SAND2021-15238C

Reduced Order Modeling of Hypersonic Aerodynamics with Grid Tailoring

David S. Ching* and Patrick J. Blonigan†
Sandia National Laboratories, Livermore, California, 94550

Francesco Rizzi
NexGen Analytics, Sheridan, WY, 82801, USA

Jeffrey A. Fike‡
Sandia National Laboratories, Albuquerque, New Mexico, 87185-0828

High-fidelity hypersonic aerodynamic simulations require extensive computational resources, hindering their usage in hypersonic vehicle design and uncertainty quantification. Projection-based reduced-order models (ROMs) are a computationally cheaper alternative to full-order simulations that can provide major speedup with marginal loss of accuracy when solving many-query problems such as design optimization and uncertainty propagation. However, ROMs can present robustness and convergence issues, especially when trained over large ranges of input parameters and/or with few training samples.

This paper presents the application of several different residual minimization-based ROMs to hypersonic flows around flight vehicles using less training data than in previous work. The ROM demonstrations are accompanied by a comparison to fully data-driven approaches including kriging and radial basis function interpolation. Results are presented for three test cases including one three-dimensional flight vehicle. We show that registration-based ROMs trained on grid-tailored solutions can compute quantities of interest more accurately than data driven approaches for a given sparse training set. We also find that the classic ℓ^2 state error metric is not particularly useful when comparing different model reduction techniques on sparse training data sets.

Nomenclature

$\boldsymbol{A}$ = Residual minimization weight matrix
C = Constant for kriging square exponential kernel
c_i = ith conservation equation weightings
c_v = Specific heat at constant volume
$\boldsymbol{D}$ = Cell volume scaling matrix
$\boldsymbol{D}_{max}$ = Basis scaling matrix
E = Total energy per unit mass
F_x = Axial force
$\boldsymbol{f}$ = Residual vector
$\tilde{\boldsymbol{f}}$ = Approximate residual vector
$\boldsymbol{h}$ = Clipping function
$\boldsymbol{I}$ = Identity matrix
l = Length scale for kriging square exponential kernel
Ma = Mach number
N = Number of states
n_t = Number of turbulent transport equations
n_{train} = Number of training points

---

*Thermal/Fluid Science and Engineering, MS 9042

†Extreme-Scale Data Science and Analytics, MS 9159, AIAA Member

‡Aerosciences, PO Box 5800, MS 0825, AIAA Member

n_u = Number of conserved quantities
n_μ = Number of parameters
p = Number of columns in reduced-basis matrix $\boldsymbol{\Phi}$
$\boldsymbol{P}_r$ = sampling matrix
$\mathbb{R}^{n \times m}$ = space of $n \times m$ real-valued matrices
r = Dummy variable for interpolation functions
T = Gas temperature
u_i = Conserved quantity i
$\tilde{u}_i$ = ith conserved quantity field for approximate state vector $\tilde{\boldsymbol{x}}$
$\tilde{u}_i'$ = ith conserved quantity field for approximate state vector $\tilde{\boldsymbol{x}}'$
v_j = Fluid velocity in jth cartesian direction
$\hat{x}_j$ = jth generalized coordinate
$\boldsymbol{x}$ = State vector
$\hat{\boldsymbol{x}}$ = Generalized coordinate vector
$\tilde{\boldsymbol{x}}$ = Approximate state vector
$\boldsymbol{X}_{\text{snap}}$ = Snapshot matrix
$\hat{z}$ = Dummy generalized coordinate vector
α_i = ith interpolation function weight
ε_{F_x} = Axial force error
$\varepsilon_{\boldsymbol{x}}$ = State error
ϵ = Radial basis function shape factor
μ = Input parameter
$\bar{\mu}$ = Input parameter mean
$\boldsymbol{\mu}$ = Vector of parameters
$\check{\boldsymbol{\mu}}$ = Vector of normalized parameters
ν = Regularization parameter
ρ = Fluid density
σ = Standard deviation
$\boldsymbol{\Phi}$ = Reduced-basis matrix
$\bar{\boldsymbol{\Phi}}$ = Unscaled reduced-basis matrix
$\boldsymbol{\Phi}_r$ = Reduced-basis matrix for residual snapshots
ϕ_t = Turbulent transport scalar t
Ψ = Radial basis function
$\boldsymbol{\Psi}$ = Test basis matrix
Ω_i = Cell Volume i
Subscripts
$+$ = Strictly positive
FOM = Full order model projection
$local$ = Local in parameter space
min = Minimum allowable value
n = Nearest neighbor
ref = Reference quantity
$wall$ = quantity on a wall
Superscripts
0 = Initial Guess

I. Introduction

Hypersonic aerodynamics simulations are an important part of designing reentry vehicles, missiles, and launch vehicles. It is critical that designers are able to accurately predict forces and heating while in the design stage without conducting expensive, time-consuming tests. Because of this, there is an increasing reliance on computational models for design and analysis of hypersonic configurations [1].

One of the main challenges associated with hypersonic aerodynamics simulations is their large computational cost: due to the many disparate length scales, typical simulation grids must be highly refined close to the body and near shock locations, leading to a very large state space. This makes it computationally expensive to even deploy simplified models such as Reynolds-Averaged Navier-Stokes (RANS) simulations. In addition, to design hypersonic vehicles, engineers need to iterate through hundreds or thousands of designs, each requiring simulations at a variety of speeds, air densities, air temperatures, angles of attack, and airfoil configurations [2], and/or varying flight conditions, vehicle geometry deformations, turbulence model parameters, and boundary layer transition locations. The large number of known and uncertain parameters makes the design of hypersonic vehicles a *many-query problem*, because it requires a large number of model evaluations with different parameters $\boldsymbol{\mu}$. It can be computationally demanding or prohibitively expensive to run full-order models (FOMs) across the parameter space to identify optimal designs and uncertainty estimates. Therefore, there is a strong need for simplified models that run within a reasonable time frame while maintaining high levels of accuracy.

Projection-based reduced-order models (ROMs) are a promising alternative to the full-order models (FOMs) for *many-query problems* [3–9] due to their reduced computational cost. These ROMS project the governing equations onto a subspace constructed using a limited number of full-order simulations. Therefore, the ROM maintains knowledge of the physical governing equations while operating in a much lower dimensional subspace than the FOM. This gives ROMs high accuracy and robustness without needing a large training set. Galerkin projection [10–12] is the basic projection-based method but can perform poorly for convection dominated nonlinear systems including compressible fluid flow simulations [13]. Least-squares Petrov-Galerkin projection (LSPG) is a ROM variant that projects the fully discretized equations onto the basis and has been shown to have improved stability and accuracy over Galerkin projection ROMs [5, 13]. Another promising approach for hypersonic flows is the ℓ^1 minimization approach proposed by Refs. [14, 15]. This approach utilizes a reduced basis comprised of snapshots and utilizes ℓ^1 projection to make ROMs robust for shock-dominated flows. To date, the approach has been demonstrated on 1D unsteady and 2D steady transonic flows.

None of the above approaches directly enforce conservation laws; these methods may violate conservation laws while minimizing an objective function. Conservative least-squares Petrov-Galerkin projection (C-LSPG) is LSPG with a constraint that enforces conservation over the computational domain or subsets of it [16, 17]. The ROMs in Ref. [17] also apply a clipper function to the state approximation to ensure that physical positivity constraints are satisfied. In this paper, we demonstrate LSPG ROMs with the clipper function to develop accurate models in hypersonic aerodynamics.

There are few examples of projection-based ROMs for hypersonic applications. Refs. [18–20] apply Galerkin projection ROMs to thermal modeling in hypersonic flows. Ref. [17] shows that C-LSPG ROMs are viable in hypersonic aerodynamic applications, demonstrating that ROMs can be over 2000 times faster than the FOM with state errors of approximately 0.1%. However, Ref. [17] considers a relatively small range of parameters, which means that a large amount of training data (at least thousands of FOM solutions) would be required to cover the entire flight envelope of a given vehicle.

A major weakness of projection-based ROMs using fixed grids is that they only allow for solutions that lie in a linear subspace. A linear basis is often insufficient to capture the dynamics in hypersonic flows, which generally have a large Kolmogorov N-width due to the change in position of shockwaves as input parameters and/or time are varied. If a ROM is tested on a problem with a shock location different from any of the FOMs used to compute the basis, the ROM will be unable to correctly model the shock location. To address this issue, authors have proposed methods including ROMs with an adaptive linear basis [21–23], ROMs with nonlinear manifolds as a basis (see [24] and the references within), or registration-based ROMs [25–29].

This work demonstrates a registration-based ROM using grid tailoring to determine the snapshot grid displacement as a function of input parameters. Our approach shares similarities with [25], as well as [28]. Following [25], we transform our snapshots such that the shockwaves are in the same position, construct a basis, then transform the basis to move the shockwaves to the correct position. Following [28], we use FOM snapshots that are solved on tailored grids, not interpolated on to them.

This work uses grid tailoring to put the shock at a specified location in the grid during the FOM simulations. Misalignment between the grid and shock can result in large variations in the surface heat flux, but grid tailoring modifies the grid to align with the shock to reduce errors [30, 31]. Grid tailoring has been used in a number of studies of hypersonic flow simulations [30–38].

Our paper expands on Ref. [17] by further developing robust ROMs for hypersonic aerodynamic applications. In the present work, we specifically address the following new directions:

- Larger parameter ranges and/or sparser snapshot sampling than in the previous work,
- Different choices of the low-dimensional state approximation used by the ROM, most notably one which leverages grid tailored FOM solutions,

- Different choices of interpolations for the initial guess used in the ROM minimization problem,
- Comparison of the accuracy of the ROM to additional interpolation.

This paper starts with a brief summary of the grid tailored finite-volume approach used for the FOM in section II. This is followed by a description of steady residual minimization ROM approaches in section III. Section IV contains studies of the ROM for three hypersonic applications, a blunt wedge in section IV.C, a 2D axisymetric flow around the nose cone of the HIFiRE-1 flight vehicle in section IV.D, and a 3D flow around the HIFiRE-1 in section IV.E. Section V wraps up the paper with some conclusions and potential directions for future work on model reduction for hypersonic aerodynamic simulations.

II. Full-order model: finite-volume discretization of hypersonic aerodynamic flows

In this paper, we consider high-Mach external aerodynamics with enthalpy not sufficiently high to drive dissociation of the gas we consider (air). Hence, we solve the perfect gas, compressible Reynolds-Averaged Navier–Stokes (RANS) equations for conserved quantities

$$\begin{pmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \\ u_6 \\ \vdots \\ u_{n_u} \end{pmatrix} = \begin{pmatrix} \rho \\ \rho v_1 \\ \rho v_2 \\ \rho v_3 \\ \rho E \\ \rho \phi_1 \\ \vdots \\ \rho \phi_{n_t} \end{pmatrix} \tag{1}$$

where ρ is density of the fluid, ρv_j is the fluid density times the fluid velocity v_j, ρE is the fluid density times the total energy per unit mass E, and ϕ_t is a set of scalars that belongs to the turbulent transport equations and n_t is the number of turbulent transport equations. For the full set of governing equations, please consult Refs. [17] and/or [39].

We consider steady-state solutions of the Navier–Stokes or RANS equations, which can be expressed as the solution of the following system.

$$\boldsymbol{f}(\boldsymbol{x}; \boldsymbol{\mu}) = \boldsymbol{0}, \tag{2}$$

where $\boldsymbol{f}(\cdot, \cdot) \in \mathbb{R}^N$ is the residual vector, which is a function of the state vector $\boldsymbol{x} \in \mathbb{R}^N$ and the vector of system parameters $\boldsymbol{\mu} \in \mathbb{R}^{n_\mu}$. The governing equations are discretized with a cell-centered finite volume method [40, 41]. Therefore, the state vector $\boldsymbol{x}$ is comprised of each conserved quantity at each cell center, while the residual vector $\boldsymbol{f}$ is comprised of the discretized conservation laws (the governing equations) at each cell center. A more detailed derivation of (2) is included in Ref. [17].

A. Grid tailoring

Past studies have shown that misalignment between the grid and a bow shock wave can lead to large errors, particularly for heat flux [42]. Grid tailoring was introduced by Ref. [30] to modify the mesh to be aligned with the shock. Note that grid tailoring moves existing nodes; it does not add nodes to the mesh, making it an r-adaptation approach.

The procedure used for the results presented in this paper is very similar to that used by Ref. [30]. Each grid tailoring has several parameters: the number of buffer cells N_b, the maximum initial spacing in the wall normal direction $w_{0,max}$, the shock criterion S_c, and the number of smoothing iterations N_{smooth}. The shock detection condition given in Eq. 3 is used to find the first cell of the shock detected along a line from the inlet to the vehicle surface, such as the lines shown in Fig. 1.

$$\text{Shock detection condition} \equiv \frac{|Ma - Ma_{inlet}|}{Ma_{inlet}} > S_c \tag{3}$$

The grid tailoring algorithm is given as:

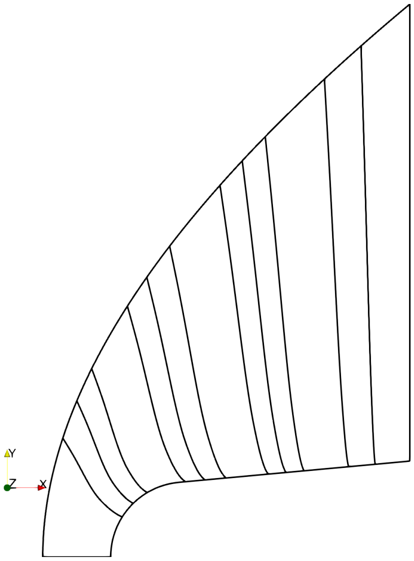

Fig. 1 Examples of mesh lines used for grid tailoring of a blunt wedge. The inlet boundary is the left edge and the blunt wedge is in the bottom right corner.

1) From each face on the body, find a mesh line connecting adjacent cell centers from the surface to the inlet, like those shown in Fig. 1
2) Find the first cell along each mesh line with the Mach number meeting the shock detection condition given in Eqn. 3. The location of that cell is assigned to a face-centered surface quantity $shock_s$.
3) Interpolate $shock_s$ from faces to nodes.
4) Apply a smoothing algorithm that successively replaces each value of $shock_s$ with the average of its neighbors and itself a total of N_{smooth} times.
5) Find mesh lines along nodes from the surface to the inlet.
6) Redistribute nodes. Set the first cell width to the smaller of $y^+ = 1$ or $w_{0,max}$. Next ensure that there are N_b cells between the inlet and the cell centered at $shock_s$. Cells between the shock and surface grow with a constant growth factor.
7) Apply nearest neighbor interpolation to transfer variable values from previous mesh.

Grid tailoring can be applied multiple times to get a smooth grid that is well aligned with the shock; in practice applying grid tailoring twice is generally sufficient. Grid tailoring should be applied after the shock has stopped moving in a pseudo-time simulation.

Fig. 2a shows the Mach number around a bow shock overlaid on the grid of a converged simulation without tailoring, and the shock is clearly misaligned with the grid. After tailoring is applied twice, the shock and grid are well aligned, as shown in Fig.2b. Furthermore, the grid has a finer resolution near the shock because the grid is used more efficiently in the tailored simulations, as there are fewer mesh cells between the inlet and the shock.

Grid tailoring adds computational cost because the residual must be sufficiently converged after each round of grid tailoring. However, grid tailoring has multiple benefits, including:

- Significantly more accurate heat flux computations
- A single grid can be used over a large Mach number range
- The grid is used more efficiently, resulting in finer mesh resolution in the shock and boundary layer.
- The boundary layer is ensured to have a specified resolution, at least for the first cells off the wall (e.g. $y^+ <= 1$).

III. Reduced-order model: steady-state residual minimization

The following section discusses each component of the steady-state reduced-order models considered in this paper.

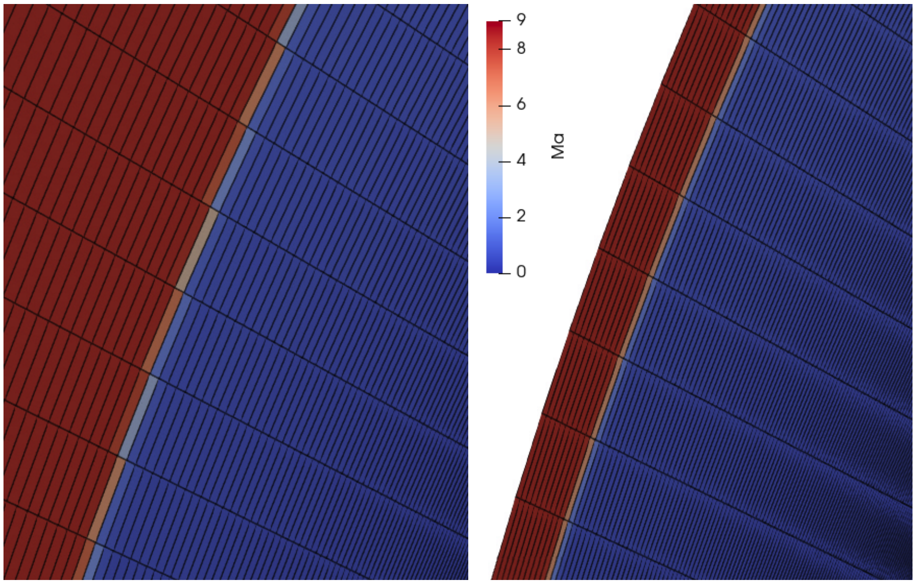

(a) Without grid tailoring **(b) With grid tailoring**

Fig. 2 Bow shock without and with grid tailoring, colored by Mach number

A. State Approximation and Basis

Projection-based reduced-order models rely on approximating the high-fidelity state as:

$$\boldsymbol{x}(\boldsymbol{\mu}) \approx \tilde{\boldsymbol{x}}(\boldsymbol{\mu}) = \boldsymbol{h}(\tilde{\boldsymbol{x}}'(\boldsymbol{\mu})) = \boldsymbol{h}(\boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}) + \boldsymbol{\Phi}\hat{\boldsymbol{x}}(\boldsymbol{\mu})), \tag{4}$$

where $\boldsymbol{\Phi} \in \mathbb{R}^{N \times p}$ is the reduced-basis matrix of dimension $p \leq N$, $\hat{\boldsymbol{x}} \in \mathbb{R}^p$ denotes the generalized coordinates, and $\boldsymbol{h}$ is a nonlinear transformation function to be defined in section III.A.3. The classic approach uses $\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{x}$, resulting in a linear affine approximation. Note that the approximation Eq. (4) can also be used for reducing the cost of surrogate models by interpolating the reduced coordinates $\hat{\boldsymbol{x}}$ rather than the full state $\boldsymbol{x}$.

This basis $\boldsymbol{\Phi}$ can be computed in a variety of ways during the offline stage, e.g., eigenmode analysis, Proper Orthogonal Decomposition (POD) [43], reduced-basis method [44, 45], or a dictionary of solution snapshots [15]. Typically, $\boldsymbol{\Phi}$ is orthonormal, and $\boldsymbol{\Phi}^T\boldsymbol{\Phi} = \boldsymbol{I}$. Note that it does not need to be orthonormal for residual minimization approaches discussed in section III.B.

The following section discusses two classes of approaches used to compute the basis $\boldsymbol{\Phi}$. We only consider linear subspace approximations in this paper due to their low computational cost and relative ease of implementation. We consider both of these in combination with grid tailoring to obtain a registration-based ROM.

1. Proper-Orthogonal Decomposition

POD modes are the most popular choice of basis for model reduction since they provide the optimal state approximation (in an L^2 sense) for a given subspace dimension p [43]. The POD modes are the columns of the left singular vector matrix $\boldsymbol{U}$ corresponding to the p leading singular values obtained from the singular value decomposition

$$\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T = \boldsymbol{D}_{max}^{-1}(\boldsymbol{X}_{\text{snap}} - \boldsymbol{x}_{\text{ref}}). \tag{5}$$

where $\boldsymbol{X}_{\text{snap}} \in \mathbb{R}^{N \times n_{\text{train}}}$ is the snapshot matrix, in which the ith column corresponds to a FOM solution $\boldsymbol{x}^i$ at parameter values $\boldsymbol{\mu}_i$. $\boldsymbol{D}_{max} \in \mathbb{R}^{N \times N}$ is a diagonal matrix with the maximum absolute value of each conserved quantity along the main diagonal. These maxima are computed over all cells in all training snapshots. The reference state $\boldsymbol{x}_{\text{ref}}$ is simply the mean of all snapshots. Note that this is one relatively straightforward and convenient choice for the reference state; other choices may prove better, and will be the subject of future studies. It should be pointed out that the snapshot mean does not have to be physically or statistically meaningful. Therefore, our choice of $\boldsymbol{x}_{\text{ref}}$ is nothing more than a convenient reference point in phase space to center our snapshot data.

To account for the scaling of the snapshot data, the basis $\boldsymbol{\Phi}$ is defined as

$$\boldsymbol{\Phi} = \boldsymbol{D}_{max}\bar{\boldsymbol{\Phi}}, \tag{6}$$

where $\bar{\boldsymbol{\Phi}}$ is the first p columns of $\boldsymbol{U}$.

The normalization by $\boldsymbol{D}_{max}$ is done because using dimensional quantities is standard practice for hypersonic CFD codes, especially for cases with non-equilibrium thermochemical effects. Normalizing snapshot data prior to computing POD modes increases numerical robustness. This is because reducing the range of variable scales makes the normalized, centered snapshot matrix Eq. (5) better conditioned than the centered, unscaled snapshot matrix $(\boldsymbol{X}_{\text{snap}} - \boldsymbol{x}_{\text{ref}})$.

Because of the diagonal matrix in Eq. (6), the basis $\boldsymbol{\Phi}$ is not orthonormal ($\boldsymbol{\Phi}^T\boldsymbol{\Phi} \neq \boldsymbol{I}$), so projections need to be done with the Moore-Penrose pseudo inverse of $\boldsymbol{\Phi}$. Since $\boldsymbol{D}_{max}$ is diagonal we can write the pseudo inverse as an ℓ^2 projection onto $\bar{\boldsymbol{\Phi}}$

$$\hat{\boldsymbol{x}}(\boldsymbol{\mu}) = (\boldsymbol{\Phi})^{+}(\boldsymbol{x}(\boldsymbol{\mu}) - \boldsymbol{x}_{\text{ref}}) = \bar{\boldsymbol{\Phi}}^T\boldsymbol{D}_{max}^{-1}(\boldsymbol{x}(\boldsymbol{\mu}) - \boldsymbol{x}_{\text{ref}}). \tag{7}$$

2. Local Dictionary

When the shockwave location varies with $\boldsymbol{\mu}$ or time for unsteady cases, POD modes can provide a poor approximation of snapshots [15]. These issues arise in part from the truncation of the POD basis to p columns: since the leading POD modes tend to be smooth, it is difficult to approximate sharp discontinuities such as shocks with them, resulting in oscillations near the shockwave. Note that this issue is similar in nature to the inability of a finite number of Fourier modes to approximate square waves.

The dictionary approach proposed by Refs. [14, 15] avoids these overshoots by not conducting any data compression. Furthermore, it uses local data, preventing snapshot data far away in parameter space (snapshots with very different shock wave configurations) from negatively impacting the accuracy of the approximation Eq. (4). We define our local dictionary basis as

$$\boldsymbol{\Phi}(\boldsymbol{\mu}) = \boldsymbol{X}_{\text{local}}(\boldsymbol{\mu}) - \boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}_n)\boldsymbol{1}^T \tag{8}$$

where $\boldsymbol{\mu}_n$ is the nearest parameter to $\boldsymbol{\mu}$ for which a FOM snapshot is available, $\boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}_n)$ is the snapshot corresponding to $\boldsymbol{\mu}_n$, $\boldsymbol{X}_{\text{local}} \in \mathbb{R}^{N\times p}$ is a dictionary of the p snapshots whose parameters $\boldsymbol{\mu}$ are closest to $\boldsymbol{\mu}_n$ by some definition of distance, and $\boldsymbol{1} \in \mathbb{R}^p$ is a vector of 1s. Since the parameters can be defined in units with vastly different scales, we compute distances in parameter space between normalized parameter vectors $\check{\boldsymbol{\mu}}$, with each normalized parameter $\check{\mu}$ defined as

$$\check{\mu} \equiv \frac{\mu - \bar{\mu}}{\sigma(\mu)} \tag{9}$$

where $\bar{\mu}$ is the mean and $\sigma(\mu)$ is the standard deviation of parameter μ in the training set.

Note that it is not guaranteed that the columns of $\boldsymbol{X}_{\text{local}}$ are linearly independent, which can lead to some numerical issues [14, 15]. These issues can be overcome by orthogonalizing the dictionary via QR-decomposition.

3. Clipping Function

One of the shortcomings of an affine linear trial subspace used to compute the approximate state is that there exists some $\hat{\boldsymbol{x}}$ such that $\tilde{\boldsymbol{x}}$ will contain some non-physical local phenomena such as regions of negative density or temperature. Refs. [46, Section VI.E] and [17] shows that ensuring $\tilde{\boldsymbol{x}}$ does not have non-physical local flow features can significantly improve the robustness of LSPG and Galerkin ROMs.

The clipping function $\boldsymbol{h}_{clip}(\tilde{\boldsymbol{x}})$ considered in this paper is designed specifically for the case of a single-specie calorically perfect gas; clipping functions for flows with thermal and/or chemical nonequilibrium are being investigated in an ongoing effort to enable ROMs for higher speed/enthalpy flight regimes.

The perfect gas clipping function $\boldsymbol{h}_{clip}(\tilde{\boldsymbol{x}})$ we consider is designed to enforce fluid density $\rho > \rho_{min}$ and temperature $T > T_{min}$ in the flowfield represented by $\tilde{\boldsymbol{x}}$. The density field $\tilde{u}_1$ is computed by

$$\tilde{u}_1 = \max(\rho_{min}, \tilde{u}_1'), \tag{10}$$

where $\rho_{min} > 0$ is a minimum density that should be very small relative to the free stream density and $\tilde{u_1}'$ is the density field computed from the linear affine approximation in Eq. (4),

$$\tilde{\boldsymbol{x}}' = \boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}) + \boldsymbol{\Phi}\hat{\boldsymbol{x}}(\boldsymbol{\mu}) \tag{11}$$

The clipper for $T > T_{min}$ is derived from the definition of the total energy per unit mass E for a perfect gas

$$c_{\text{v}}T = E - \frac{1}{2}(v_j v_j). \tag{12}$$

where c_v is the constant volume specific heat and v_j is the fluid velocity. Since $c_{\text{v}} > 0$, a lower bound for temperature, T_{min}, can be enforced by setting

$$\tilde{u}_5 = \max\left(\tilde{u}_1 c_v T_{min} + \frac{1}{2\tilde{u}_1}\left[\tilde{u}_2^2 + \tilde{u}_3^2 + \tilde{u}_4^2\right]\ ,\ \tilde{u}_5'\right), \tag{13}$$

Note the presence of $\tilde{u}_1$ rather than $\tilde{u}_1'$, this term is included because the density clipping function must be applied before the energy clipping function. Therefore, $\boldsymbol{h}_{clip}(\tilde{\boldsymbol{x}})$ is of the form $\boldsymbol{h}_5(\boldsymbol{h}_1(\tilde{\boldsymbol{x}}'))$, where $\tilde{z} = \boldsymbol{h}_1(\tilde{\boldsymbol{x}}')$ applies Eq. (10) to $\tilde{\boldsymbol{x}}'$, and $\tilde{\boldsymbol{x}} = \boldsymbol{h}_5(\tilde{z})$ applies Eq. (13) to $\tilde{z}$.

4. Grid approximation

When simulations are grid tailored, the grid is different for each FOM solution snapshot, so running a ROM requires generating a new grid. A grid can be generated based on the parameters by either computing POD modes of the grid displacements from the training set or by creating a local dictionary of grid displacements. In this paper we consider a local dictionary with a reference state of zero, since the snapshots are of grid displacements, not node positions.

$$\boldsymbol{\Phi}_{grid}(\boldsymbol{\mu}) = \boldsymbol{X}_{\text{grid,local}}(\boldsymbol{\mu}) \tag{14}$$

Where $\boldsymbol{X}_{\text{grid,local}} \in \mathbb{R}^{3n_{grid} \times p_{grid}}$ is a matrix comprised of p_{grid} vectors of grid displacements in the x-,y- and z-directions, n_{grid} is the number of nodes in the grid, and p_{grid} is the number of snapshots contained in the grid displacement dictionary. As for the state dictionary, the grid displacement dictionary columns are the p_{grid} nearest neighbors in parameter space to $\boldsymbol{\mu}_n$, the closet point in parameter space to $\boldsymbol{\mu}$ with a snapshot.

While a POD basis can be formed for the grid displacements, it is found that the local dictionary approach works just as well or better for the small amount of training data considered in this paper. Therefore, we proceed with the local dictionary, since it does not require computing an SVD of grid displacement data.

5. Basis variables

We consider bases computed from snapshots comprised of conserved or primitive variables, where primitive variables are defined as

$$\begin{pmatrix} \rho \\ v_1 \\ v_2 \\ v_3 \\ T \\ \phi_1 \\ \vdots \\ \phi_{n_t} \end{pmatrix} \tag{15}$$

Since the FOM residual $\boldsymbol{f}$ is a function of conserved variables, an approximate state computed with primitive variables needs to be converted to conserved variables. To simplify notation, this transformation, combined with the appropriate clipper function, is defined as $\boldsymbol{h}(\tilde{\boldsymbol{x}})$. Note that in the case where primitive variables are used for the basis, temperature can be clipped directly

$$T = \max\left(T_{min}, T'\right), \tag{16}$$

B. Residual Minimization Reduced-Order Modeling Approaches

This paper considers ROMs that can be expressed as a residual minimization problem. In the case of a steady simulation, we substitute the approximation $\boldsymbol{x} \leftarrow \tilde{\boldsymbol{x}}$ into the FOM steady-state equations Eq. (2), and subsequently minimize the residual in a weighted ℓ^P-norm, i.e.,

$$\hat{\boldsymbol{x}} = \underset{\hat{\boldsymbol{z}} \in \mathbb{R}^p}{\arg\min} \| \boldsymbol{A} \boldsymbol{f}(\boldsymbol{h}(\boldsymbol{x}_{\mathrm{ref}}(\boldsymbol{\mu}) + \boldsymbol{\Phi}\hat{\boldsymbol{z}}; \boldsymbol{\mu})) \|_P. \tag{17}$$

where $\boldsymbol{A} = \boldsymbol{I}$, for example.

The nonlinearity of $\boldsymbol{f}$ requires evaluating $\boldsymbol{f}$ and the corresponding Jacobian $\partial \boldsymbol{f} / \partial \boldsymbol{x}$ for each nonlinear solver iteration. The computational cost of this scales with N, which limits the computational efficiency of the ROM. To ensure that the ROM incurs an N-independent operation count, the weighting matrix $\boldsymbol{A}$ should be sparse in the sense that it has a small number of nonzero columns. In this case, one can set $\boldsymbol{A} = (\boldsymbol{P}_r \boldsymbol{\Phi}_r)^+ \boldsymbol{P}_r$ and $\boldsymbol{A} = \boldsymbol{P}_r$ in the case of gappy POD and collocation, respectively. Here, $\boldsymbol{P}_r \in \{0,1\}^{n_{p,r} \times N}$ denotes a sampling matrix comprising selected rows of the $N \times N$ identity matrix, while $\boldsymbol{\Phi}_r \in \mathbb{R}^{N \times p_r}$ is a p_r-dimensional reduced-basis matrix constructed for the residual $\boldsymbol{f}$. Employing the gappy POD approximation results in the GNAT reduced-order model [6].

Note that we do not use hyper-reduction for the results presented in this paper; applying hyper-reduction to the grid-tailored ROM approaches is left for future work. Fortunately, cost reductions can be achieved for steady ROMs without hyper-reduction since they do not use the same pseudo-time-stepping approach as the FOM, which requires hundreds or thousands of iterations [17]. The direct solver approaches used for the ROM requires far fewer evaluations of $\boldsymbol{f}$ and the corresponding Jacobian $\partial \boldsymbol{f} / \partial \boldsymbol{x}$, resulting in substantial cost reductions, as shown in [17] and section IV.

1. Least-Squares Petrov–Galerkin Projection

LSPG corresponds to Eq. (17) with $P = 2$, or an ℓ^2 minimization of the discrete residual. The first-order optimality conditions can be determined analytically and correspond to a projection equation of the form

$$\boldsymbol{\Psi}^T \boldsymbol{A} \boldsymbol{f}(\boldsymbol{h}(\boldsymbol{x}_{\mathrm{ref}}(\boldsymbol{\mu}) + \boldsymbol{\Phi}\hat{\boldsymbol{x}}; \boldsymbol{\mu})) = \boldsymbol{0}, \tag{18}$$

for which the optimal choice of test basis $\boldsymbol{\Psi} \in \mathbb{R}^{N \times p}$ is

$$\boldsymbol{\Psi} = \boldsymbol{A} \left. \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} \right|_{\boldsymbol{h}(\tilde{\boldsymbol{x}})} \left. \frac{\partial \boldsymbol{h}}{\partial \boldsymbol{x}} \right|_{\tilde{\boldsymbol{x}}'} \boldsymbol{\Phi}. \tag{19}$$

This is in contrast to a Galerkin ROM, for which $\boldsymbol{\Psi} = \boldsymbol{\Phi}$. The Petrov-Galerkin projection is better suited for systems with asymmetric Jacobian matrices $\partial \boldsymbol{f} / \partial \boldsymbol{x}$, most notably advection-dominated systems including Navier-Stokes solvers for practical aerodynamic applications. It is important to note that the test basis for LSPG also includes the Jacobian of the variable transformation and clipping function $\boldsymbol{h}$. This implies that one should choose $\boldsymbol{h}(\boldsymbol{x})$ to be differentiable with respect to $\boldsymbol{x}$, however we consider a piecewise differentiable $\boldsymbol{h}$ without issue in this study.

We also considered C-LSPG, a version of LSPG with a conservation constraint proposed in Refs. [16] that has previously been applied to steady hypersonic aerodynamics [17]. However, we found that C-LSPG has robustness issues for a number of the applications considered in this paper, therefore we leave a more in-depth investigation of C-LSPG to future work.

C. Interpolation methods for model reduction

One fully data-driven alternative to the ROMs presented in section III.B is to use interpolation to compute the generalized coordinates $\hat{\boldsymbol{x}}$. We consider several different interpolation schemes in this paper to 1) provide initial guesses for the ROMs and 2) provide a comparison for the accuracy of the ROMs. The convergence properties and accuracy of a ROM can be strongly dependent on the initial guess $\hat{\boldsymbol{x}}^0(\boldsymbol{\mu})$. Interpolation methods are negligible in computational cost relative to ROMs and are therefore excellent candidates to compute $\hat{\boldsymbol{x}}^0(\boldsymbol{\mu})$. At the same time, the low cost of interpolation methods on low-dimensional subspaces make them the clear choice when their accuracy is similar to that of the ROM.

Each interpolation scheme can be written as

$$\hat{\boldsymbol{x}}(\boldsymbol{\mu}) = g(\hat{\boldsymbol{x}}(\boldsymbol{\mu}_i), \check{\boldsymbol{\mu}}_i, \check{\boldsymbol{\mu}}) \tag{20}$$

where $g(\cdot,\cdot,\cdot)$ is the interpolating function and $\hat{\boldsymbol{x}}(\boldsymbol{\mu}_i)$ are the generalized coordinate values at the training set parameters $\boldsymbol{\mu}_i$. Note that that parameter space is normalized as for the dictionary by Eq. (9) to improve the conditioning of the interpolation scheme.

For a potentially non-orthonormal basis, $\boldsymbol{\Phi}$, each $\hat{\boldsymbol{x}}(\boldsymbol{\mu}_i)$ is computed by

$$\hat{\boldsymbol{x}}^0(\boldsymbol{\mu}_i) = \boldsymbol{\Phi}^+(\boldsymbol{x}(\boldsymbol{\mu}_i) - \boldsymbol{x}_{\text{ref}}(\boldsymbol{\mu}_i)), \tag{21}$$

where $(\ )^+$ is the Moore-Pentrose pseudo-inverse. For a local dictionary basis Eq. (8), this reduces to

$$\hat{x}_i(\boldsymbol{\mu}_j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}, \tag{22}$$

where $\hat{x}_i$ is the ith component of $\hat{\boldsymbol{x}}$.

A variety of interpolation methods are available and are tested in this study for interpolating the generalized coordinates. A summary of each method is given here.

1. Inverse Distance Interpolation

Two previous works used inverse distance interpolation to construct the initial guess [17, 47]. Inverse distance interpolation constructs a smooth interpolating function using weights linearly dependent on the inverse distance between each sampling point and training data. The interpolating function is

$$\hat{\boldsymbol{x}}(\boldsymbol{\mu}) = \sum_{i=1}^{n_{\text{train}}} \alpha_i \hat{\boldsymbol{x}}(\boldsymbol{\mu}_i), \tag{23}$$

where n_{train} is the number of parameter points in the training set. The weights α_i are found with Equation 24

$$\alpha_i = \frac{1/\|\check{\boldsymbol{\mu}} - \check{\boldsymbol{\mu}}\|_2}{\sum_{j=1}^{n_{\text{train}}} 1/\|\check{\boldsymbol{\mu}} - \check{\boldsymbol{\mu}}\|_2}. \tag{24}$$

2. Kriging

Kriging, also called Gaussian process regression, is a nonparametric interpolation method that predicts distributions at sampling points by performing inference over function distributions. A detailed explanation of kriging is given in [48]. Kriging does not have hyperparameters, but requires a kernel function to define similarity between points. The kernel function used in this work is the widely used squared exponential kernel

$$k(r) = C \exp\left(-\frac{\|\check{\boldsymbol{\mu}}_i - \check{\boldsymbol{\mu}}_j\|^2}{2l^2}\right), \tag{25}$$

where the characteristic length scale l and constant value C are learned during the inference step. Kriging has been used in Ref. [49] to construct an initial guess for steady ROMs.

3. Radial Basis Function interpolation

Radial basis function interpolation uses a weighted sum of radial basis functions to construct a smooth interpolating function for each generalized coordinate $\hat{x}_j(\boldsymbol{\mu})$, given by

$$\hat{x}_j(\boldsymbol{\mu}) = \sum_{i=1}^{n_{\text{train}}} \alpha_i \Psi(\|\check{\boldsymbol{\mu}} - \check{\boldsymbol{\mu}}_i\|_2), \tag{26}$$

where $\Psi(r)$ is the radial basis function. For this study, a quadratic basis function is used,

$$\Psi(r) = \left[\left(\frac{r}{\epsilon}\right)^2 + 1\right]^{0.5} \tag{27}$$

where ϵ is the shape parameter. The weights α_i for the jth generalized coordinate $\hat{x}_j$ are found by solving the system

$$\begin{bmatrix} \Psi\|\check{\boldsymbol{\mu}}_0 - \check{\boldsymbol{\mu}}_0\|_2 & \Psi\|\check{\boldsymbol{\mu}}_1 - \check{\boldsymbol{\mu}}_0\|_2 & \dots & \Psi\|\check{\boldsymbol{\mu}}_{n_{\text{train}}} - \check{\boldsymbol{\mu}}_0\|_2 \\ \Psi\|\check{\boldsymbol{\mu}}_0 - \check{\boldsymbol{\mu}}_1\|_2 & \Psi\|\check{\boldsymbol{\mu}}_1 - \check{\boldsymbol{\mu}}_1\|_2 & \dots & \Psi\|\check{\boldsymbol{\mu}}_{n_{\text{train}}} - \check{\boldsymbol{\mu}}_1\|_2 \\ \vdots & \vdots & \ddots & \vdots \\ \Psi\|\check{\boldsymbol{\mu}}_0 - \check{\boldsymbol{\mu}}_{n_{\text{train}}}\|_2 & \Psi\|\check{\boldsymbol{\mu}}_1 - \check{\boldsymbol{\mu}}_{n_{\text{train}}}\|_2 & \dots & \Psi\|\check{\boldsymbol{\mu}}_{n_{\text{train}}} - \check{\boldsymbol{\mu}}_{n_{\text{train}}}\|_2 \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_{n_{\text{train}}} \end{bmatrix} = \begin{bmatrix} \hat{x}_j(\mu_0) \\ \hat{x}_j(\mu_1) \\ \vdots \\ \hat{x}_j(\mu_{n_{\text{train}}}) \end{bmatrix} \tag{28}$$

where n_{train} is the number of points included in the training set.

The shape parameter ϵ in Eqn. 27 controls the width of the Gaussian basis function. For $\epsilon \longrightarrow 0$, the basis functions are very narrow and the interpolation function is zero everywhere except close to the known points x_i. For $\epsilon \longrightarrow \infty$, the matrix in Eqn. 28 is ill-conditioned. In practice, the shape parameter should be set as high as possible while maintaining a condition number sufficiently below numerical precision [50]. In this work, the shape parameter is fixed at 20. Computing the weights α_i can be done in the offline stage, keeping the online computational costs low.

IV. Reduced-order model numerical experiments

The following section discusses the common features of the numerical experiments conducted to test the accuracy, performance, and robustness of the various ROM techniques discussed in the previous sections. These experiments make use of two codes being developed at Sandia National Laboratories, namely SPARC and *Pressio*[*]. SPARC (Sandia Parallel Aerodynamics and Reentry Code) is a compressible CFD code focused on aerodynamics and aerothermodynamics problems. It solves the compressible Navier–Stokes and Reynolds-Averaged Navier–Stokes (RANS) equations on structured and unstructured grids using a cell-centered finite volume discretization scheme [39]. Its target use cases are transonic flows to support gravity bomb analyses and hypersonic flows for re-entry vehicle analyses.

Pressio is an open-source C++11 header-only library aimed at enabling parallel, scalable, and performant ROM capabilities to be adopted by any C++ application with a minimally intrusive API *Pressio*. The main design principle behind *Pressio* is to require an application to only expose, for a given state $\boldsymbol{x}$, time t, and parameters $\boldsymbol{\mu}$, the velocity vector $\boldsymbol{f}(\boldsymbol{x}, t; \boldsymbol{\mu})$ and the action of the Jacobian matrix $\partial \boldsymbol{f}(\boldsymbol{x}, t; \boldsymbol{\mu})/\partial \boldsymbol{x}$. Exposing from SPARC the required functionalities was relatively easy, since it involved the creation of a new adapter class and no changes to the original SPARC code. We remark that while in this work we limit our attention to LSPG, the *same* interface developed in SPARC can now be used to run *any* of the ROM methods supported in *Pressio*.

A. Full-Order Models

We demonstrate the application of grid-tailored ROMs on three test cases: a two-dimensional blunt wedge, a two-dimensional HiFIRE-1 nose cone, and a three-dimensional HiFIRE-1 cone and body. Unless otherwise stated, all cases were run until the residual ℓ^2 norm decreased by 7 orders of magnitude.

We consider cases with and without grid tailoring. Grid tailoring is applied twice for all cases; further details including the number of smoothing iterations are detailed in the following sections. The difference in QoIs between the grid tailored and non-tailored simulations is analyzed by considering a relative difference,

$$\frac{|QoI_{tailored} - QoI_{non-tailored}|}{|QoI_{tailored}|}. \tag{29}$$

B. Reduced-Order Models

We apply ROMs to the three cases mentioned above. For each case, different ROM setups are compared for a fixed training set. ROMs are considered for:

- Training FOMs with and without grid tailoring
- LSPG and C-LSPG
- Inverse distance, kriging and RBF interpolation for initial guesses

The nonlinear least-squares problem arising from LSPG is solved via a QR-based Gauss-Newton, while the one stemming from C-LSPG is solved via normal equations. The Gauss-Newton solver used is provided by *Pressio*, and is run until the relative residual ℓ^2 norm from the residual defined in Equation 18 falls below 10^{-6} or after 100 iterations, unless otherwise specified. Note that 100 iterations is a relatively large number of iterations selected to prevent the ROM nonlinear solver from running infinitely long.

---

[*] https://github.com/Pressio

The selection of the weighting matrix $\boldsymbol{A}$ in (17) is crucial for the accuracy and speed of LSPG and C-LSPG. Similarly to Ref. [17], we set $\boldsymbol{A} = \boldsymbol{D} \in \mathbb{R}^{N\times N}$, where $\boldsymbol{D} \in \mathbb{R}^{N\times N}$ is defined as

$$\boldsymbol{D}_{i+jn_u,i+jn_u} = c_i|\Omega_j|, \quad i \in \mathbb{N}(n_u),\ j \in \mathbb{N}(N_\Omega). \tag{30}$$

This is a diagonal matrix whose elements correspond to the size of each cell volume $|\Omega_k|$ and a conservation equation weight c_i. Each c_i should be chosen to normalize the contributions from the discrete residuals of each governing equation. In practice, a good choice of c_i is the reciprocal of the maximum value of the ith conserved quantity, as in (6). It is found that this choice of $\boldsymbol{A}$ vastly improves the convergence rate of Gauss–Newton iteration, along with the accuracy of the ROM solution, as observed in Ref. [17].

The clipping functions (10) and (13) are applied with $\rho_{min} = 10^{-6}$ and $T_{min} = 100$.

The ROM state initial guess and grid are generated by using an interpolation method on the dictionary of grid displacements $\boldsymbol{\Phi}_{grid}(\boldsymbol{\mu})$ and the state dictionary $\boldsymbol{\Phi}(\boldsymbol{\mu})$, respectively. In all the cases presented here, the same interpolation method is used for both the state initial guess and grid. In this work, the ROM grid is fixed and not modified in the residual minimization.

The interpolations were implemented in Python. A custom inverse distance interpolation scheme was used, the scikit-learn package was used for kriging [51], and the scipy package was used for RBF interpolation [52].

1. Error Metrics

We measure the accuracy of the ROM with the following error metrics. Firstly, the state ℓ^2 error, defined as

$$\varepsilon_{\boldsymbol{x}} = \frac{\|\boldsymbol{x}(\boldsymbol{\mu}) - \tilde{\boldsymbol{x}}(\boldsymbol{\mu})\|_2}{\|\boldsymbol{x}(\boldsymbol{\mu})\|_2}, \tag{31}$$

where $\boldsymbol{x}(\boldsymbol{\mu})$ and $\tilde{\boldsymbol{x}}(\boldsymbol{\mu})$ are the full state computed with the FOM and some approximation with (4), respectively. Although this metric does not have a physical interpretation, it provides a means of verifying the LSPG ROM, since the state ℓ^2 error should decrease as p is increased. Note that the state ℓ^2 error is very close to the ℓ^2 error of conserved energy, ρE, since ρE is much larger in magnitude than the other conserved quantities. The vector $\tilde{\boldsymbol{x}}(\boldsymbol{\mu})$ is usually the ROM solution, but we also compute $\tilde{\boldsymbol{x}}(\boldsymbol{\mu})$ for other states as well, such as the interpolated state used as an initial guess for the ROM.

Secondly, we compute the axial force, defined as

$$\varepsilon_{F_x} = \frac{|F_x(\boldsymbol{\mu}) - \tilde{F}_x(\boldsymbol{\mu})|}{|F_x(\boldsymbol{\mu})|}, \tag{32}$$

where $F_x(\boldsymbol{\mu})$ and $\tilde{F}_x(\boldsymbol{\mu})$ are the integrals of axial force computed with the FOM and some corresponding approximation with (4), respectively. Similarly, the integrated heat flux error is defined as

$$\varepsilon_{Q_{wall}} = \frac{|Q_{wall}(\boldsymbol{\mu}) - \tilde{Q}_{wall}(\boldsymbol{\mu})|}{|Q_{wall}(\boldsymbol{\mu})|}, \tag{33}$$

where $Q_{wall}(\boldsymbol{\mu})$ is the integrated heat flux over the surface computed with the FOM and $\tilde{Q}_{wall}(\boldsymbol{\mu})$ is the integrated heat flux of an approximation.

C. 2D Wedge

The two-dimensional blunt wedge is shown in Figure 3. This case is a simple geometry that is representative of the nose or leading edge of many hypersonic vehicles. The corresponding free stream conditions considered in this study are listed in Table 1.

The blunt wedge mesh, shown in Fig. 3 has 100,000 cells, corresponding to a state-space size of 500,000 with $n_u = 5$ since we are considering laminar flow. The flow is solved using pseudo-time stepping with a backward Euler time step and scheduled increases in CFL number. The convergence criteria are a reduction in relative residual by 7 orders of magnitude or 5,000 pseudo-time steps. Convergence is achieved in around 400-500 steps for the range of freestream velocities and densities considered.

Figure 4 shows the flow field for three different simulations at different freestream Mach numbers and densities. Note that the position of the bow shock changes significantly with freestream Mach number. For a fixed freestream Mach number, the boundary layer thickens as density decreases. In Fig. 5, we plot the axial force, F_x, to show that

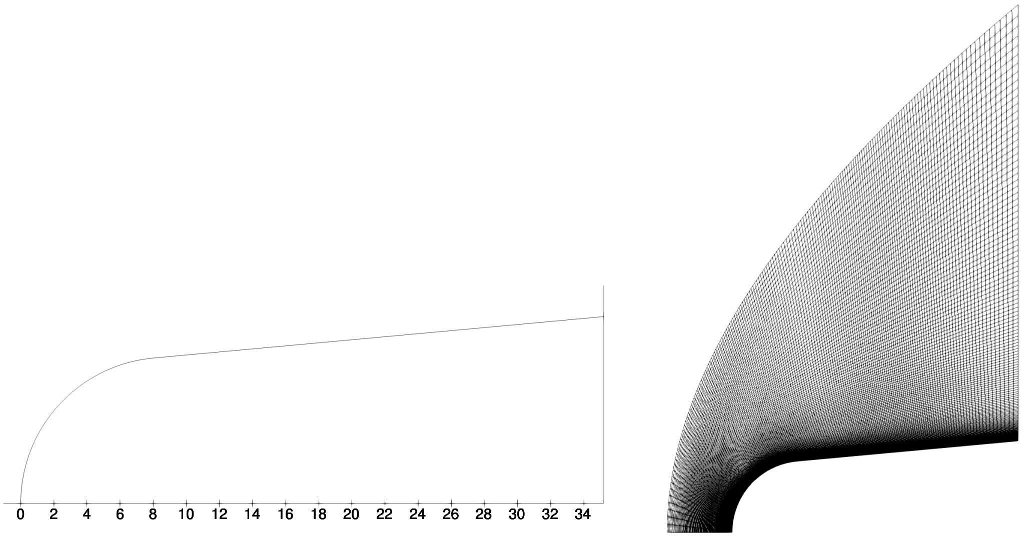

(a) Geometry, units in mm **(b) Mesh**

Fig. 3 2D Blunt Wedge

Table 1 Free stream flow conditions for the two-dimensional blunt wedge.

Density	0.035 to 0.085 kg/m^3
Velocity	885.23 to 2803.21 m/s
Mach Number	3.0 to 9.5
Angle of attack	0.0°
Temperature	216.66 K
Reynolds Number	1.5 $\times 10^7$ $1/m$

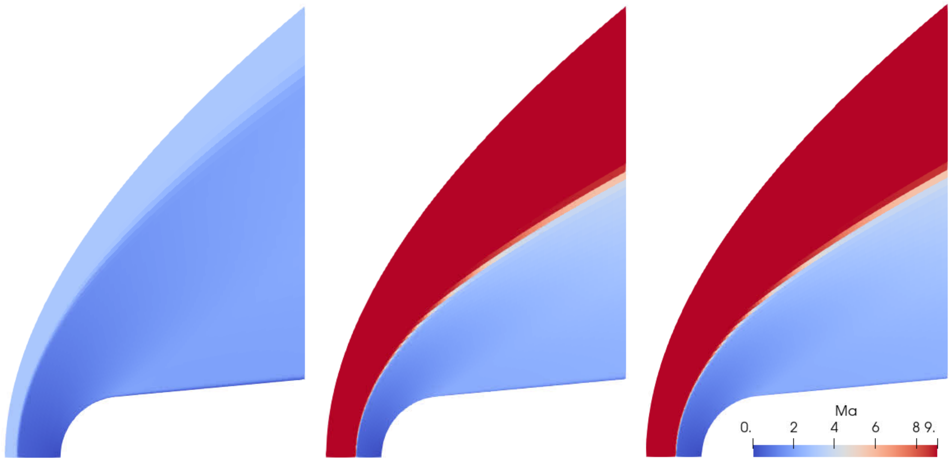

(a) $Ma_\infty = 3.0$, $\rho_\infty = 0.035 kg/m^3$ **(b)** $Ma_\infty = 9.0$, $\rho_\infty = 0.035 kg/m^3$ **(c)** $Ma_\infty = 9.0$, $\rho_\infty = 0.085 kg/m^3$

Fig. 4 Representative visualizations of the blunt wedge simulations for three choices of free-stream Mach number and density. The other freestream properties are those listed in Table 1. The figures are color-coded by the local Mach number Ma.

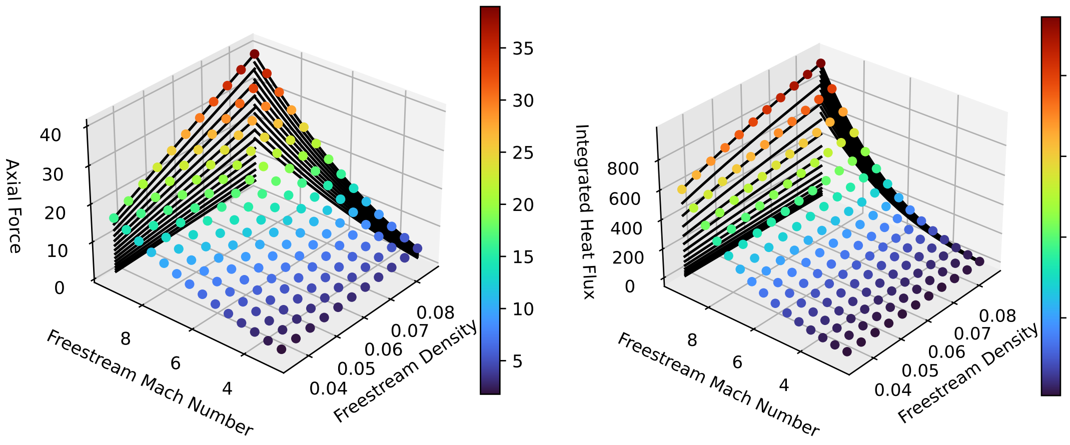

(a) Axial force F_x

(b) Integrated heat flux Q_{wall}

Fig. 5 Axial force and integrated heat flux as a function of freestream density ρ_∞ and Mach number Ma_∞ from 2D wedge.

Table 2 Grid tailoring parameters for the two-dimensional wedge case.

Parameter	1st tailoring	2nd tailoring
Number of buffer cells N_b	5	5
Maximum initial spacing $w_{0,max}$	None	None
Shock criterion S_c	0.01	0.01
Number of smoothing iterations N_{smooth}	10	5

(some) quantities of interest (QoIs) vary nonlinearly over the range of freestream densities and velocities. Specifically, it appears that F_x varies linearly with freestream density for a fixed freestream Mach number, but nonlinearly with freestream Mach number for a fixed freestream density.

A separate suite of simulations are performed with grid tailoring, where each simulation is grid tailored twice with the tailoring parameters shown in Table 2. The grid tailored simulations achieve convergence in less than 3000 iterations. Fig. 6 shows the Mach number flow field from a grid tailored simulation. The grid nodes are moved so that the bow shock location is aligned with the grid, resulting in a shock location at approximately 5 cells from the inlet. Some slight oscillations are observed along the inlet near the top of the domain because the bow shock location prior to tailoring is not a smooth field, and even after the smoothing operation some oscillations in the inlet location remain. Applying grid tailoring twice makes the inlet location smoother and aligns the grid with the bow shock better than just one grid tailoring, but some oscillations remain.

Fig. 7 shows the relative difference for axial force and integrated heat flux for each simulation. The maximum relative difference in axial force is 0.0013, while the maximum relative difference in integrated heat flux is 0.32. The results are in agreement with previous works observing that grid tailoring affects the surface heat flux more than pressure and shear stress [42].

1. Reduced-order model accuracy

The training set for the blunt wedge case is a set of 12 FOMs on a 3×4 grid in parameter space. Specifically, the parameters of the training set are the combinations of $\rho_\infty = 0.04,\ 0.06,\ \&\ 0.08\ kg/m^3$ and $Ma_\infty = 3.0,\ 5.0,\ 7.0,\ \&\ 9.0$. The test set consists of 142 other parameters combinations spaced on a regular grid in parameter space. All ROMs are run until the relative residual ℓ^2 norm falls below 10^{-5} or after 100 iterations.

As discussed in section III.A, the most popular basis for projection-based ROMs are POD modes computed over the entire training set as the basis. A localized basis can be more accurate by preventing data far away in parameter space from affecting a solution. This is particularly useful for simulations with hyperbolic governing equations, such as

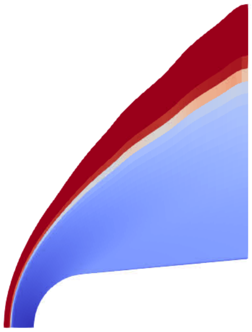

Fig. 6 Full-order simulation of blunt wedge with grid tailoring at $Ma_\infty = 9.0$, $\rho_\infty = 0.085kg/m^3$ The other freestream properties are those listed in table 1. The figure is color-coded by the local Mach number Ma.

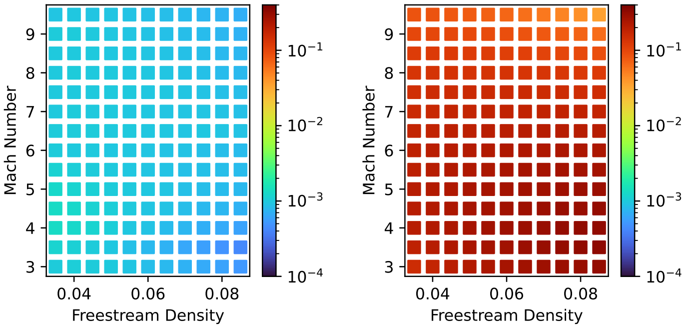

(a) Axial force F_x **(b) Integrated heat flux Q_{wall}**

Fig. 7 Relative difference in QoIs between grid tailored and non-tailored simulations, where the relative difference is defined in Eq. (29)**.**

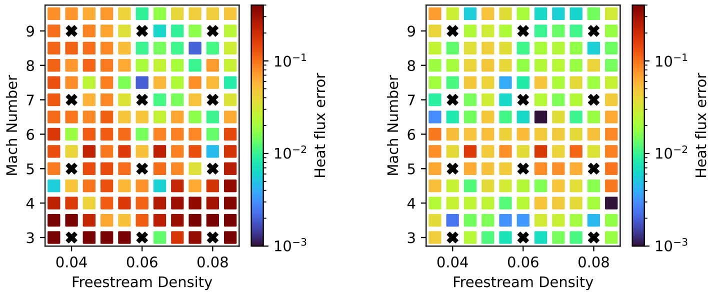

(a) POD basis **(b) Local dictionary basis**

Fig. 8 Heat flux $\varepsilon_{Q_{wall}}$ using global POD basis and local dictionary basis. Both cases use grid tailoring, 3 modes (p = 3), 4 nearest grid snapshots, conserved variables, LSPG, and RBF interpolation. Training cases are indicated by black X's.

Table 3 QoI errors from selected blunt wedge initial guesses (IG) and ROMs shown in Figs. 9 and 10 comparing grid-tailored and non-tailored ROMs. The percentage of ROMs in which a QoI is improved relative to the initial guess is abbreviated as %I. "% Converged" shows the percentage of cases that converged. The label <X% denotes the percentage of cases with errors below X%.

Grid	Model	% Converged	ε_{F_x}			$\varepsilon_{Q_{wall}}$		
			<5%	<1%	%I	<5%	<2%	%I
Tailored	IG		99%	72%	-	61%	20%	
Tailored	ROM	100%	100%	94%	68%	69%	30%	58%
Non-tailored	IG		65%	35%	-	44%	32%	
Non-tailored	ROM	87%	69%	27%	54%	28%	23%	37%

supersonic flow simulations [14]. We have observed that for the hypersonic cases tested here, using a local dictionary basis is generally more accurate than a global POD basis, consistent with the findings of Refs. [14, 15, 25]. Since the dictionary approach is not novel to this work, we show just one test case using a global POD basis. Fig. 8 shows a comparison of integrated heat flux error $\varepsilon_{Q_{wall}}$ for a case that uses 3 global POD modes and a case that uses a local dictionary with the 4 nearest snapshots (p = 3) as the basis. The local dictionary case has a lower error throughout most of parameter space. In the rest of this work, all the ROMs presented use a dictionary basis. Note that in cases where all the snapshots are used, the dictionary is a global dictionary and the basis spans the same subspace as an untruncated POD basis. In those cases, the untruncated POD basis has been observed to produce identical ROM results to the dictionary.

Grid Tailoring We demonstrate the effect of grid tailoring on the initial guess and ROMs in the following section. The ROM setups use conserved variables, LSPG, p = 8, and RBF interpolation.

We first examine the initial guesses of axial force shown in Figs. 9a and 9d, and it is obvious that the grid tailored surrogates have much higher accuracy. Interestingly, for the grid tailored initial guess, the axial force error along constant density (vertical lines) have very low error when there are training cases at the same density. The bow shock moves as the Mach number changes, so surrogates without grid tailoring cannot accurately resolve the bow shock, which can result in high errors. The heat flux error $\varepsilon_{Q_{wall}}$ from the initial guesses in Figs. 9b and 9e also show that interpolating over the grid tailored initial guesses is much more accurate than using the non-tailored initial guesses.

The first and third rows of Table 3 show data from the initial guesses. The grid tailored initial guess has 72% of cases with ε_{F_x} below 1%, while only 35% of cases with a non-tailored initial guess have ε_{F_x} below 1%. For the heat flux error $\varepsilon_{Q_{wall}}$, applying grid tailoring increases the number of cases with errors below 5% from 44% to 61%. Grid tailoring reduced the number of cases with $\varepsilon_{Q_{wall}}$ below 1%, but also drastically reduced the number of cases with high $\varepsilon_{Q_{wall}}$.

The state errors $\varepsilon_{\boldsymbol{x}}$ of the initial guesses with and without grid tailoring are shown in Figs. 9c and 9f. The state error is much lower for the grid tailored ROMs throughout almost the entire parameter space. However, it is important to note

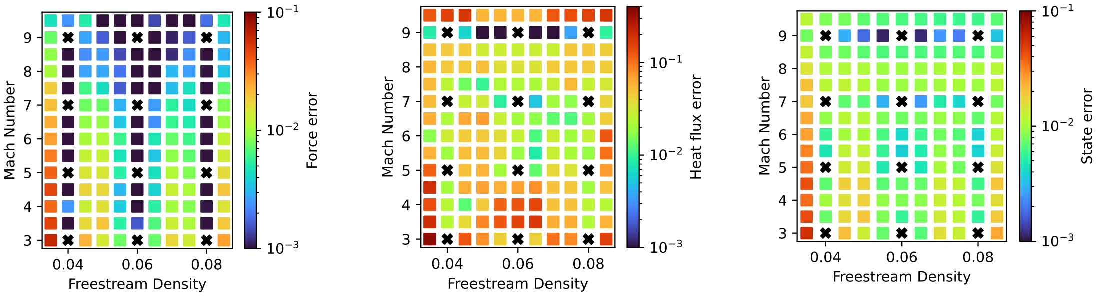

(a) ε_{F_x}**, Tailored** **(b)** $\varepsilon_{Q_{wall}}$**, Tailored** **(c)** ε_{x}**, Tailored**

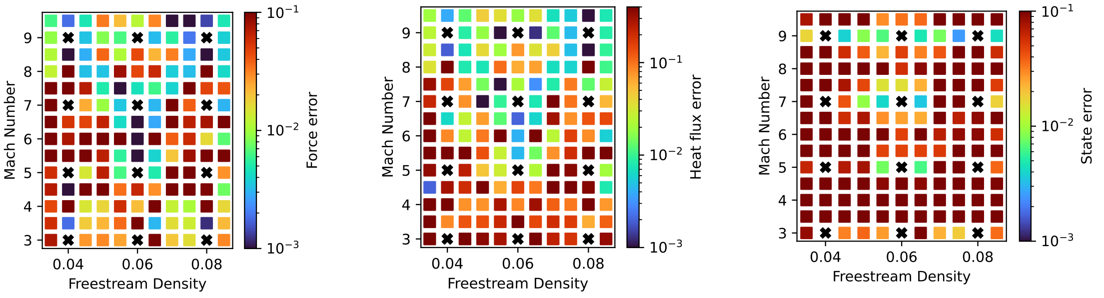

(d) ε_{F_x}**, Non-tailored** **(e)** $\varepsilon_{Q_{wall}}$**, Non-tailored** **(f)** ε_{x}**, Non-tailored**

Fig. 9 Blunt wedge errors ε_{F_x}**,** $\varepsilon_{Q_{wall}}$**, and** ε_{x} **for selected initial guesses demonstrating the effect of grid tailoring. All cases use conserved variables,** $p = 8$**, and RBF interpolation. Training cases are indicated by black X's.**

Table 4 State errors from selected blunt wedge initial guesses (IG) and ROMs shown in Table 3.

Grid	Model	% Converged	Median state error	
			Model	Proj.
Tailored	IG		0.0089	0.0046
Tailored	ROM	100%	0.0071	0.0046
Non-tailored	IG		0.1099	0.0055
Non-tailored	ROM	87%	0.0918	0.0055

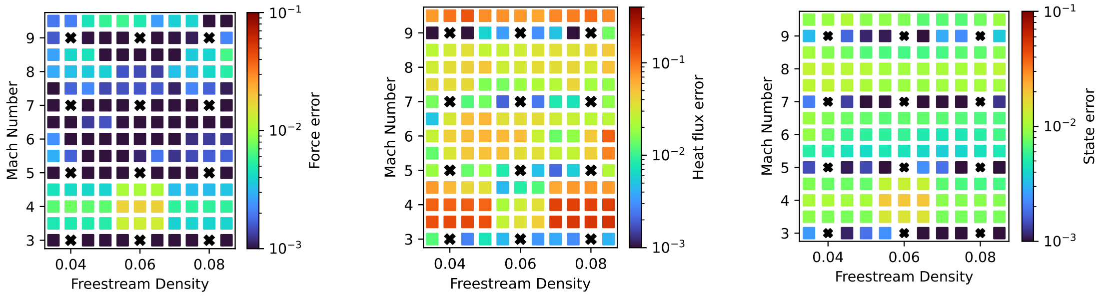

(a) ε_{F_x}, Tailored **(b) $\varepsilon_{Q_{wall}}$, Tailored** **(c) $\varepsilon_{\boldsymbol{x}}$, Tailored**

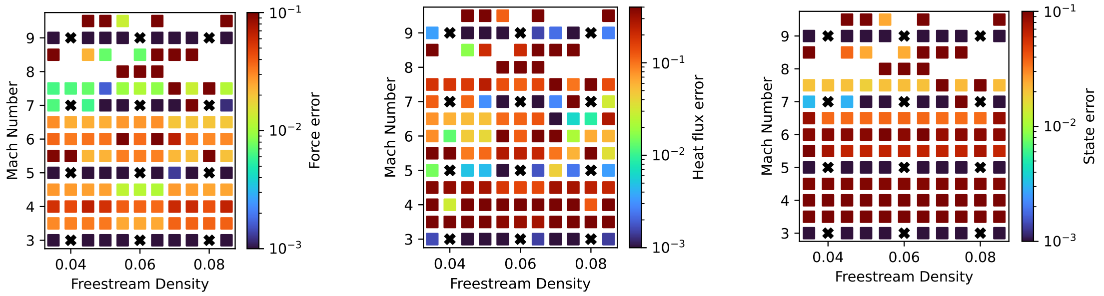

(d) ε_{F_x}, Non-tailored **(e) $\varepsilon_{Q_{wall}}$, Non-tailored** **(f) $\varepsilon_{\boldsymbol{x}}$, Non-tailored**

Fig. 10 Blunt wedge errors ε_{F_x}, $\varepsilon_{Q_{wall}}$, and $\varepsilon_{\boldsymbol{x}}$ for selected ROMs demonstrating the effect of grid tailoring. All cases use conserved variables, $p = 8$, LSPG, and RBF interpolation. Training cases are indicated by black X's. The cases with no error shown failed.

that the state error is not clearly correlated with either QoI error. Some correlations can be observed, particularly for force error, but we cannot in general assume that cases with low state error have low QoI errors. For this reason, we do not believe that state error alone is a good metric by which to analyze ROM accuracy.

The ROM QoI errors are next examined in Fig. 10 and the second and fourth rows of Tables 3. The axial force error ε_{F_x} of the grid tailored ROM in Fig. 10a is significantly lower than the initial guess errors in Fig. 9a throughout most of parameter space, particularly at medium and low Mach numbers. In contrast, the non-tailored ROM resulted in 13% of the high Mach number cases failing to converge. A number of cases with Mach numbers below 7.5 have lower errors than the initial guess, but almost all the high Mach number cases which converged still give high errors. The statistics in Table 3 show that 100% of the grid tailored ROMs have axial force errors ε_{F_x} below 1%, but only 27% of the non-tailored ROMs do. 69% of the grid tailored ROMs have heat flux error $\varepsilon_{Q_{wall}}$ below 5%, but only 28% of the non-tailored ROMs have this. 68% of the grid tailored ROMs have a lower axial force error than their initial guesses and 58% have a lower heat flux error, so the grid tailored ROMs are a generally an improvement over the initial guess. However, the non-tailored ROMs improve the QoI errors only 54% of the time for axial force and 37% of the time for heat flux, so it is generally better to just use the initial guess as a surrogate when grid tailoring is not applied.

Table 4 shows the median state errors of the model (initial guess or ROM) and the median projected state error. The median errors are used instead of the mean because a number of the ROM cases failed to converge, so the mean cannot be computed. The projected state error is the lowest possible state error achievable with a given basis. The median state error for both grid tailored and non-tailored ROMs is lower than the corresponding initial guess, but the median state error decreases only marginally with the non-tailored ROMs. This again shows that state error alone is not a great metric for ROM accuracy, it should be used in conjunction with assorted QoI errors.

Basis variable set We next demonstrate the effect of conservative and primitive variable bases. An examination of the initial guess errors in Fig. 11 shows that the primitive variable initial guess is more accurate than conserved variables at predicting all three errors. The axial force error is only slightly more accurate, but the heat flux error and state errors are significantly more accurate when a primitive variable basis is used. However, Fig. 12 shows that the conservative variable ROM has significantly lower axial force errors ε_{F_x}, while the primitive variable ROM has lower heat flux errors

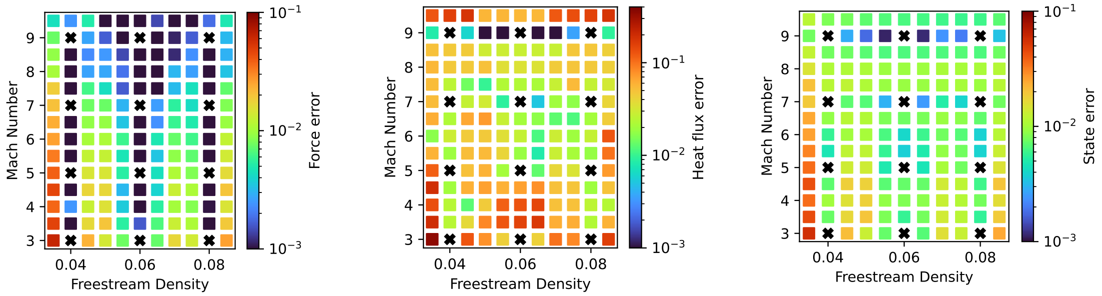

(a) ε_{F_x}, Conserved variables **(b) $\varepsilon_{Q_{wall}}$, Conserved variables** **(c) ε_{x}, Conserved variables**

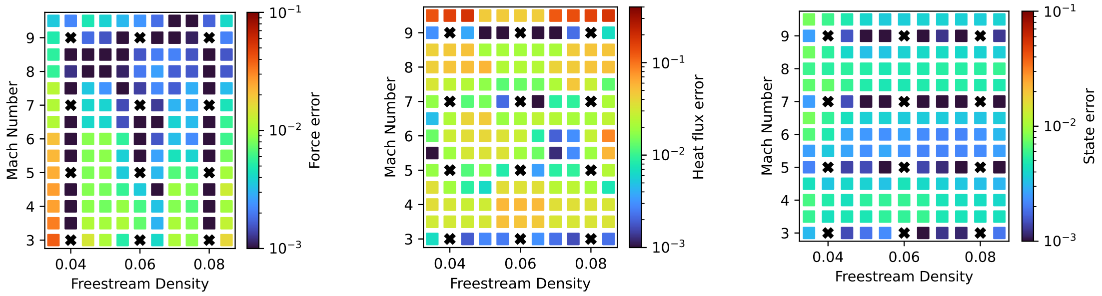

(d) ε_{F_x}, Primitive variables **(e) $\varepsilon_{Q_{wall}}$, Primitive variables** **(f) ε_{x}, Primitive variables**

Fig. 11 Blunt wedge errors ε_{F_x}, $\varepsilon_{Q_{wall}}$, and ε_{x} for selected initial guesses demonstrating conserved vs primitive variables. All cases use grid tailoring, $p = 8$, and RBF interpolation. Training cases are indicated by black X's.

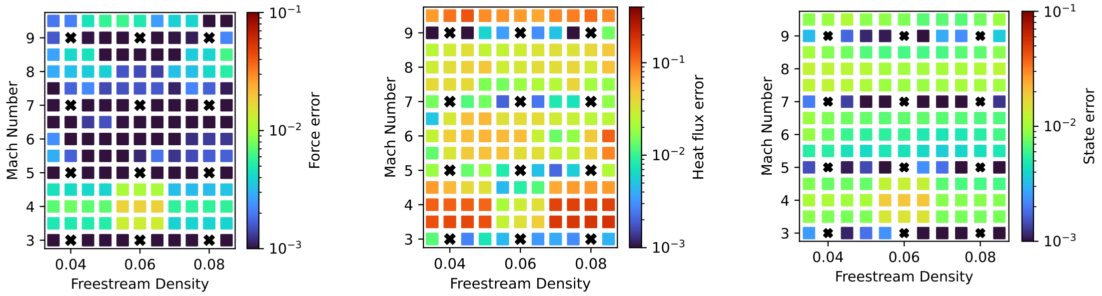

(a) ε_{F_x}, Conserved variables **(b) $\varepsilon_{Q_{wall}}$, Conserved variables** **(c) ε_{x}, Conserved variables**

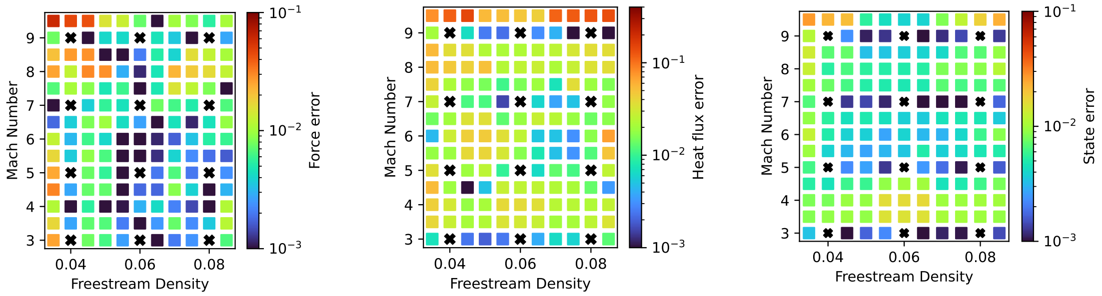

(d) ε_{F_x}, Primitive variables **(e) $\varepsilon_{Q_{wall}}$, Primitive variables** **(f) ε_{x}, Primitive variables**

Fig. 12 Blunt wedge QoI errors for selected ROM solutions demonstrating conserved vs primitive variables. All cases use grid tailoring, LSPG, $p = 8$, and RBF interpolation. Training cases are indicated by black X's.

Table 5 QoI errors from selected blunt wedge initial guesses (IG) and ROMs shown in Figs. 11 and 12 comparing primitive vs conserved variables. All cases use grid tailoring, LSPG, $p = 8$, and RBF interpolation. The percentage of ROMs in which a QoI is improved relative to the initial guess is abbreviated as %I. "% Converged" shows the percentage of cases that converged. The label <X% denotes the percentage of cases with errors below X%.

Variables	Model	% Converged	ε_{F_x}			$\varepsilon_{Q_{wall}}$		
			<5%	<1%	%I	<5%	<2%	%I
Conserved	IG		99%	72%	-	61%	20%	
Conserved	ROM	100%	100%	94%	68%	69%	30%	58%
Primitive	IG		100%	85%	-	89%	37%	
Primitive	ROM	100%	99%	77%	47%	90%	43%	61%

Table 6 State errors from selected blunt wedge initial guesses (IG) and ROMs shown in Table 5.

Variables	Model	% Converged	Median state error	
			Model	Proj.
Conserved	IG	99%	0.0089	0.0046
Conserved	ROM	100%	0.0071	0.0046
Primitive	IG	100%	0.0037	0.0032
Primitive	ROM	100%	0.0051	0.0032

$\varepsilon_{Q_{wall}}$ and state errors $\varepsilon_{\mathbf{x}}$. This shows that the choice of variables used for a basis in a ROM or surrogate model can have a major effect on the accuracy, but different QoIs may require different choices of variables for the highest accuracy. It may not always be possible to find a single ROM setup that is best for minimizing all errors, and aiming to minimize the state error can lead to a ROM setup with higher QoI errors.

The statistics in Table 5 reinforce this. The conserved variable ROM has the best accuracy at predicting axial force with 100% of cases having ε_{F_x} below 1%. However, the primitive variable ROM has 90% of cases having $\varepsilon_{Q_{wall}}$ below 5%, versus 69% for the conserved variable ROM. Interestingly, the primitive variable ROM has a slightly higher median state error than the initial guesses, as shown in Table 6. This can occur because there is a nonlinear relationship between the residual and the approximate state, so minimizing the residual does not guarantee that the state error decreases. Despite this, the axial force and heat flux error are reduced by the ROM in 47% and 61% of cases, respectively, showing again that judging the ROM on state error alone can be misleading.

Initial guess interpolation method A comparison of the three interpolation methods described in section III.C are shown in Fig. 13. When the axial force error ε_{F_x} from the initial guesses are compared (Figs. 13a, 13b, & 13c), it can be seen that RBF interpolation is the most accurate interpolation, followed by Kriging, and then inverse distance interpolation. In all three cases, the ROM improves ε_{F_x}; the RBF interpolation and Kriging lead to almost identical results, while the ROM solution from inverse distance interpolation is less accurate. Table 7 shows that the same trends hold for other QoIs. The heat flux error and state errors (see Table 8) are lowest for RBF interpolation and highest for inverse distance interpolation.

Kriging is a popular interpolation method that has been shown to be an accurate predictor for many datasets, but it generally requires large datasets to give accurate results [53]. It may be possible to increase the accuracy by using a different kernel function than the squared exponential kernel, but the optimal kernel would be problem-dependent. Inverse distance interpolation has been shown in a variety of studies to have poor accuracy in comparison with other methods [54, 55], but is a popular interpolation method due to its simplicity. However, using an inaccurate interpolation method can give misleading results about the relative accuracy of the ROM. For example, if we examine the percentage of cases in which the ε_{F_x} is improved by a ROM, Table 7 shows that 98% of cases with inverse distance interpolation have lower ROM error than the initial guess, whereas only 68% of cases with RBF interpolation are improved. Therefore, if ROM results are compared only to an initial guess using RBF interpolation, it may appear that the ROM improves

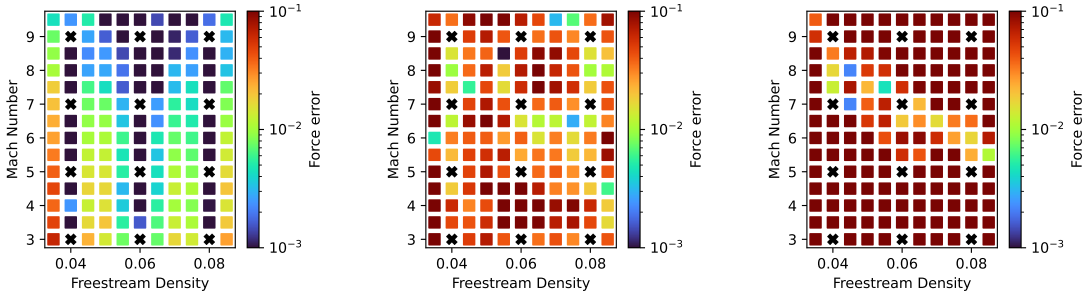

(a) Initial guess ε_{F_x}, RBF interpolation

(b) Initial guess ε_{F_x}, Kriging interpolation

(c) Initial guess ε_{F_x}, Inverse distance interpolation

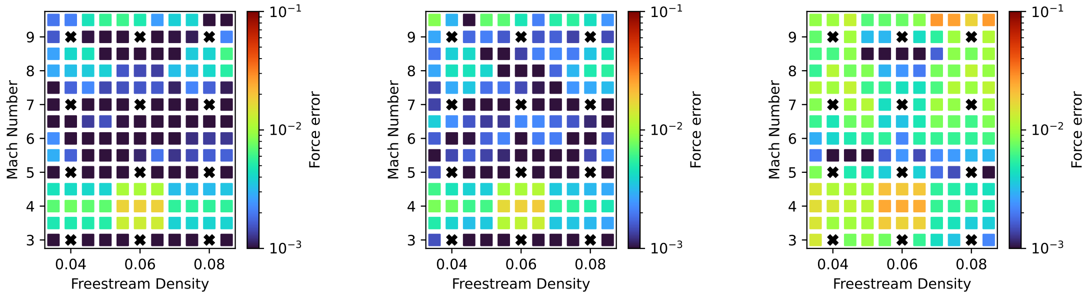

(d) ROM ε_{F_x}, RBF interpolation

(e) ROM ε_{F_x}, Kriging interpolation

(f) ROM ε_{F_x}, Inverse distance interpolation

Fig. 13 Blunt wedge axial force error ε_{F_x} for selected initial guesses and ROM solutions demonstrating various interpolation methods. All cases use grid tailoring, conserved variables, LSPG, and $p = 8$. Training cases are indicated by black X's.

Table 7 QoI errors from selected blunt wedge initial guesses (IG) and ROMs shown in Fig. 13 comparing interpolation methods. All cases use grid tailoring, conserved variables, LSPG, and $p = 8$. The percentage of ROMs in which a QoI is improved relative to the initial guess is abbreviated as %I. "% Converged" shows the percentage of cases that converged. The label <X% denotes the percentage of cases with errors below X%.

Interpolation	Model	% Converged	ε_{F_x}			$\varepsilon_{Q_{wall}}$		
			<5%	<1%	%I	<5%	<2%	%I
RBF	IG		99%	72%	-	61%	20%	
RBF	ROM	100%	100%	94%	68%	69%	30%	58%
Kriging	IG		61%	6%	-	51%	11%	
Kriging	ROM	100%	100%	94%	99%	30%	4%	24%
Inverse Dist.	IG		15%	2%	-	8%	3%	
Inverse Dist.	ROM	100%	100%	77%	98%	8%	3%	68%

Table 8 State errors from selected blunt wedge initial guesses (IG) and ROMs shown in Table 7.

Interpolation	Model	% Converged	Median state error	
			Model	Proj.
RBF	IG		0.0089	0.0046
RBF	ROM	100%	0.0071	0.0046
Kriging	IG		0.0428	0.0046
Kriging	ROM	100%	0.0073	0.0046
Inverse Dist.	IG		0.2172	0.0046
Inverse Dist.	ROM	100%	0.0079	0.0046

Table 9 Errors from selected blunt wedge initial guesses (IG) and ROMs shown in Figs. 14 and 15. All cases use grid tailoring, LSPG, conserved variables, and RBF interpolation. The percentage of ROMs in which a QoI is improved relative to the initial guess is abbreviated as %I. % Converged shows the percentage of cases that converged. The label <X% denotes the percentage of cases with errors below X%.

p	Model	% Converged	ε_{F_x}			$\varepsilon_{Q_{wall}}$			Median state error	
			<5%	<1%	%I	<5%	<2%	%I	Model	Proj.
3	IG		80%	37%	-	66%	33%		0.0173	0.0063
3	ROM	100%	100%	94%	80%	84%	38%	58%	0.0103	0.0063
8	IG		99%	72%	-	61%	20%		0.0089	0.0046
8	ROM	100%	100%	94%	68%	69%	30%	58%	0.0071	0.0046
11	IG		99%	75%	-	72%	28%		0.0090	0.0040
11	ROM	100%	100%	100%	68%	69%	38%	57%	0.0074	0.0040

errors by a large margin, whereas a well-designed surrogate model might have lower errors than the ROM in much of parameter space.

Basis size The last study we perform with the blunt wedge examines the impact of the basis size on accuracy. Fig. 14 shows plots of ε_{F_x} and $\varepsilon_{Q_{wall}}$ from initial guess surrogates using $p = 3,\ 8,\ \&\ 11$. The ε_{F_x} is highest for $p = 3$ and very similar for $p = 8$ and $p = 11$. However, it is unclear from Fig. 14 which surrogate is most accurate at predicting $\varepsilon_{Q_{wall}}$. The surrogates with $p = 8$ and $p = 11$ result in high errors for extrapolative cases, but the surrogate with $p = 3$ has a number of cases with high errors for some interpolative cases.

The statistics in Table 9 show the initial guess with $p = 11$ is marginally more accurate for ε_{F_x} than $p = 8$. However, the statistics for $\varepsilon_{Q_{wall}}$ do not present an obvious best case. The initial guess with $p = 11$ has the highest number of cases with $\varepsilon_{Q_{wall}} < 5\%$, but $p = 3$ has the highest number of cases with $\varepsilon_{Q_{wall}} < 2\%$.

The ROM solutions in Fig. 15 show that a large basis is best for predicting ε_{F_x}, but the smallest basis is most accurate at predicting $\varepsilon_{Q_{wall}}$. Table 9 backs up these conclusions, showing that 100% of cases with $p = 11$ have ε_{F_x} below 1%, and 84% of cases with $p = 3$ have $\varepsilon_{Q_{wall}}$ below 5%.

Flow visualizations Finally, we consider flow visualizations of the ROM results for the wedge at $Ma_\infty = 8.0$, $\rho_\infty = 0.035 kg/m^3$. These conditions are ones for which the grid tailored ROM is accurate, 1-2% error for both QoIs. This is more accurate than the initial guess, for which $\varepsilon_{F_x} \approx 1 - 2\%$ and $\varepsilon_{Q_{wall}} \approx 6 - 7\%$.

Fig. 16 shows that the tailored grid for the ROM matches the grid computed by the FOM very well. Additionally, the flow fields are very similar for the FOM, ROM and RBF initial guess, although it appears that the freestream Mach number is slightly under-predicted by RBF and the ROM. Without grid tailoring, the ROM slightly overpredicts the freestream Mach number, and there is a large spike in Mach number just upstream of the rounded tip of the wedge. This spike is not present in the grid tailored ROM, an example of the relative robustness of the grid tailored ROM.

Differences between the ROM with and without grid tailoring can also be seen in the surface quantities. Fig. 17

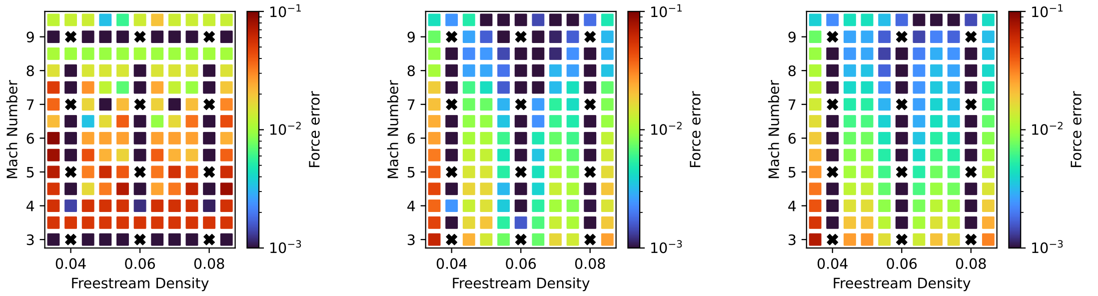

(a) ε_{F_x}, $p = 3$ **(b)** ε_{F_x}, $p = 8$ **(c)** ε_{F_x}, $p = 11$

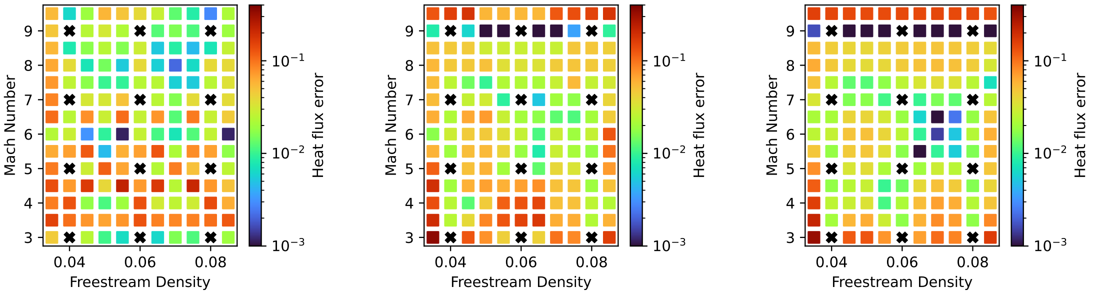

(d) $\varepsilon_{Q_{wall}}$, $p = 3$ **(e)** $\varepsilon_{Q_{wall}}$, $p = 8$ **(f)** $\varepsilon_{Q_{wall}}$, $p = 11$

Fig. 14 Blunt wedge QoI errors ε_{F_x} and $\varepsilon_{Q_{wall}}$ for selected initial guesses demonstrating effect of number of modes. All cases use grid tailoring, LSPG, conserved variables, and RBF interpolation. Training cases are indicated by black X's.

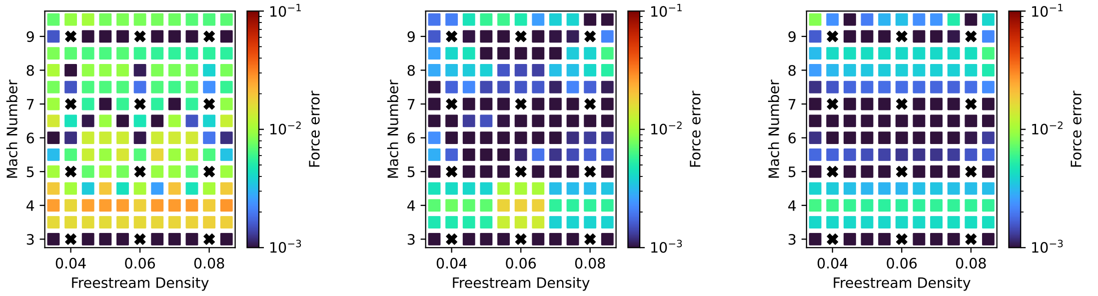

(a) ε_{F_x}, $p = 3$ **(b)** ε_{F_x}, $p = 8$ **(c)** ε_{F_x}, $p = 11$

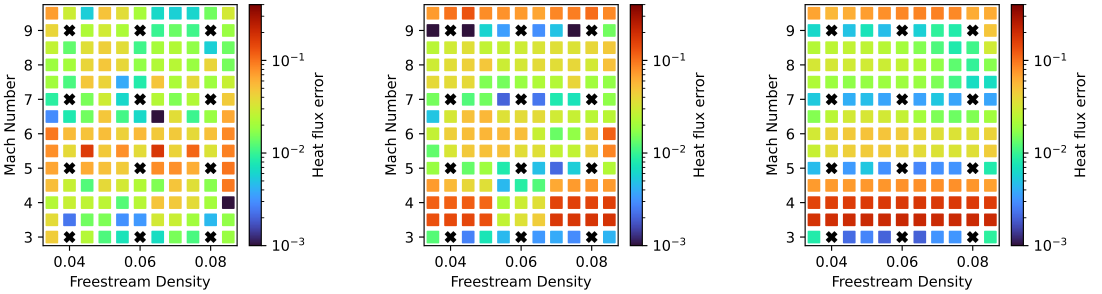

(d) $\varepsilon_{Q_{wall}}$, $p = 3$ **(e)** $\varepsilon_{Q_{wall}}$, $p = 8$ **(f)** $\varepsilon_{Q_{wall}}$, $p = 11$

Fig. 15 Blunt wedge QoI errors ε_{F_x} and $\varepsilon_{Q_{wall}}$ for selected ROMs demonstrating effect of number of modes. All cases use grid tailoring, LSPG, conserved variables, and RBF interpolation. Training cases are indicated by black X's.

(a) FOM **(b) IG** **(c) ROM**

(d) Grid tailored FOM **(e) Grid tailored IG** **(f) Grid tailored ROM**

Fig. 16 Mach contours for the wedge at $Ma_\infty = 8.0$, $\rho_\infty = 0.035 kg/m^3$. ROM results are for the $p = 3$ LSPG ROM with a conserved variable basis and RBF initial guess. IG is short for "initial guess".

Table 10 Average simulation times and ROM speedup

Configuration	Simulation	Number of cores	Average time	Speedup
Blunt Wedge	FOM	16	101.0 s	-
	ROM	16	3.7 s	27.5

shows heat flux on the wedge surface, along with the heat flux error of the ROMs and initial guesses. We see that the grid tailored ROM is more accurate than the initial guess over the entire surface. The improvement in accuracy is largest for heat flux, the ROM reduces the maximum error, near the front of the wedge, from around 8% to 3%. The ROM reduces the error below 1% on much of the wedge, while the initial guess error is around 4-5% in the same region. In contrast, the ROM without tailoring is considerably less accurate than the initial guess; it exceeds 10% over the entire surface while the initial guess was around 5% on most of the surface.

Overall, these visualizations show the ability of the grid tailored ROM to compute flow fields and surface quantities to higher levels of accuracy than a simple interpolation like RBF. They also show the importance of using grid tailoring in conjunction with residual minimization; if the shock is not properly resolved, the ROM provides a poorer approximation of the surface quantities on the wedge. This is because grid tailoring positions the shock accurately, lowering the residuals near the shock wave. This causes the residual minimization at the heart of LSPG to focus more on reducing residuals in other regions of the flow, such as near the wall.

2. Reduced-order model performance

The average simulation times and ROM speedup are given in Table 10. The ROM speedup is defined as the ratio of the FOM simulation time to the ROM simulation time. The ROMs in Table 10 have the setups identified to produce the best results, specifically a conserved variable basis with 3 modes, a LSPG solver, and RBF interpolation for the initial guess.

The reasons a speed-up is obtained despite the lack of hyper-reduction is mostly due to the use of a direct nonlinear solver in the ROM in place of a pseudo-time-stepping approach like the one used for the FOM, as discussed in section III.B. The LSPG ROM only requires 25 nonlinear iterations at most for the cases considered, while the FOM

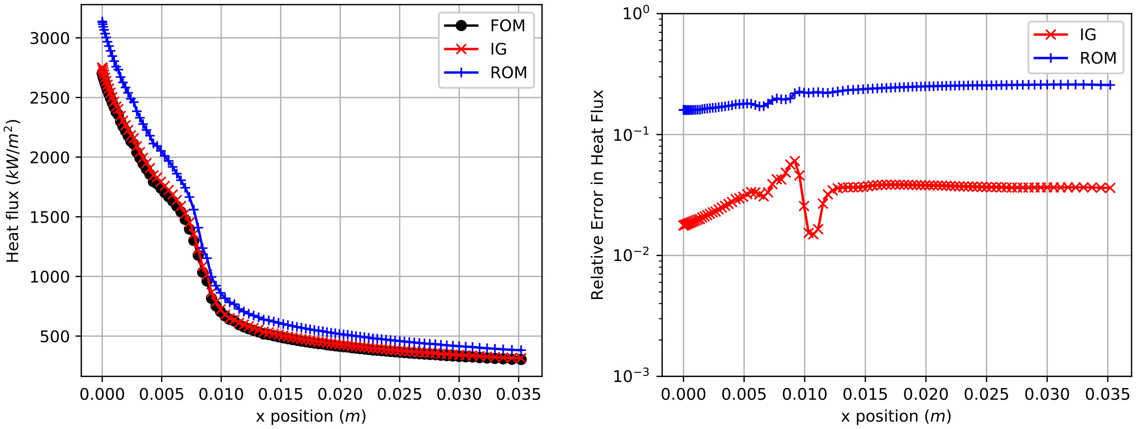

(a) Not tailored

(b) Not tailored

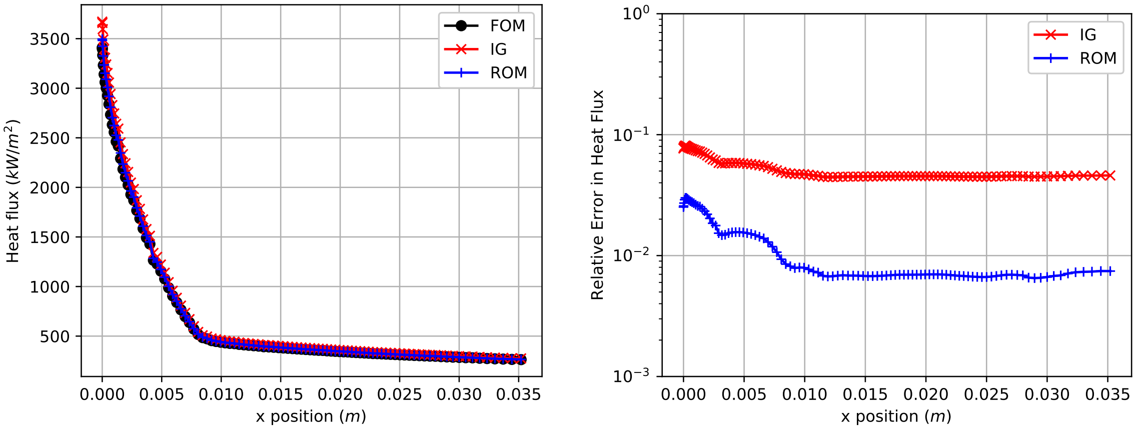

(c) Tailored

(d) Tailored

Fig. 17 Surface heat flux and corresponding errors for the wedge at $Ma_\infty = 8.0$, $\rho_\infty = 0.035 kg/m^3$ and the $p = 3$ LSPG ROM with a conserved variable basis and RBF initial guess. IG is short for "initial guess".

Table 11 Free stream flow conditions for the two-dimensional HIFiRE-1.

Density	0.040 to 0.080 kg/m^3
Velocity	1622.91 to 2803.22 m/s
Mach Number	5.5 to 9.5
Angle of attack	0.0°
Temperature	226.46 K
Reynolds Number	1.5 $\times 10^7$ $1/m$

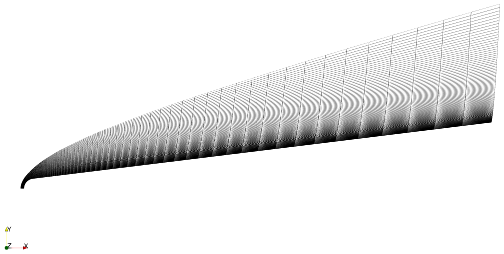

Fig. 18 Baseline mesh for 2D HiFIRE without grid tailoring.

requires 400-500 iterations per grid, or 1,200-1,500 iterations total. Fewer nonlinear solver iterations means fewer evaluations of the residual $\boldsymbol{f}$ and the corresponding Jacobian $\partial \boldsymbol{f}/\partial \boldsymbol{x}$, the cost of which scale with state size N, so any reduction in residual and Jacobian evaluations will substantially reduce the cost of the ROM.

3. Summary

The following is a summary of lessons learned from the blunt wedge case:

- Grid-tailored ROMs have better convergence and accuracy than non-tailored ROMs
- Conserved variable ROMs achieve higher accuracy at predicting axial force, while primitive variable ROMs have more accurate heat flux predictions.
- A large basis achieves higher accuracy for axial force, but a local basis may be better for predicting heat flux.

These conclusions hold only for the blunt wedge case; we find in examining other cases that some of the conclusions change based on the specific case.

D. 2D HIFiRE-1 Nose

The second flow configuration is a two-dimensional simulation of the HIFiRE-1 (Hypersonic International Flight Research Experiment) nose cone [56]. The nose cone has a 7° taper angle with a rounded nose and a length of 11.8 cm from the nose to the end of the cone. The geometry is axisymmetric with zero angle of attack, so the simulations require a mesh only on one side of the centerline, shown in Fig. 18. The 2D HiFIRE-1 mesh is much more refined in the direction away from the wall in comparison to the 2D wedge case. The mesh has 32,512 cells, but the spacing in the wall tangent direction is coarser, so there are more cells between the body and the inlet. Simulations without grid tailoring have a maximum y+ of 0.81.

Simulations are run at the range of flight conditions shown in Table 11.These conditions include the baseline case for which the grid was designed, run 30 of the experimental results presented in Ref. [57]. Two suites of FOMs are run at various freestream Mach numbers and densities, one set with no grid tailoring, and one where simulations are tailored twice, with the tailoring parameters given in Table 12. In both cases, the FOM is run until the relative residual is reduced by 7 orders of magnitude. Simulations are run using the S-A turbulence model [58], so there are six conserved variables.

Fig. 19 shows four simulations at two different Mach numbers with and without grid tailoring. The shock location moves slightly closer to the surface as the Mach number is increased and grid tailoring moves the inlet closer to the shock, but the shock is reasonably well aligned with the grid even without tailoring. The grid tailoring parameters used

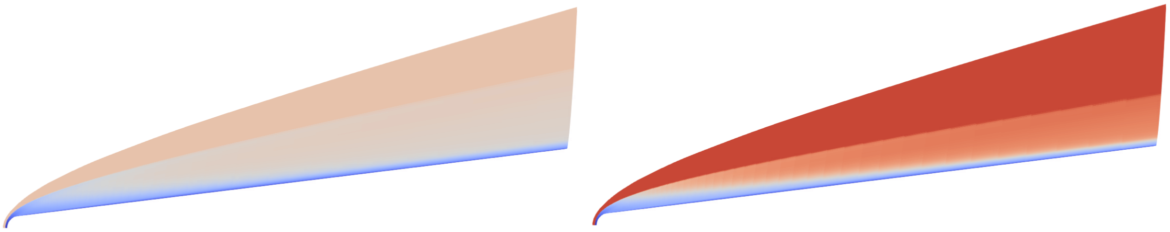

(a) $Ma = 6.0$, **non-tailored** (b) $Ma = 9.0$, **non-tailored**

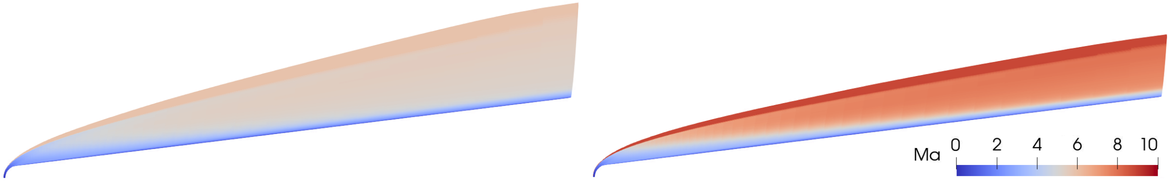

(c) $Ma = 6.0$, **tailored** (d) $Ma = 9.0$, **tailored**

Fig. 19 Two-dimensional HIFiRE-1 simulations with and without grid tailoring with varying inlet Mach numbers. Simulations are run at $\rho_\infty = 0.06$ kg/m^3. Figures are colored by Mach number.

Table 12 Grid tailoring parameters for the two-dimensional HIFiRE-1.

Parameter	1st tailoring	2nd tailoring
Number of buffer cells N_b	10	10
Maximum initial spacing $w_{0,max}$	10^{-6} m	10^{-6} m
Shock criterion S_c	0.01	0.01
Number of smoothing iterations N_{smooth}	5	5

are shown in Table 12.

Fig. 20 shows the relative differences between the tailored and non-tailored FOMs. The maximum relative difference in axial force is 0.00039, while the maximum relative difference in integrated heat flux is 0.0037. The QoIs are almost identical in the grid-tailored and non-tailored simulations because the meshes are very refined and the shock is well aligned with the grid prior to any tailoring. However, the differences can be significant close to the stagnation point. For example, the relative difference of the maximum heat flux on the surface is between 0.0017 and 0.129. The location of maximum surface heat flux occurs on the blunt nose close to the stagnation point. Since the blunt nose is small, errors in maximum surface heat flux do not significantly affect the integrated heat flux errors.

Fig. 21 shows the axial force and integrated heat flux at various freestream Mach numbers and densities. There are obvious nonlinearities in the integrated heat flux, and the axial force also exhibits nonlinear behavior.

1. Reduced-order model accuracy

The accuracy of the ROM is studied by considering a small, four snapshot training set comprised of FOM solutions near each corner of parameter space at $Ma = 6.5,\ 8.5$ and $\rho = 0.048,\ 0.072\ kg/m^3$. Since the training set has 4 snapshots, a POD basis with $p = 3$ will have no information compression; it can reconstruct all 4 snapshots with machine precision accuracy. Because of this, the ROM results are virtually identical for POD and the dictionary with $p = 3$. In this case, this means that the dictionary is the superior option, since it does not require the offline computation of a singular value decomposition.

The ROM errors presented in Table 13 reveal some similarities but several key differences to the behaviors observed for the 2D wedge case. The main similarity is that RBF is overall the most accurate interpolant in both state error and QoI. However, kriging is slightly more accurate for heat flux; for the same grid tailored primitive variable basis, 42% of the kriging cases have errors less than 5%, versus 38% for RBF. As observed for the 2D wedge case, inverse distance interpolation is substantially less accurate than either Kriging or RBF.

Unlike the 2D wedge case, grid tailoring does not make a substantial difference in the QoI accuracy of the initial

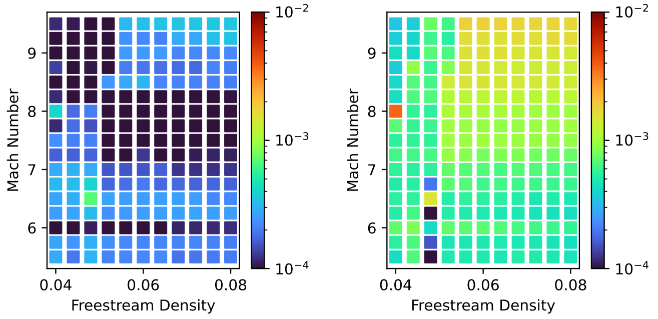

(a) Axial force error **(b) Integrated heat flux**

Fig. 20 Axial force and integrated heat flux percentage differences between non-tailored and tailored FOMs from 2D HIFiRE-1.

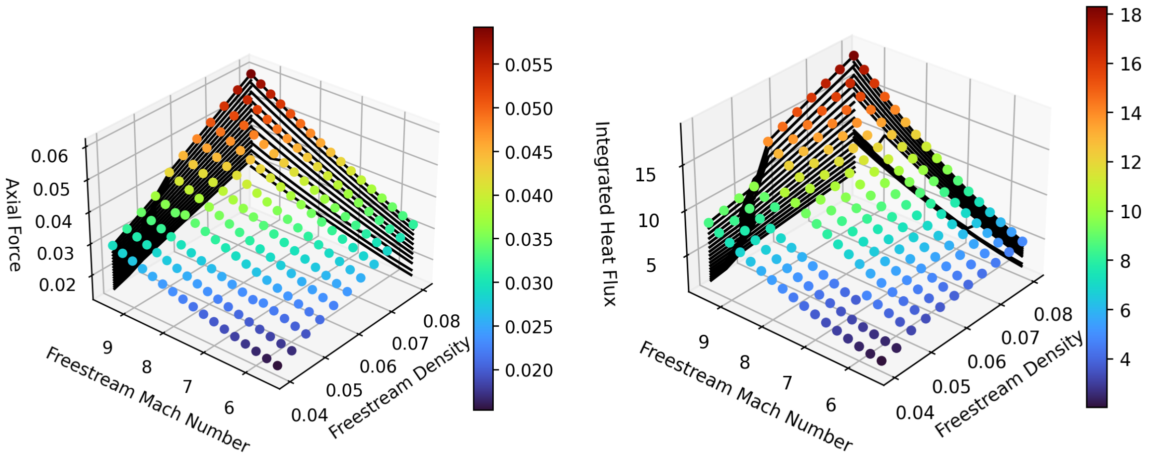

(a) Axial force **(b) Integrated heat flux**

Fig. 21 Axial force and integrated heat flux from 2D HiFIRE-1 tailored simulations.

Table 13 Errors from 2D HIFiRE-1 initial guesses.

Grid	Variables	Interpolation	ε_{F_x}		$\varepsilon_{Q_{wall}}$		Median
			<5%	<1%	<5%	<2%	state error
Tailored	Conserved	RBF	96%	29%	26%	8%	0.0264
Tailored	Primitive	RBF	99%	40%	38%	16%	0.0204
Tailored	Primitive	Inverse dist.	19%	3%	13%	5%	0.0716
Tailored	Primitive	Kriging	84%	26%	42%	19%	0.0254
Non-tailored	Conserved	RBF	96%	28%	22%	7%	0.0481
Non-tailored	Primitive	RBF	98%	38%	39%	17%	0.0310

Table 14 QoI errors from 2D HIFiRE-1. The percentage of ROMs in which the QoI is improved relative to the initial guess is abbreviated as %I.

Grid	Variables	ROM	Interpolation	% Converged	ε_{F_x}			$\varepsilon_{Q_{wall}}$		
					<5%	<1%	%I	<5%	<2%	%I
Tailored	Conserved	LSPG	RBF	100%	38%	15%	16%	30%	14%	45%
Tailored	Primitive	LSPG	RBF	100%	100%	50%	55%	44%	20%	61%
Tailored	Primitive	LSPG	Inverse Dist.	100%	86%	37%	93%	42%	20%	81%
Tailored	Primitive	LSPG	Kriging	100%	100%	49%	71%	45%	16%	54%
Non-tailored	Primitive	LSPG	RBF	100%	73%	17%	19%	22%	5%	23%
Non-tailored	Conserved	LSPG	RBF	100%	28%	7%	4%	17%	9%	30%

Table 15 State errors from 2D HIFiRE-1. The percentage of ROMs in which the QoI is improved relative to the initial guess is abbreviated as %I.

Grid	Variables	ROM	Interpolation	% Converged	State Error	
					Model	Proj.
Tailored	Conserved	LSPG	RBF	100%	0.0164	0.0107
Tailored	Primitive	LSPG	RBF	100%	0.0185	0.0135
Tailored	Primitive	LSPG	Inverse Dist.	100%	0.0214	0.0135
Tailored	Primitive	LSPG	Kriging	100%	0.0187	0.0135
Non-tailored	Primitive	LSPG	RBF	100%	0.0362	0.0267
Non-tailored	Conserved	LSPG	RBF	100%	0.0693	1.9092

guess interpolations. The number of cases with low QoI errors ε_{F_x} and $\varepsilon_{Q_{wall}}$ are within 1-2% of one another for tailored and non-tailored cases, independent of the basis variables used. However, the median state error is relatively large for both primitive and conserved variable bases without tailoring. This discrepancy in QoI errors and state errors is because the state error includes the error in approximating the bow shock wave, while the QoI errors are only impacted by the flow field near the vehicle surface. Another difference with the 2D wedge case is that the primitive variable basis is slightly more accurate than the conserved variable basis for axial force. It is substantially more accurate for heat flux, with 38% of cases with errors under 5% versus 29% for the conserved variable basis.

As seen for the initial guesses, the results in Tables 14 and 15 are similar to those for the 2D wedge in a few respects. One similarity is that C-LSPG is not as consistently accurate or robust as LSPG for this case. As observed for the 2D wedge, the ROM residual does not converge below the targeted reduction of 6 orders of magnitude for many cases. Because of this, most C-LSPG ROMs are less accurate than the corresponding initial guess. A second similarity is that RBF is still the best initial guess interpolation, leading to higher accuracy in the initial guess itself as well as the ROM. Additionally, a kriging initial guess leads to ROMs of similar accuracy to the RBF initial guess, but inverse distance initial guesses lead to less accurate ROMs.

The main difference with the 2D wedge case is that the primitive variable basis is far better for the ROMs. Although both bases result in ROMs with lower median state errors $\varepsilon_{\boldsymbol{x}}$ than the initial guess, the primitive variable ROM is much more accurate in ε_{F_x}, 100% versus 38% of cases with less than 5% error, and somewhat more accurate in $\varepsilon_{Q_{wall}}$, 44% versus 30% of cases with less than 5% error. The primitive variable ROM increases accuracy over the initial guess for both QoIs over 50% of the time, while the conserved variable ROM only improves on the initial guess in 16% of cases for ε_{F_x} and 45% of cases for $\varepsilon_{Q_{wall}}$. Taken together, these numbers make the case that the primitive variable basis leads to more accurate ROMs than the conserved variable basis for the 2D HIFiRE-1.

A number of error surface plots are included to reinforce and elaborate on the takeaways from the tables. Fig. 22 shows RBF initial guess state errors $\varepsilon_{\boldsymbol{x}}$. We focus on RBF since it is the most accurate interpolant for this case. As in Table 13, these plots show that $\varepsilon_{\boldsymbol{x}}$ is lower for the primitive variable basis, and that grid tailoring leads to lower state errors. However, the difference in state error level is smaller for the primitive variable basis than the conserved variable

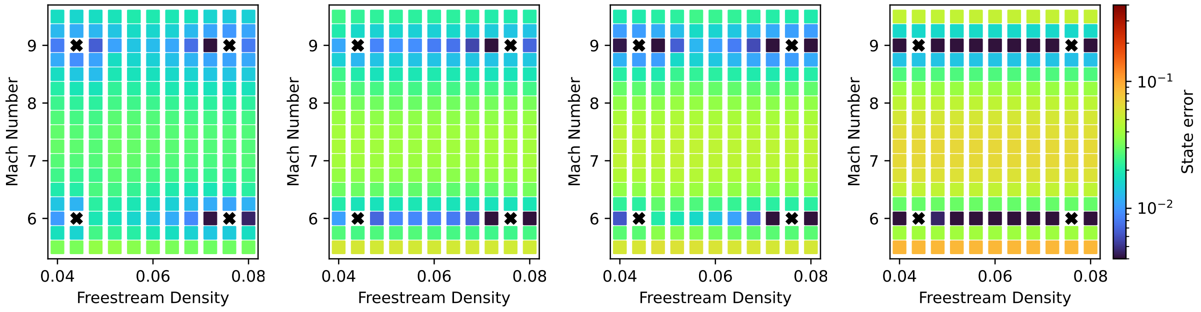

(a) Tailored, Primitive variables

(b) Tailored, Conserved variables

(c) Non-tailored, Primitive variables

(d) Non-tailored, Conserved variables

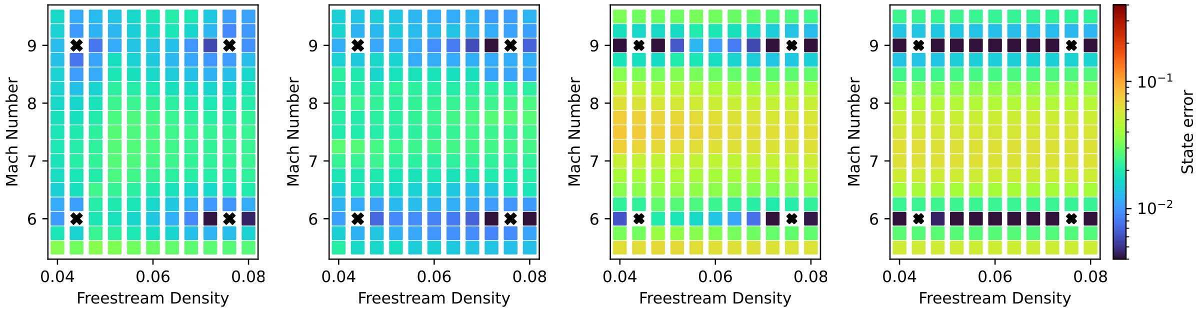

(e) Tailored, Primitive variables, LSPG

(f) Tailored, Conserved variables, LSPG

(g) Non-tailored, Primitive variables, LSPG

(h) Non-tailored, Conserved variables, LSPG

Fig. 22 2D HIFiRE-1 State errors ε_x for selected initial guesses and ROM solutions. All cases use RBF interpolation. Training cases are indicated by black X's.

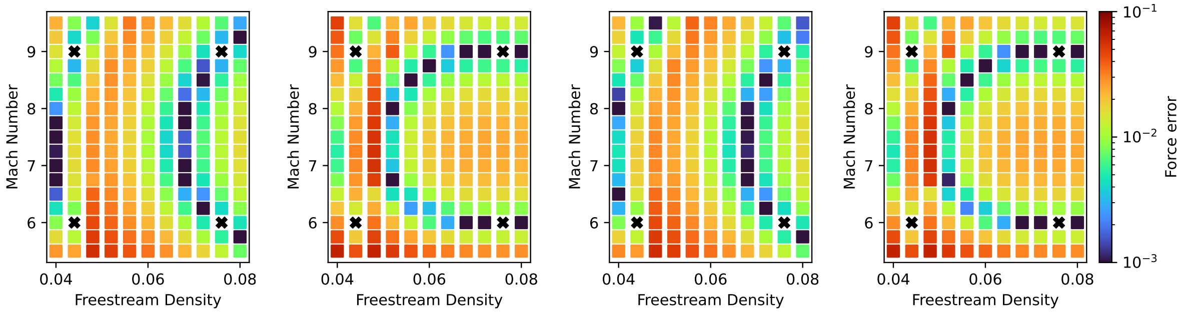

(a) Tailored, Primitive variables

(b) Tailored, Conserved variables

(c) Non-tailored, Primitive variables

(d) Non-tailored, Conserved variables

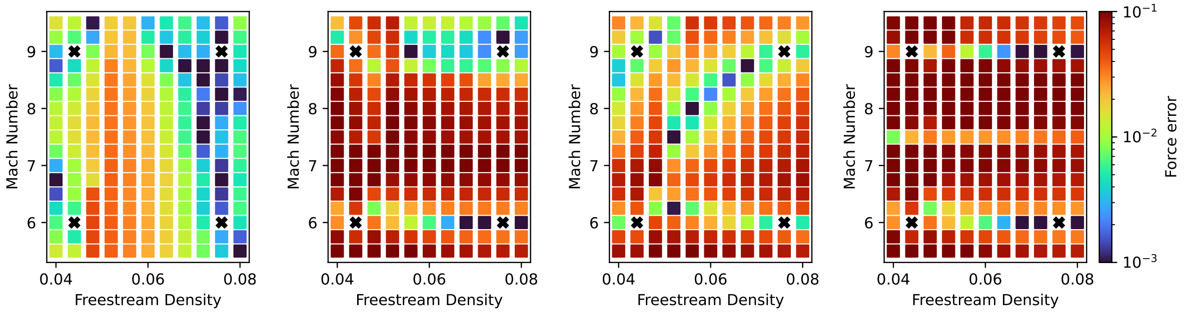

(e) Tailored, Primitive variables, LSPG

(f) Tailored, Conserved variables, LSPG

(g) Non-tailored, Primitive variables, LSPG

(h) Non-tailored, Conserved variables, LSPG

Fig. 23 2D HIFiRE-1 axial force errors ε_{F_x} for selected initial guesses and ROM solutions. All cases use RBF interpolation. Training cases are indicated by black X's.

basis.

Figure 22 also shows how the $\varepsilon_{\boldsymbol{x}}$ varies throughout parameter space. All cases have lower $\varepsilon_{\boldsymbol{x}}$ near the training points, and $\varepsilon_{\boldsymbol{x}}$ are at their largest near the minimum Mach number of 5.5 and in between the training Mach numbers of 6.0 and 9.0. For a given Mach number, $\varepsilon_{\boldsymbol{x}}$ is roughly constant, independent of density. This is because the bow shock wave location changes with Mach number, not density, and therefore the linear basis is best suited to represent states at the same Mach number as the training data. It is interesting that this same trend is observed with and without grid tailoring; this shows that grid tailoring reduces the error at cases with Mach numbers that differ from the training set, but does not entirely eliminate the limited accuracy with which linear bases resolve shockwaves.

Fig. 22 also shows the corresponding ROM state errors. These plots show how the error is the same or less than the initial guess when using LSPG. The grid tailored LSPG ROMs are more accurate than the non-tailored LSPG ROMs. Grid tailored LSPG ROMs with primitive and conserved variable bases have similar or improved $\varepsilon_{\boldsymbol{x}}$ for most cases. In particular, $\varepsilon_{\boldsymbol{x}}$ is lower for the ROMs outside of the training set, showing the ability of the ROM to extrapolate.

Visualizing QoI error surfaces reveal interesting trends in where the ROM is accurate. Fig. 23 shows that the RBF interpolation only computes accurate axial forces in a few areas of parameter space. The most accurate predictions occur in two C-shaped troughs, one occuring at low freestream density, and one at higher densities. There does not appear to be any obvious correlation between the $\varepsilon_{\boldsymbol{x}}$ and ε_{F_x} fields, showing that ROM accuracy cannot be judged by state error alone. Finally, as observed for the state error, grid tailoring does not make a large difference for the initial guess.

Fig. 23 shows that grid tailored primitive ROMs are more accurate than the initial guess at higher densities, higher mach numbers, and at the lowest densities and Mach numbers. This again shows the predictive capability of the ROM, in the form of high accuracy at the edge or outside of the training set. In a region encompassing densities of 0.05 and 0.06 and between Mach 5.5 and 8.5, it appears that ROM is not much better or slightly less accurate than the initial guess. Note that this also corresponds to the region with the highest QoI errors, including some near 5%. Overall, the grid tailored LSPG ROM with a primitive variable basis was the most accurate relative to the initial guesses and the other ROMs.

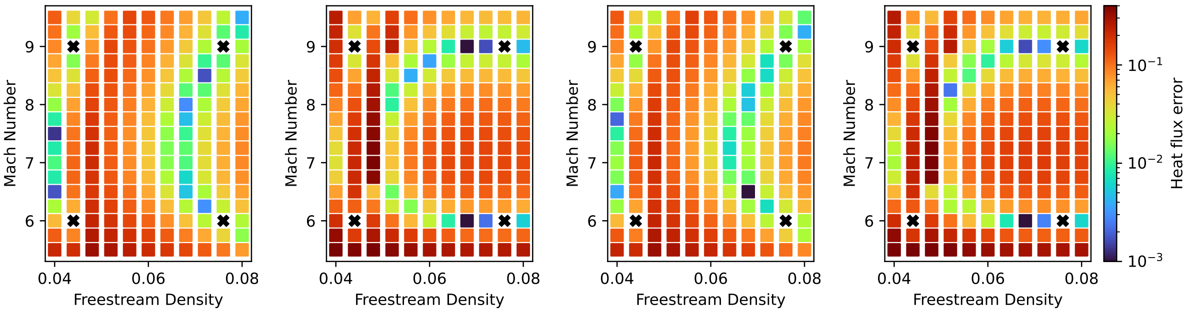

(a) Tailored, Primitive variables

(b) Tailored, Conserved variables

(c) Non-tailored, Primitive variables

(d) Non-tailored, Conserved variables

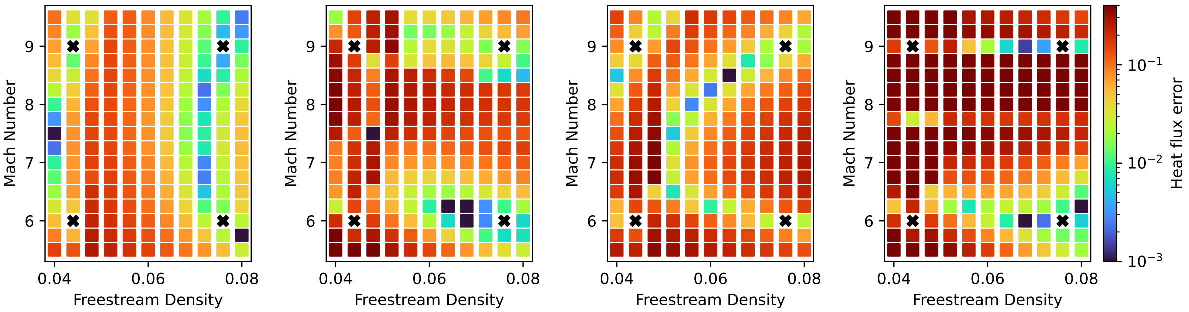

(e) Tailored, Primitive variables, LSPG

(f) Tailored, Conserved variables, LSPG

(g) Non-tailored, Primitive variables, LSPG

(h) Non-tailored, Conserved variables, LSPG

Fig. 24 2D HIFiRE-1 heat flux errors $\varepsilon_{Q_{wall}}$ for selected initial guesses and ROM solutions. All cases use RBF interpolation. Training cases are indicated by black X's.

The grid tailored conserved variable ROM computes accurate axial forces in much smaller regions of parameter space than the primitive variable ROM. Fig. 23f shows that low values of ε_{F_x} are obtained in a small region with high densities and Mach numbers, but that the grid tailored, conserved variable ROM does significantly worse than the initial guess everywhere else. The non-grid tailored ROMs are both worse than the initial guess throughout most of parameter space, with some exceptions. There is a small, diagonal region in which ε_{F_x} is very low for the non-tailored, primitive variable ROM, but the ROM is much worse than the initial guess elsewhere. Overall, the most accurate and robust ROM is the grid tailored LSPG with a primitive variable basis.

As observed for axial force, the initial guess and ROM only compute accurate values of integrated heat flux over small areas of parameter space. This can be seen in Fig. 24. As observed for axial force, the grid tailored ROM with a primitive variable basis has the largest regions of low errors $\varepsilon_{Q_{wall}}$ and improves upon the corresponding initial guess for cases with high densities. The grid tailored ROM with a conserved variable basis improves upon its initial guess in some regions, but is notable worse in others. The same observation can be made for both non-tailored ROMs. Judging by $\varepsilon_{Q_{wall}}$, the most accurate and robust ROM is the grid tailored LSPG with a primitive variable basis.

Finally, we consider some visualizations for the HIFiRE-1 at $Ma_\infty = 9.5$, $\rho_\infty = 0.072 kg/m^3$ to highlight the advantages of grid tailoring for ROMs. This case is outside of the training set, so it shows the predictive capability of the ROM. Fig. 25 shows that while the grid tailored ROM captures the shock accurately, the non-tailored case is unable to, as evidenced by the region of high Mach number upstream of the shock.

The region near the shock is not the only place in which the non-tailored ROM struggles. Fig. 26 shows that while the grid tailored ROM mostly improves upon the initial guess when computing heat flux, the non-tailored ROM is worse. Also it is interesting to note that the initial guess gives similar results with and without tailoring. This likely due to the similar boundary layer resolution with and without tailoring. Overall, these results show that grid tailoring is crucial for the robustness of the ROM when strong shock waves are present in the flow field.

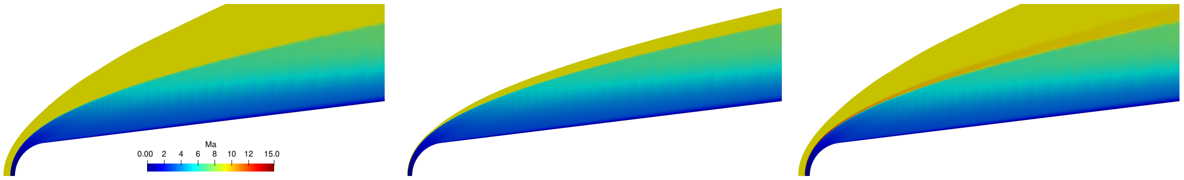

(a) FOM **(b) Grid Tailored FOM** **(c) IG**

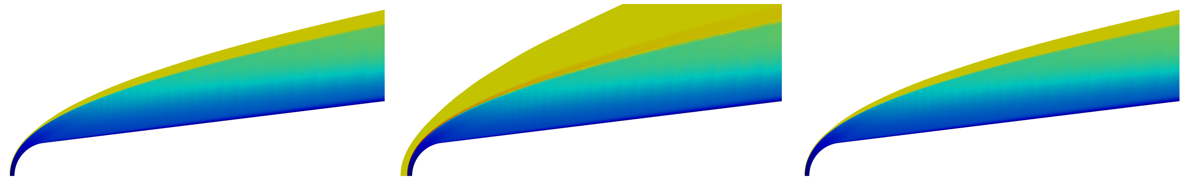

(d) Grid Tailored IG **(e) ROM** **(f) Grid Tailored ROM**

Fig. 25 Mach contours for the 2D HIFiRE-1 at $Ma_\infty = 9.5$, $\rho_\infty = 0.072kg/m^3$. ROM results are for the $p = 3$ LSPG ROM with a primitive variable basis and RBF initial guess. IG is short for "initial guess".

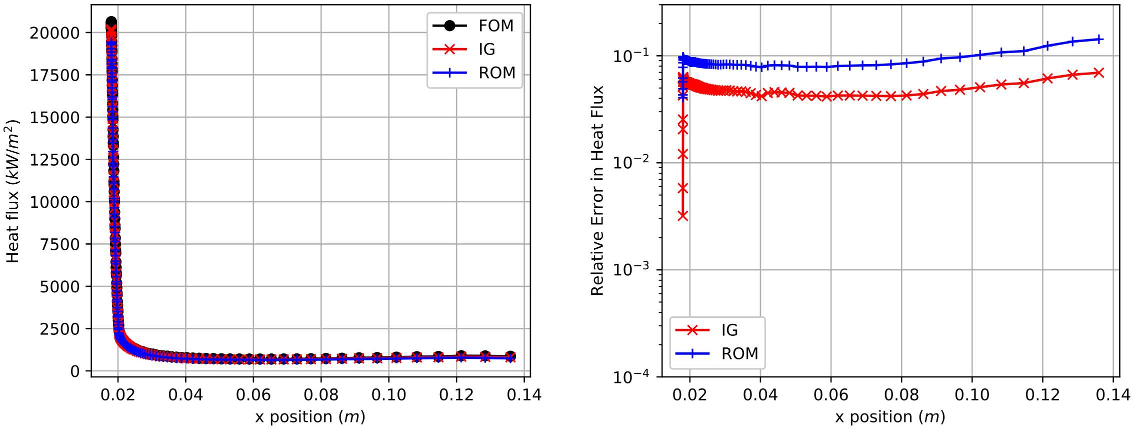

(a) Not tailored **(b) Not tailored**

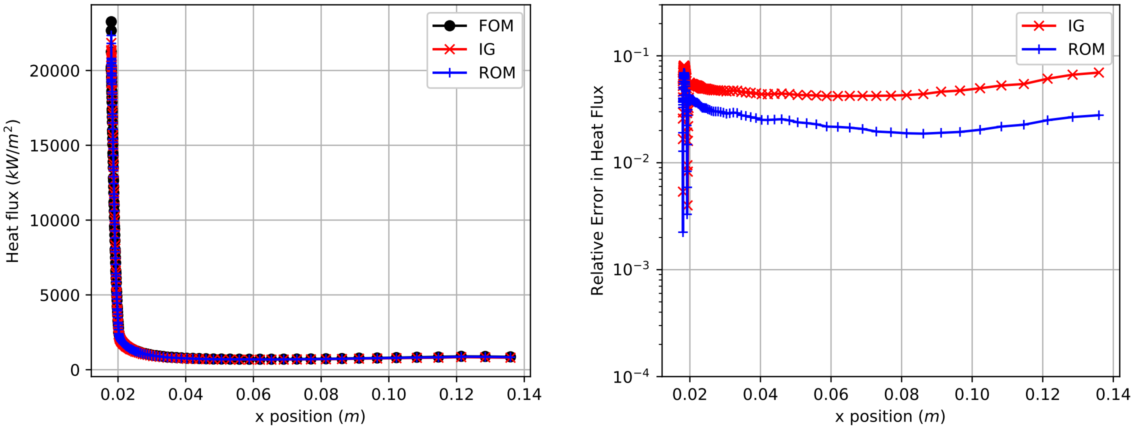

(c) Tailored **(d) Tailored**

Fig. 26 Heat flux and corresponding errors for the 2D HIFiRE-1 at $Ma_\infty = 9.5$, $\rho_\infty = 0.072kg/m^3$ and $p = 3$ LSPG ROM with a primitive variable basis and RBF initial guess. IG is short for "initial guess".

Table 16 Average simulation times and ROM speedup

Configuration	Simulation	Number of cores	Average time	Speedup
2D HiFIRE-1	FOM	32	104.5 s	-
	ROM	32	3.2 s	33.0

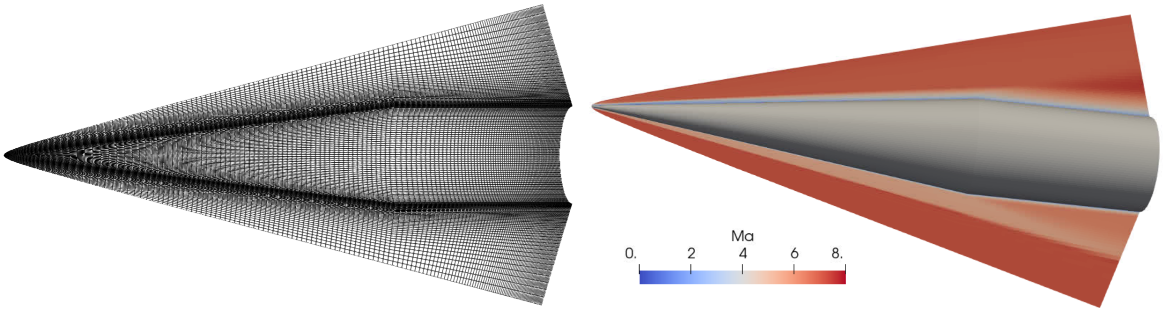

(a) 3D HiFIRE-1 mesh **(b) Mach Contours**

Fig. 27 Mach number from 3D HiFIRE-1 FOM at $Alt = 25km$, $Ma = 7$, $AoA = 6°$.

2. *Reduced-order model performance*

The average simulation times and ROM speedup are given in Table 16. The ROM speedup is defined as the ratio of the FOM simulation time to the ROM simulation time. The ROMs in Table 16 have the setups identified to produce the best results, specifically a primitive variable basis with 3 modes, a LSPG solver, and RBF interpolation for the initial guess. This speed up is only slightly larger than that obtained for the blunt wedge in section IV.C.

3. *Summary*

The following is a summary of lessons learned from the 2D HIFiRE-1 case:

- The more refined grid near the wall and shock leads to similar state errors $\varepsilon_{\boldsymbol{x}}$ for grid tailored and non-tailored interpolations.
- Overall, grid tailored ROMs are more accurate than non-tailored ROMs.
- Low state errors $\varepsilon_{\boldsymbol{x}}$ do not necessarily correspond to low QoI errors ε_{F_x} and $\varepsilon_{Q_{wall}}$.
- Grid tailored ROMs with a primitive variable basis are more accurate than those with a conserved variable basis in state and QoI errors. However, there are large regions in which the ROM and initial guess had errors exceeding 5% for both QoIs.

While a few of these lessons are consistent with those from the 2D wedge case, there a few differences, most notably the clear superiority of primitive variable basis for the state error and both QoI errors for the 2D HIFiRE-1.

E. 3D HIFiRE-1 Forebody

The third case considered in this paper is a three-dimensional simulation of the HiFIRE-1 vehicle nose and body. A three-dimensional computational domain allows us to vary the angle of attack (AoA), so for this case we vary three parameters: altitude, Mach number, and AoA. Fig. 27a shows a view of the mesh and Fig. 27b shows the Mach number from one of the FOMs. The mesh has 1,572,864 cells and the maximum y+ in non-tailored simulations is 1.03. All simulations are run until the relative residual is reduced by 9 orders of magnitude. Simulations are run using the S-A turbulence model assuming a fully turbulent boundary layer.

Grid tailoring is applied twice for the three-dimensional HiFIRE-1, and Table 18 shows the grid tailoring parameters. The relative difference varies between 8.8×10^{-5} to 8.3×10^{-4} for axial force, and 9.8×10^{-5} to 0.0013 for integrated heat flux. Grid tailoring makes little difference to the QoIs for this case because the non-tailored grid is already well aligned with the shock and the grid is refined near the wall. Similar to the 2D HiFIRE-1 case, the maximum surface heat flux is much more affected by grid tailoring, with a relative difference as high as 0.031.

The 3D HiFIRE-1 case is run at varying altitude, Mach number, and AoA with ranges given in Table 17. To generate

Table 17 Free stream flow conditions for the three-dimensional HiFIRE-1.

Altitude	25 to 28 km
Mach Number	7.0 to 9.0
Angle of attack	0.0° to 6.0°
Temperature	216.66 K
Reynolds Number	$1.5 \times 10^7\ 1/m$

Table 18 Grid tailoring parameters for the three-dimensional HiFIRE-1.

Parameter	1st tailoring	2nd tailoring
Number of buffer cells N_b	5	5
Maximum initial spacing $w_{0,max}$	10^{-6} m	10^{-6} m
Shock criterion S_c	0.001	0.001
Number of smoothing iterations N_{smooth}	10	5

the training and test sets, 27 FOMs were run on a 3× 3×3 grid in parameter space, and an additional 12 FOMs were run with random parameter values inside the specified ranges, for a total of 39 FOMs. The training set consists of the eight FOMs at the corners of the parameter space and the other 31 FOMs form the test set.

1. Reduced-order model accuracy

We test a variety of ROM setups with various basis variables, interpolation methods, dictionary sizes, and grid-tailoring. Table 19 shows results from the initial guess surrogates, and Tables 20 and 21 show error statistics from the ROM solutions. The most accurate initial guess surrogate for the two QoIs is non-tailored with primitive variables and RBF interpolation. The initial guesses with grid tailoring, primitive variables, RBF interpolation, and 4 or 7 modes are nearly as accurate for ε_{F_x} and have the same percentage of cases with $\varepsilon_{Q_{wall}} < 5\%$ and $\varepsilon_{Q_{wall}} < 2\%$. Note that interpolation surrogates with and without grid-tailoring are expected to have almost identical errors for a well-resolved grid. The initial guess interpolants are identical with and without grid-tailoring, so if the mesh resolution is such that QoIs are almost identical between the grid-tailored and non-tailored simulations, the initial guess surrogates will also be nearly identical.

Consistent with our previous findings, RBF interpolation is the most accurate, followed by kriging, then inverse distance interpolation. We also observe that the initial guesses with conserved variables have high errors.

The ROM results in Tables 20 and 21 show that the ROM in the second row with grid-tailoring, primitive variables, LSPG, $p = 7$, and RBF interpolation is most accurate at predicting both the axial force and integrated heat flux. The axial force error ε_{F_x} is improved over the initial guess in 80% of cases, and the heat flux error $\varepsilon_{Q_{wall}}$ is improved

Table 19 Errors from 3D HiFIRE-1 initial guesses.

Grid	Variables	p	Interpolation	ε_{F_x}		$\varepsilon_{Q_{wall}}$		Median
				<5%	<1%	<5%	<2%	state error
Tailored	Primitive	4	RBF	90%	13%	61%	19%	0.0135
Tailored	Primitive	7	RBF	87%	13%	61%	19%	0.0132
Tailored	Primitive	7	Kriging	77%	16%	48%	32%	0.0184
Tailored	Primitive	7	Inverse dist.	13%	6%	3%	0%	0.0754
Tailored	Conserved	7	RBF	35%	0%	0%	0%	0.0423
Non-tailored	Primitive	7	RBF	97%	13%	61%	19%	0.0169
Non-tailored	Conserved	7	RBF	35%	0%	0%	0%	0.0600

Table 20 QoI errors from 3D HiFIRE-1. The percentage of ROMs in which the QoI is improved relative to the initial guess is abbreviated as %I. Note that the "% Converged" is omitted because all ROMs considered for this case converged.

Grid	Variables	ROM	p	Interp-olation	ε_{F_x}			$\varepsilon_{Q_{wall}}$		
					<5%	<1%	%I	<5%	<2%	%I
Tailored	Primitive	LSPG	4	RBF	100%	38%	80%	74%	54%	67%
Tailored	Primitive	LSPG	7	RBF	100%	38%	80%	74%	58%	77%
Tailored	Primitive	LSPG	7	Kriging	100%	38%	87%	77%	45%	80%
Tailored	Primitive	LSPG	7	Inv. dist.	100%	35%	90%	58%	23%	97%
Tailored	Conserved	LSPG	7	RBF	48%	3%	80%	19%	3%	67%
Non-tailored	Primitive	LSPG	7	RBF	39%	10%	26%	42%	16%	32%
Non-tailored	Conserved	LSPG	7	RBF	22%	3%	22%	19%	6%	32%

Table 21 State errors from 3D HiFIRE-1.

Grid	Variables	ROM	p	Interp-olation	State Error	
					Model	Proj.
Tailored	Primitive	LSPG	4	RBF	0.0213	0.0131
Tailored	Primitive	LSPG	7	RBF	0.0194	0.0129
Tailored	Primitive	LSPG	7	Kriging	0.0195	0.0129
Tailored	Primitive	LSPG	7	Inv. dist.	0.0327	0.0129
Tailored	Conserved	LSPG	7	RBF	0.0333	0.0250
Non-tailored	Primitive	LSPG	7	RBF	0.0694	0.0167
Non-tailored	Conserved	LSPG	7	RBF	0.0913	0.0412

Table 22 Average simulation times and ROM speedup

Configuration	Simulation	Number of cores	Average time	Speedup
3D HiFIRE-1	FOM	64	41.6 min	-
	ROM	64	0.41 min	102

in 77% of cases. The ROM with $p = 4$ has similar QoI errors as the comparable case with $p = 7$, but it has higher state errors $\varepsilon_{\boldsymbol{x}}$. Note that the ROM with Kriging interpolation initial guesses has similar accuracy for axial force and slightly more cases with hear flux errors below 5%; however it has fewer cases with errors below 2%. The ROMs with conserved variables are much less accurate, as are the non-tailored ROMs. Out of the non-tailored ROMs, the LSPG ROM with a primitive variable basis is most accurate for ε_{F_x}, and the LSPG ROM with a primitive variable basis has lowest errors of $\varepsilon_{Q_{wall}}$. However, in all the non-tailored setups, the percentage of ROMs that improve each QoI error is below 50%, implying that the ROMs are ineffective and a well-designed surrogate model is superior to non-tailored ROMs in the majority of cases.

Fig. 28 shows the initial guess and ROM errors for the non-tailored case with LSPG, a $p = 7$ primitive variable basis, and RBF interpolation. All three error metrics have less accurate ROM solutions than the initial guess for a large majority all of the cases. There are four cases which have low state errors in the initial guess; they are the four edge cases for which there are training cases at the same Mach number and AoA. Since the shock location does not move when the freestream density changes, those four cases have shocks at the same location as some of the training snapshots, so they are expected to have low state errors $\varepsilon_{\boldsymbol{x}}$ with non-tailored simulations. Those four cases also have low QoI errors in the ROM solutions.

Fig. 29 shows results for the most accurate grid-tailored case, which uses LSPG, a $p = 7$ primitive variable basis, and RBF interpolation. Both axial force and heat flux are much better predicted with the ROM than the initial guess. The ROM state error $\varepsilon_{\boldsymbol{x}}$ remains almost unchanged or increases slightly from the initial guess even though the QoI errors decreased, showing again that the state error alone is not a good indicator of ROM accuracy.

Finally, we consider flow visualizations of the ROM results at $Alt = 26.5km$, $Ma = 8$, and $AoA = 3°$. Fig. 30 shows Mach contours and heat flux contours for the FOM, initial guess, and ROM with and without grid tailoring. Interestingly, the grid tailored ROM has a smoother looking mesh than the corresponding FOM. This is because the limited grid displacement data that is used to interpolate the mesh for the ROM does not contain the oscillations seen in the FOM for this flight condition. The most obvious difference between the flow fields can be seen in the bow shock structure. In Figs. 30c and 30e, the shock on top of the vehicle is poorly approximated as multiple discontinuities, in contrast to the grid tailored initial guess and ROM, which capture the upper shock well.

Inspecting the flow near the nose reveals more issues with the shockwave. Figs. 31b and 31c show that interpolation and the ROM are unable to accurately capture the bow shock without grid-tailoring. The bow shock is approximated as two discontinuities, most noticeably downstream of the rounded tip. This "staircase" approximation of the shock is similar to results shown in [25] for ROMs without any mesh deformation. In contrast, Figs. 31e and 31f both approximate the shock well, with a single discontinuity in the correct position.

Since grid tailoring positions the shock accurately, the residual minimization at the heart of LSPG is more focused on residuals in other regions of the flow, such as near the wall. Fig. 32 shows that the ROM reduced errors in heat flux on most of the vehicle surface, with one exception just downstream of the nose tip. By lowering the error on the bottom of the vehicle, where the heat flux is highest (see Fig. 32a), the ROM is able to improve the accuracy of integrated heat flux substantially. Although the error is reduced on top, it is still relatively high (>10%) along much of the top centerline, especially just aft of the nose tip. This corresponds to a region with low heat flux, as shown in Fig. 32a, likely due to an adverse pressure gradient brought about by the nonzero angle of attack. Despite this minor shortcoming, these results show the advantages of incorporating residual minimization in model reduction; if implemented correctly, with the right basis, residual minimization uses the governing equations to choose the best possible solution using the limited data available to it.

2. Reduced-order model performance

The average simulation times and ROM speedup are given in Table 22. The ROM speedup is defined as the ratio of the FOM simulation time to the ROM simulation time. The ROMs in Table 22 have the setups identified to produce the best results, specifically a primitive variable basis with 7 modes, a LSPG solver, and RBF interpolation for the initial

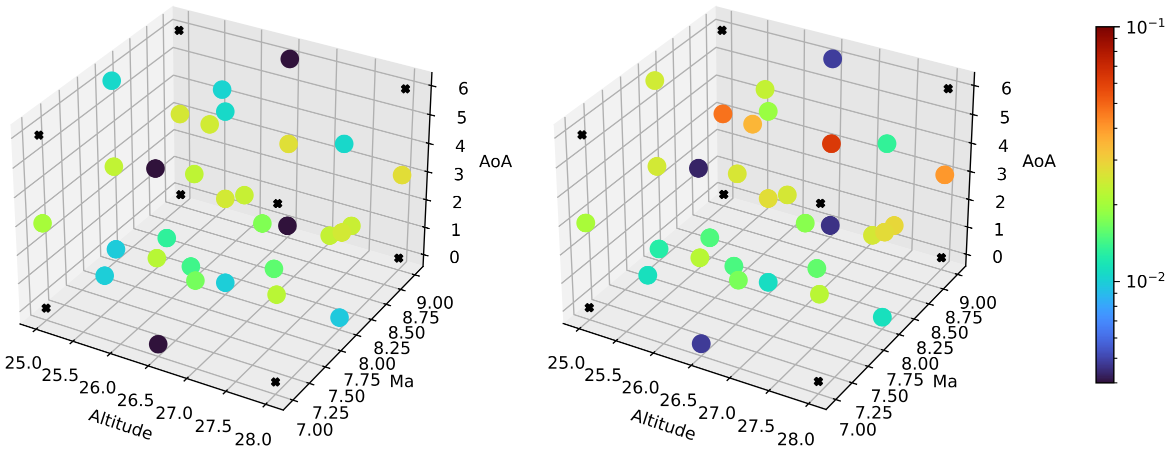

(a) ε_{x} with RBF Interpolation

(b) ε_{x} with Primitive variable LSPG

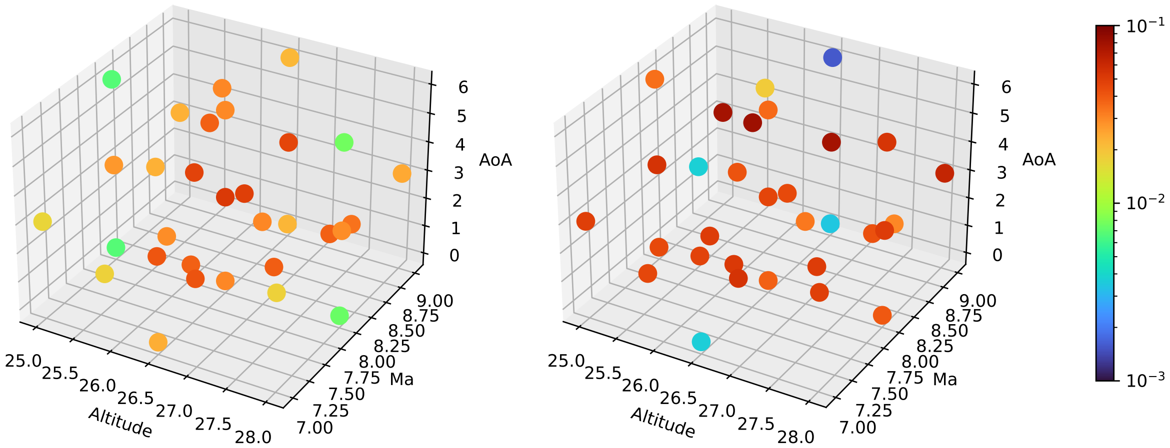

(c) ε_{F_x} with RBF Interpolation

(d) ε_{F_x} with Primitive variable LSPG

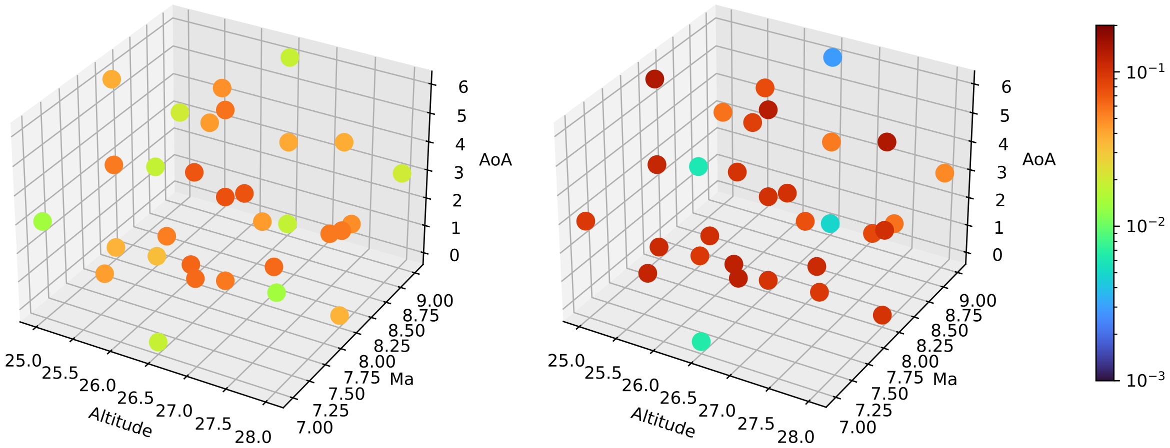

(e) $\varepsilon_{Q_{wall}}$ with RBF Interpolation

(f) $\varepsilon_{Q_{wall}}$ with Primitive variable LSPG

Fig. 28 Errors ε_{x}, ε_{F_x}, and $\varepsilon_{Q_{wall}}$ for the 3D HiFIRE-1 projectile using non-tailored simulations with primitive variables, LSPG, $p = 7$, and RBF interpolation. The figures on the left show the initial guess errors and the figures on the right show ROM solution errors. Training cases are indicated by black X's.

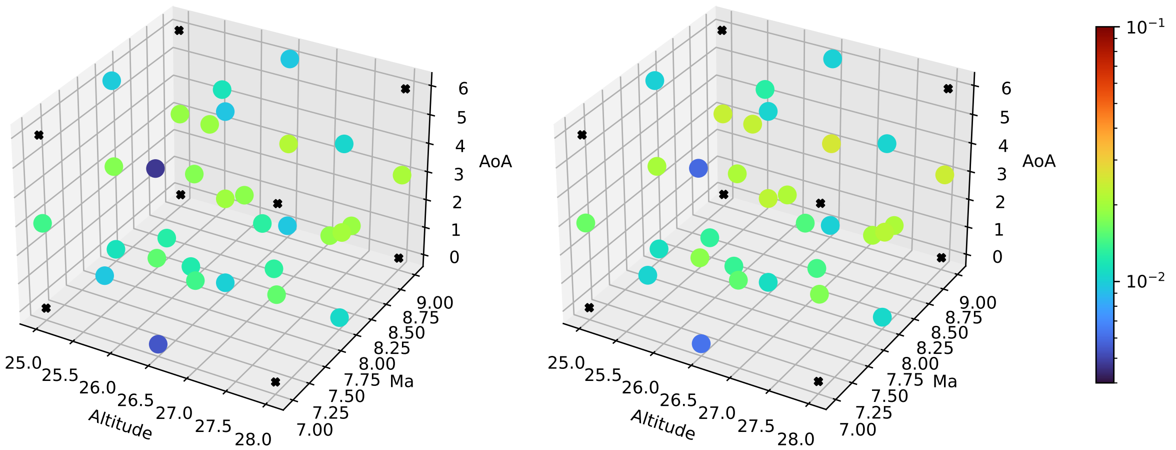

(a) ε_x with RBF Interpolation

(b) ε_x with Primitive variable LSPG

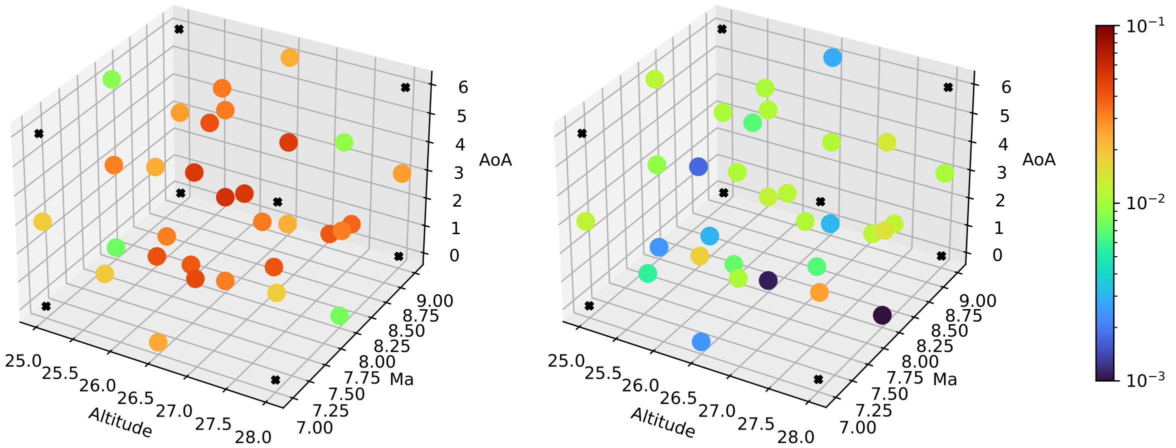

(c) ε_{F_x} with RBF Interpolation

(d) ε_{F_x} with Primitive variable LSPG

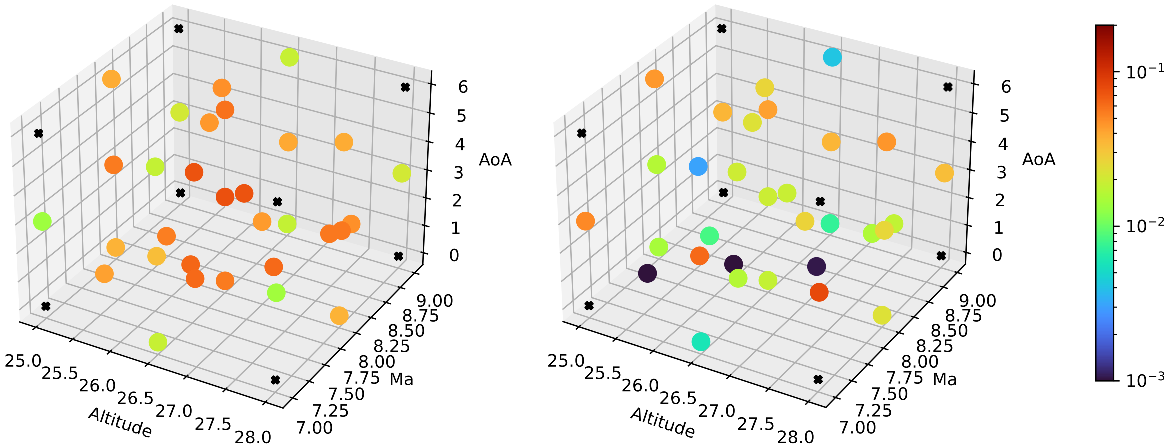

(e) $\varepsilon_{Q_{wall}}$ with RBF Interpolation

(f) $\varepsilon_{Q_{wall}}$ with Primitive variable LSPG

Fig. 29 Errors ε_x, ε_{F_x}, and $\varepsilon_{Q_{wall}}$ for the 3D HiFIRE-1 projectile using grid-tailored simulations with primitive variables, LSPG, $p = 7$, and RBF interpolation. The figures on the left show the initial guess errors and the figures on the right show ROM solution errors. Training cases are indicated by black X's.

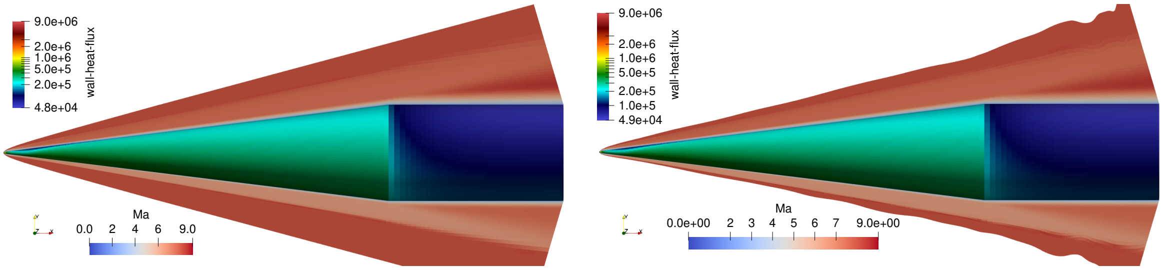

(a) FOM **(b) Grid Tailored FOM**

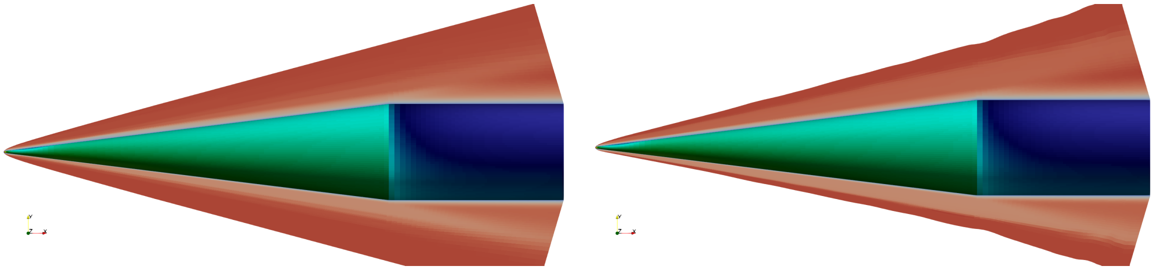

(c) IG **(d) Grid Tailored IG**

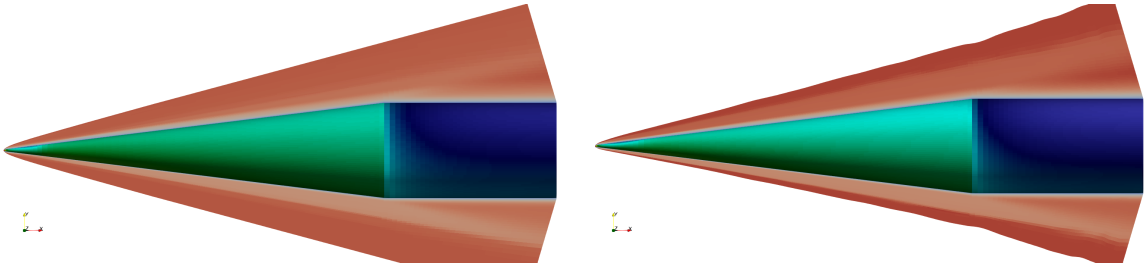

(e) ROM **(f) Grid Tailored ROM**

Fig. 30 Mach number and wall heat flux for the 3D HiFIRE-1 at $Alt = 26.5km$, $Ma = 8$, $AoA = 3°$. ROM results are for C-LSPG with a $p = 7$ primitive variable basis. The initial guess (IG) is with RBF interpolation.

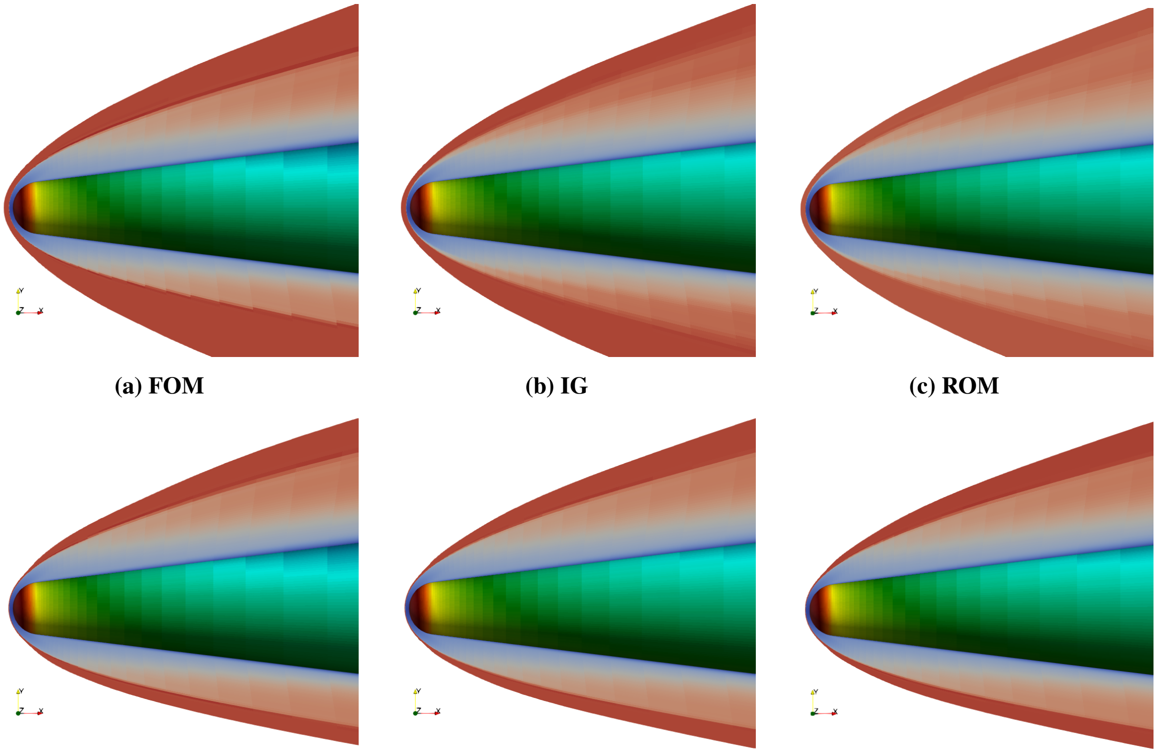
(a) FOM (b) IG (c) ROM

(d) Grid Tailored FOM (e) Grid Tailored IG (f) Grid Tailored ROM

Fig. 31 Mach number and wall heat flux for the nose region of 3D HiFIRE-1 at $Alt = 26.5km$, $Ma = 8$, $AoA = 3°$. ROM results are for LSPG with a $p = 7$ primitive variable basis. The initial guess (IG) is with RBF interpolation.

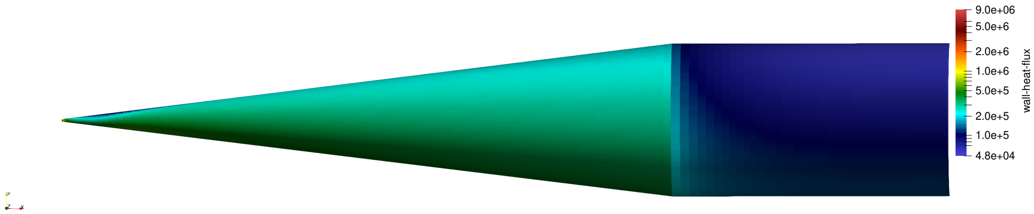

(a) FOM

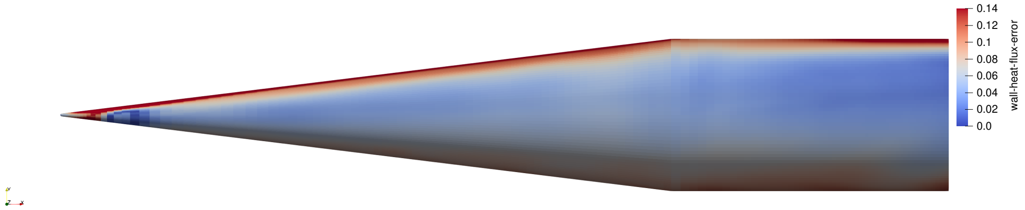

(b) IG

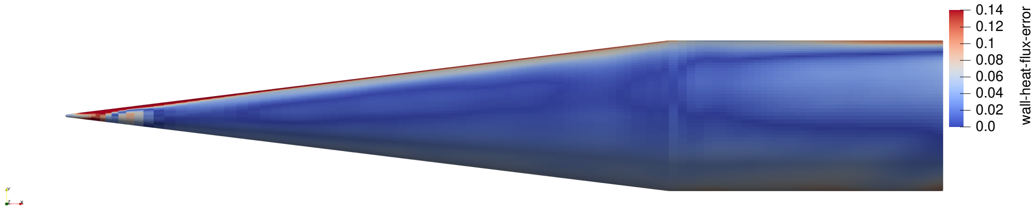

(c) ROM

Fig. 32 Wall heat flux errors for the grid tailored 3D HiFIRE-1 at $Alt = 26.5km$, $Ma = 8$, $AoA = 3°$. ROM results are for LSPG with a $p = 7$ primitive variable basis. The initial guess (IG) is with RBF interpolation.

guess. Note that this case has a larger speed-up than the two 2D cases, likely due to it having a larger computational mesh.

3. Summary

Some conclusions drawn from the 3D HiFIRE-1 study are:

- Grid tailored ROMs are much more effective than non-tailored ROMs
- Primitive variables lead to more accurate initial guesses and ROMs
- RBF interpolation is more accurate, followed by closely by Kriging, then inverse distance interpolation

F. Discussion

In all three cases, grid tailored ROMs are superior to non-tailored ROMs. The non-tailored ROMs are ineffective in all three cases, while grid-tailored ROMs improve QoI errors over the best initial guess surrogates. In two of the cases, the grid-tailored ROMs show dramatic improvements to both axial force and integrated heat flux. In the 2D HIFiRE-1 case, only slight improvements are observed for the most accurate ROM setup. This case is particularly challenging for the ROMs because the grid is very refined, so the shock occurs over a small distance. Even slight perturbations in the tailored grid can change which cells contain the shock, in which case the shock is not at the same location in all the training snapshots. This issue can be mitigated with improvements in the grid tailoring algorithm to ensure that the number of cells between the bow shock and inflow boundary is constant. Improvements in grid tailoring will also enable the application of ROMs to more complicated geometries such as vehicles with fins, which have secondary shock structures that are not aligned with grid.

The improvements in accuracy and robustness from grid tailoring occurs because grid tailoring allows the basis to approximate the shock accurately, lowering the residuals near the shock wave relatively to those in a non-tailored ROM. The lower residuals near the shock allow the LSPG residual minimization to focus more on reducing residuals in other regions of the flow, such as near the wall, increasing the accuracy of wall QoIs. This shows the importance of an accurate basis in the presence of convective features including shocks, and has ramifications for fluid dynamic applications beyond hypersonic vehicle aerodynamics.

Another insight provided about the basis is that a local dictionary is as good or better than POD modes for approximating the solution state with limited training data. This ensures that the basis is not polluted by non-local flow features and eliminates the need to compute singular value decompositions offline. Additionally, the sparseness of the training data considered in this study likely makes the local dictionary more robust than observed in Refs. [14, 15]. Sampling at fewer points in parameter space makes it more likely that each solution snapshot is quite different than the others, making it highly unlikely that the solutions are linearly dependent on one another.

These insights were made possible by considering multiple error metrics. All three cases show that ℓ^2 state error alone is not a good judge of overall ROM accuracy. It is observed that relatively low state errors do not necessarily correspond to low QoI errors, as in Ref. [17]. This shows the importance of comparing the different ROM setups by inspecting multiple QoIs in addition to state errors, as is done in this paper.

Our previous work that applied projection-ROMs to hypersonic aerodynamics showed that ROMs can improve QoI predictions over the initial guess interpolation [17]. However, in that work inverse distance interpolation is used for the initial guess surrogates, and in all three cases here inverse distance interpolation was the worst interpolation method, followed by kriging. Using inverse distance or kriging interpolation can give misleading results; if the interpolation method has poor accuracy, the ROM has more potential to improve over the initial guess. The overall goal of using ROMs is to provide a simulation method that is more accurate than a well-designed surrogate model, so ROMs should use and be compared with accurate surrogate models. We suggest that future works use RBF interpolation as a baseline interpolation method or examine multiple interpolation methods to compute the initial guess. In addition to providing a good initial guess for the ROM, RBF interpolation might be of acceptable accuracy as a surrogate model itself. RBF interpolation is more accurate than the ROM for non-tailored FOM data, and is competitive with the grid tailored ROM for some QoIs and certain areas of parameter space in the three cases considered.

Studying individual cases sometimes led to different conclusions. In the blunt wedge case, it was found that conserved variables led to the most accurate axial force predictions while primitive variables led to the most accurate integrated heat flux. In the other two cases, primitive variable ROMs were superior to conserved variable ROMs. It is possible that there is no single optimal ROM setup that can be used for any hypersonic flow configuration and for any QoI, but that the optimal ROM setup for a given configuration depends on the boundary conditions, grid, equations, parameters, and QoIs.

Table 23 Average simulation times and ROM speedup

Configuration	Simulation	Number of cores	Average time	Speedup
Blunt Wedge	FOM	16	101.0 s	-
	ROM	16	3.7 s	27.5
2D HiFIRE-1	FOM	32	104.5 s	-
	ROM	32	3.2 s	33.0
3D HiFIRE-1	FOM	64	41.6 min	-
	ROM	64	0.41 min	102

The case with the best ROM performance and accuracy was the 3D case, which considered variations over three parameters. The other two cases only considered variations over two parameters. It is possible that ROMs have greater benefits when the parameter space has higher dimension. Further studies are need to determine whether this is because of the more complex flow fields in 3D, the additional parameter, or a combination of these factors.

The average simulation times and ROM speedup for all three cases are given in Table 23. The ROM speedup is defined as the ratio of the FOM simulation time to the ROM simulation time. The ROMs in Table 23 have the setups identified to produce the best results. Specifically, the blunt wedge ROM uses an 8 mode conserved variable basis, LSPG, and RBF interpolation. The 2D HiFIRE-1 uses a 3 mode primitive variable basis, LSPG, and RBF interpolation. The 3D HiFIRE-1 uses a 7 mode primitive variable basis, LSPG, and RBF interpolation.

All the ROMs in Table 23 show significant speedup over the FOMs even though none of the ROMs use hyperreduction, with the blunt wedge having the minimum speedup of 27.5. The speedup is lowest for the 2D cases and increases with the computational size of a configuration. The reasons a speed-up is obtained despite the lack of hyper-reduction is mostly due to the use of a direct nonlinear solver in the ROM, which required far few iterations than the pseudo-time-stepping approach like the one used for the FOM, as discussed in section III.B. One of the reasons for the ROM needing fewer nonlinear iterations comes from the smaller state space of the ROM. For example, the 3D HIFiRE-1 has 1,572,864 cells and solves 6 equations, so it has a state space of dimension 9,437,184 in which to find a solution that minimizes the residual. In contrast, the ROM has 7 modes, so it only needs to find a solution in a state space of dimension 7. Optimizers are more efficient at finding solutions in lower dimensions, so the ROMs require far fewer evaluations of the residual and corresponding Jacobian than FOMs, thereby resulting in considerable speedup despite not using hyperreduction.

V. Conclusion

This paper presents ROM approaches for steady hypersonic aerodynamic simulations. Specifically, we investigate LSPG ROMs, both with and without a conservation constraint. Building on earlier work on LSPG for steady hypersonics, we explore several modifications including the use of local bases, grid tailoring, primitive variable bases, and more advanced interpolation techniques for the initial guess.

Firstly, we find that grid tailoring and local bases are key components of accurate and robust ROMs. The local basis ensures that solution snapshots with irrelevant flow features are not included in a given ROM basis. Using grid tailoring provides a mechanism to accurately capture shock waves that move as input parameters are varied. Grid tailoring is observed to increase ROM accuracy for all three test cases considered, especially the three dimensional, three parameter case.

While some modifications almost always improve the accuracy of LSPG, the utility of other modifications are problem dependent. Using Grid tailoring and RBF interpolation initial guesses consistently improves ROM accuracy, while the best choice of basis variables seems to be problem dependent. Primitive variables are better overall for two out of the three test cases. In the one outlier case, the blunt wedge, the primitive basis improves accuracy for integrated heat flux, but degrades it for axial force. These results suggest that there may not be a single optimal ROM setup for any hypersonic flow configuration and for any QoI. Rather, the optimal ROM setup for a given configuration may be problem dependent or even vary over parameter space.

In evaluating these results, we find that ℓ^2 state error alone is not a good judge of overall ROM accuracy; in all three test cases low state error do not necessarily correspond to low QoI errors. This shows the importance of evaluating ROM accuracy by inspecting multiple QoIs in additional to state errors. This finding is consistent with previous findings on ROMs for steady hypersonic aerodynamics.

There are many directions to choose from to follow-on the work presented in this paper. One of the crucial directions is to improve grid tailoring, as improved grid tailoring will improve the accuracy and robustness of the ROM as well as the FOM itself. One key improvement of grid tailoring for the ROM would be to reduce the discrepancy between the actual number of buffer cells and the targeted number; the ROM will be most accurate when the target is always hit. This could be achieved by improving the shock detection and smoothing algorithms. Improvements to the smoothing algorithms will also reduce the oscillations of the grid observed at the inflow boundary for some cases. Another key improvement would be the generalization of grid tailoring to cases with secondary shock structures, such as those observed on vehicles with fins. This would enable deployment of the ROM to more complicated flight vehicle geometries. This work showed significant speedup without hyperreduction, but future works should implement hyperreduction for grid tailored meshes. Additionally, more work is needed to examine which ROM setups are best across a wide suite of hypersonic configurations. Future works could expand parameter ranges, examine other interpolation methods, and introduce ROMs in higher dimension parameter spaces.

Acknowledgments

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

We are also grateful to Derek Dinzl and the SPARC team for their guidance on how to implement the Pressio interface and how to use grid tailoring.

References

[1] Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., and Mavriplis, D., "CFD vision 2030 study: a path to revolutionary computational aerosciences," Tech. Rep. NASA/CR–2014-218178, NASA, March 2014.

[2] Yondo, R., Andrés, E., and Valero, E., "A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses," *Progress in Aerospace Sciences*, Vol. 96, 2018, pp. 23–61.

[3] Lieu, T., Farhat, C., and Lesoinne, M., "POD-based aeroelastic analysis of a complete F-16 configuration: ROM adaptation and demonstration," *46th AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics and materials conference*, 2005, p. 2295.

[4] Amsallem, D., and Farhat, C., "Interpolation method for adapting reduced-order models and application to aeroelasticity," *AIAA journal*, Vol. 46, No. 7, 2008, pp. 1803–1813.

[5] Carlberg, K., Bou-Mosleh, C., and Farhat, C., "Efficient non-linear model reduction via a least-squares Petrov–Galerkin projection and compressive tensor approximations," *International Journal for Numerical Methods in Engineering*, Vol. 86, No. 2, 2011, pp. 155–181. https://doi.org/10.1002/nme.3050, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.3050.

[6] Carlberg, K., Farhat, C., Cortial, J., and Amsallem, D., "The GNAT method for nonlinear model reduction: Effective implementation and application to computational fluid dynamics and turbulent flows," *Journal of Computational Physics*, Vol. 242, 2013, pp. 623 – 647. https://doi.org/10.1016/j.jcp.2013.02.028, URL http://www.sciencedirect.com/science/article/pii/S0021999113001472.

[7] Bui-Thanh, T., Willcox, K., and Ghattas, O., "Parametric reduced-order models for probabilistic analysis of unsteady aerodynamic applications," *AIAA journal*, Vol. 46, No. 10, 2008, pp. 2520–2529.

[8] LeGresley, P., and Alonso, J., "Investigation of non-linear projection for pod based reduced order models for aerodynamics," *39th Aerospace Sciences Meeting and Exhibit*, 2001, p. 926.

[9] Rovas, D. V., "Reduced-basis output bound methods for parametrized partial differential equations," Ph.D. thesis, Massachusetts Institute of Technology, 2003.

[10] Rowley, C. W., Colonius, T., and Murray, R. M., "Model reduction for compressible flows using POD and Galerkin projection," *Physica D: Nonlinear Phenomena*, Vol. 189, No. 1, 2004, pp. 115 – 129. https://doi.org/https://doi.org/10.1016/j.physd.2003.03.001, URL http://www.sciencedirect.com/science/article/pii/S0167278903003841.

[11] Terragni, F., Valero, E., and Vega, J. M., "Local POD plus Galerkin projection in the unsteady lid-driven cavity problem," *SIAM Journal on Scientific Computing*, Vol. 33, No. 6, 2011, pp. 3538–3561.

[12] Barone, M. F., Kalashnikova, I., Segalman, D. J., and Thornquist, H. K., "Stable Galerkin reduced order models for linearized compressible flow," *Journal of Computational Physics*, Vol. 228, No. 6, 2009, pp. 1932–1946.

[13] Carlberg, K., Barone, M., and Antil, H., "Galerkin v. least-squares Petrov–Galerkin projection in nonlinear model reduction," *Journal of Computational Physics*, Vol. 330, 2017, pp. 693 – 734. https://doi.org/10.1016/j.jcp.2016.10.033, URL http://www.sciencedirect.com/science/article/pii/S0021999116305319.

[14] Abgrall, R., Amsallem, D., and Crisovan, R., "Robust model reduction by L^1-norm minimization and approximation via dictionaries: application to nonlinear hyperbolic problems," *Advanced Modeling and Simulation in Engineering Sciences*, Vol. 3, No. 1, 2016, p. 1. https://doi.org/10.1186/s40323-015-0055-3, URL https://doi.org/10.1186/s40323-015-0055-3.

[15] Abgrall, R., and Crisovan, R., "Model reduction using L^1-norm minimization as an application to nonlinear hyperbolic problems," *International Journal for Numerical Methods in Fluids*, Vol. 87, No. 12, 2018, pp. 628–651. https://doi.org/10.1002/fld.4507, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/fld.4507.

[16] Carlberg, K., Choi, Y., and Sargsyan, S., "Conservative model reduction for finite-volume models," *Journal of Computational Physics*, Vol. 371, 2018, pp. 280 – 314. https://doi.org/10.1016/j.jcp.2018.05.019, URL http://www.sciencedirect.com/science/article/pii/S002199911830319X.

[17] Blonigan, P. J., Rizzi, F., Howard, M., Fike, J. A., and Carlberg, K. T., "Model Reduction for Steady Hypersonic Aerodynamics via Conservative Manifold Least-Squares Petrov–Galerkin Projection," *AIAA Journal*, Vol. 59, No. 4, 2021, pp. 1296–1312. https://doi.org/10.2514/1.J059785, URL https://doi.org/10.2514/1.J059785.

[18] Falkiewicz, N. J., and Cesnik, C. E. S., "Proper Orthogonal Decomposition for Reduced-Order Thermal Solution in Hypersonic Aerothermoelastic Simulations," *AIAA Journal*, Vol. 49, No. 5, 2011, pp. 994, 1010. https://doi.org/10.2514/1.J050701.

[19] 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, 2369. https://doi.org/10.2514/1.J055499.

[20] Mignolet, M. P., Culler, A. J., McNamara, J. J., Matney, A., and Spottswood, S. M., "Panel response prediction through reduced order models with application to hypersonic aircraft," *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015, p. 1630.

[21] Carlberg, K., "Adaptive h-refinement for reduced-order models," *International Journal for Numerical Methods in Engineering*, Vol. 102, No. 5, 2015, pp. 1192–1210. https://doi.org/https://doi.org/10.1002/nme.4800, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.4800.

[22] Etter, P. A., and Carlberg, K. T., "Online adaptive basis refinement and compression for reduced-order models via vector-space sieving," *Computer Methods in Applied Mechanics and Engineering*, Vol. 364, 2020, p. 112931. https://doi.org/https://doi.org/10.1016/j.cma.2020.112931, URL https://www.sciencedirect.com/science/article/pii/S0045782520301146.

[23] Peherstorfer, B., "Sampling Low-Dimensional Markovian Dynamics for Preasymptotically Recovering Reduced Models from Data with Operator Inference," *SIAM Journal on Scientific Computing*, Vol. 42, No. 5, 2020, pp. A3489–A3515. https://doi.org/10.1137/19M1292448, URL https://doi.org/10.1137/19M1292448.

[24] Lee, K., and Carlberg, K. T., "Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders," *Journal of Computational Physics*, Vol. 404, 2020, p. 108973. https://doi.org/10.1016/j.jcp.2019.108973, URL http://www.sciencedirect.com/science/article/pii/S0021999119306783.

[25] Nair, N. J., and Balajewicz, M., "Transported snapshot model order reduction approach for parametric, steady-state fluid flows containing parameter-dependent shocks," *International Journal for Numerical Methods in Engineering*, Vol. 117, No. 12, 2019, pp. 1234–1262. https://doi.org/10.1002/nme.5998, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.5998.

[26] Mojgani, R., and Balajewicz, M., "Physics-aware registration based auto-encoder for convection dominated PDEs," , 2020. URL https://arxiv.org/pdf/2006.15655.

[27] Taddei, T., "A Registration Method for Model Order Reduction: Data Compression and Geometry Reduction," *SIAM Journal on Scientific Computing*, Vol. 42, No. 2, 2020, pp. A997–A1027. https://doi.org/10.1137/19M1271270, URL https://doi.org/10.1137/19M1271270.

[28] Ferrero, A., Taddei, T., and Zhang, L., "Registration-based model reduction of parameterized two-dimensional conservation laws," , 2021. URL https://arxiv.org/pdf/2105.02024, submitted to Elsevier, May 2021.

[29] Rim, D., Peherstorfer, B., and Mandli, K. T., "Manifold Approximations via Transported Subspaces: Model reduction for transport-dominated problems," , 2020.

[30] Saunders, D., Yoon, S., and Wright, M., "An approach to shock envelope grid tailoring and its effect on reentry vehicle solutions," *45th AIAA Aerospace Sciences Meeting and Exhibit*, 2007, p. 207.

[31] Martin, A., and Boyd, I., "Mesh tailoring for strongly coupled computation of ablative material in nonequilibrium hypersonic flow," *10th AIAA/ASME Joint Thermophysics and Heat Transfer Conference*, 2010, p. 5062.

[32] Neitzel, K., and Boyd, I. D., "Influence of Turbulence Modeling On Aftbody Surface Heating Prediction For A Hypersonic Entry Capsule," *44th AIAA Thermophysics Conference*, 2013, p. 2774.

[33] Johnson, H., Alba, C., Candler, G., MacLean, M., Wadhams, T., and Holden, M., "Boundary-layer stability analysis of the hypersonic international flight research transition experiments," *Journal of Spacecraft and Rockets*, Vol. 45, No. 2, 2008, pp. 228–236.

[34] Alba, C. R., Johnson, H. B., Bartkowicz, M. D., Candler, G. V., and Berger, K. T., "Boundary-layer stability calculations for the HIFiRE-1 transition experiment," *Journal of Spacecraft and Rockets*, Vol. 45, No. 6, 2008, pp. 1125–1133.

[35] Martin, A., and Boyd, I. D., "Strongly coupled computation of material response and nonequilibrium flow for hypersonic ablation," *Journal of Spacecraft and Rockets*, Vol. 52, No. 1, 2015, pp. 89–104.

[36] Gronvall, J. E., Johnson, H. B., and Candler, G. V., "Boundary-layer stability analysis of the high enthalpy shock tunnel transition experiments," *Journal of Spacecraft and Rockets*, Vol. 51, No. 2, 2014, pp. 455–467.

[37] Reddy, D. S. K., Saikia, B., and Sinha, K., "Effect of high-enthalpy air chemistry on stagnation point heat flux," *Journal Of Thermophysics And Heat Transfer*, Vol. 28, No. 2, 2014, pp. 356–359.

[38] Reddy, S., and Sinha, K., "Analysis of High-enthalpy Air-chemistry and its Effect on Stagnation Point Heat Flux," *43rd AIAA Thermophysics Conference*, 2012, p. 3002.

[39] 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," *23rd AIAA Computational Fluid Dynamics Conference*, Vol. 1, AIAA, Denver, CO, 2017. https://doi.org/10.2514/6.2017-4407, URL https://arc.aiaa.org/doi/abs/10.2514/6.2017-4407.

[40] LeVeque, R. J., *Finite volume methods for hyperbolic problems*, Vol. 31, Cambridge university press, 2002. https://doi.org/10.1017/CBO9780511791253.

[41] Eymard, R., Gallouët, T., and Herbin, R., "Finite volume methods," *Handbook of numerical analysis*, Vol. 7, 2000, pp. 713–1018. https://doi.org/10.1016/S1570-8659(00)07005-8.

[42] Candler, G., Barnhardt, M., Drayna, T., Nompelis, I., Peterson, D., and Subbareddy, P., "Unstructured grid approaches for accurate aeroheating simulations," *18th AIAA Computational Fluid Dynamics Conference*, 2007, p. 3959.

[43] Holmes, P., Lumley, J., and Berkooz, G., *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, 1996. https://doi.org/10.1017/CBO9780511622700.

[44] Prud'homme, C., Rovas, D. V., Veroy, K., Machiels, L., Maday, Y., Patera, A. T., and Turinici, G., "Reliable Real-Time Solution of Parametrized Partial Differential Equations: Reduced-Basis Output Bound Methods ," *Journal of Fluids Engineering*, Vol. 124, No. 1, 2001, pp. 70–80. https://doi.org/10.1115/1.1448332, URL https://doi.org/10.1115/1.1448332.

[45] Rozza, G., Huynh, D. B. P., and Patera, A. T., "Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations," *Archives of Computational Methods in Engineering*, Vol. 15, No. 3, 2008, pp. 229–275. https://doi.org/10.1007/BF03024948.

[46] Huang, C., Duraisamy, K., and Merkle, C., "Investigations and Improvement of Robustness of Reduced-Order Models of Reacting Flow," *AIAA Scitech 2019 Forum*, Vol. 1, AIAA, San Diego, CA, 2019. https://doi.org/10.2514/6.2019-2012, URL https://arc.aiaa.org/doi/abs/10.2514/6.2019-2012.

[47] Washabaugh, K. M., Zahr, M. J., and Farhat, C., "On the use of discrete nonlinear reduced-order models for the prediction of steady-state flows past parametrically deformed complex geometries," *54th AIAA Aerospace Sciences Meeting*, 2016, p. 1814.

[48] Rasmussen, C. E., and Williams, C. K. I., *Gaussian processes for machine learning*, The MIT Press, 2006.

[49] Zimmermann, R., and Görtz, S., "Non-linear reduced order models for steady aerodynamics," *Procedia Computer Science*, Vol. 1, No. 1, 2010, pp. 165–174.

[50] Schaback, R., "Error estimates and condition numbers for radial basis function interpolation," *Advances in Computational Mathematics*, Vol. 3, No. 3, 1995, pp. 251–264.

[51] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Édouard Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, Vol. 12, No. 85, 2011, pp. 2825–2830. URL http://jmlr.org/papers/v12/pedregosa11a.html.

[52] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, Vol. 17, 2020, pp. 261–272. https://doi.org/10.1038/s41592-019-0686-2.

[53] Mühlenstädt, T., and Kuhnt, S., "Kernel interpolation," *Computational statistics & data analysis*, Vol. 55, No. 11, 2011, pp. 2962–2974.

[54] Wu, C.-Y., Mossa, J., Mao, L., and Almulla, M., "Comparison of different spatial interpolation methods for historical hydrographic data of the lowermost Mississippi River," *Annals of GIS*, Vol. 25, No. 2, 2019, pp. 133–151.

[55] Amidror, I., "Scattered data interpolation methods for electronic imaging systems: a survey," *Journal of electronic imaging*, Vol. 11, No. 2, 2002, pp. 157–176.

[56] Dolvin, D., "Hypersonic international flight research and experimentation (HIFiRE) fundamental science and technology development strategy," *15th AIAA International Space Planes and Hypersonic Systems and Technologies Conference*, 2008, p. 2581.

[57] Wadhams, T., Mundy, E., MacLean, M., and Holden, M., "Ground test studies of the HIFiRE-1 transition experiment part 1: experimental results," *Journal of Spacecraft and Rockets*, Vol. 45, No. 6, 2008, pp. 1134–1148.

[58] Spalart, P., and Allmaras, S., "A one-equation turbulence model for aerodynamic flows," *30th aerospace sciences meeting and exhibit*, 1992, p. 439.