.1017/S0962492921000064>.
- [19] 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.

- [20] Klock, R. J., “Efficient Numerical Simulation of Aerothermoelastic Hypersonic Vehicles,” Ph.D. thesis, University of Michigan, 2017.
- [21] 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.
- [22] 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 2024.
- [23] Clausen, J., Brunini, V., Collins, L., Knaus, R., Kucala, A., Lin, S., Matula, N., Moser, D., Phillips, M., Ranseggnola, T., Subia, S., Vasyliv, Y., Voskuilen, T., Smith, T., Tejan-Kella, Z., and Lamb, J., “SIERRA Multimechanics Module: Aria User Manual - Version 5.12.” Tech. Rep. SAND2023-12589, Sandia National Laboratories, 2023.
- [24] Howard, M., Bradley, A., Bova, S. W., Overfelt, J., Wagnild, R., Dinzl, D., Hoemmen, M., and Klinvex, A., “Towards Performance Portability in a Compressible CFD Code,” *AIAA Aviation Forum*, 2017, paper 2017-4407. <https://doi.org/10.2514/6.2017-4407>.
- [25] O’Leary-Roseberry, T., Luo, D., and Villa, U., “hippyflow: Dimension Reduced Surrogate Construction for Parametric PDE Maps in Python,” <https://github.com/hippylib/hippyflow>, Accessed 2025.
- [26] Meinicke, A. C., Cesnik, C. E., Blocher, B. L., and Clemens, N. T., “Experimental Validation of Thermoelastic Neural Network-Based Inverse Model on a Slender Hypersonic Vehicle,” *AIAA Aviation Forum and Ascend*, 2025, paper 2025-3400. <https://doi.org/10.2514/6.2025-3400>.
- [27] Morelli, E. A., “Real-Time Aerodynamic Parameter Estimation Without Air Flow Angle Measurements,” *Journal of Aircraft*, Vol. 49, No. 4, 2012, pp. 1064–1074. <https://doi.org/10.2514/1.C031568>.