

An Assessment of the Laminar Hypersonic Double-Cone Experiments in the LENS-XX Tunnel

Jaideep Ray* and Patrick Blonigan†
Sandia National Laboratories, Livermore, CA 94550, USA

Eric T. Phipps‡ and Kathryn Maupin§
Sandia National Laboratories, Albuquerque, NM 87185

In this paper, we investigate two experimental datasets of laminar hypersonic flows, over a double-cone geometry, acquired in CUBRC’s LENS-XX expansion tunnel. These datasets have yet to be modeled accurately. A previous paper suggested that this could partly be due to mis-specified inlet conditions. The authors of this paper solved a Bayesian inverse problem to infer the inlet conditions of the LENS-XX test-section and found that in one case they lay outside the uncertainty bounds specified in the experimental dataset. However, the inference was performed using approximate surrogate models. In this paper, we revisit the experimental datasets and perform inversions for the tunnel test-section inlet conditions with a Navier-Stokes simulator. The inversion is deterministic and can provide uncertainty bounds on the inlet conditions under a Gaussian assumption. We find that deterministic inversion yields inlet conditions that do not agree with what was stated in the experiments. We also present an *a posteriori* method to check the validity of the Gaussian assumption for the posterior distribution. This paper contributes to ongoing work on the assessment of datasets from challenging experiments conducted in extreme environments, where the experimental apparatus is pushed to the margins of its design and performance envelopes*.

I. Nomenclature

h'_0	=	non-dimensional stagnation enthalpy of the flow
PDF	=	probability density function
P'_{Pitot}	=	non-dimensional Pitot pressure of the flow
$p(x)$	=	surface pressure on the double-cone
$q(x)$	=	surface heat-flux on the double-cone

*Corresponding author; jairay@sandia.gov; Technical Staff, Extreme-Scale Data Science and Analytics, MS 9152, Member, AIAA

†Technical Staff, Extreme-Scale Data Science and Analytics, MS 9152, Member, AIAA

‡Technical Staff, Scalable Algorithms, MS 1318

§Technical Staff, Optimization & UQ, MS 1318

*A version of this paper was presented at AIAA SciTech 2023; see Ref [45] for details

x	=	distance along the axis of the double-cone
U_∞	=	inlet velocity being inferred
U_{norm}	=	inlet velocity, used for normalization; set to 10^3 m/s
ρ_∞	=	inlet density being inferred
ρ_{norm}	=	inlet density, used for normalization; set to 10^{-4} kg/m ³
CI	=	Credibility Interval
M_∞	=	inlet Mach number
Re	=	Unit Reynolds number
T_∞, T_{v_∞}	=	inlet temperatures, translational and vibrational
y	=	model parameters to be inferred, i.e., an inlet flow condition $\{\rho_\infty, U_\infty\}$
MAP	=	Maximum <i>a posteriori</i> estimate i.e., most probable estimate
MCMC	=	Markov chain Monte Carlo

II. Introduction

EXPERIMENTAL datasets, meant for model validation, contain measurements of the phenomena that are the subject of the experiment, as well as a detailed quantitative description of the experimental environment and apparatus (so that the experiment may be modeled). The error/uncertainty in measurements that quantify the phenomena, as well as the experimental environment, are also supplied. Model validation efforts usually take the measurements and their (quantified) errors at face value. However, when simulating extreme environments, e.g., high enthalpy flows in shock tunnels with large cross-sections, one is forced to employ novel techniques, instrumentation, and apparatus, whose interactions may not be known very well, or which might be required to work at the edges of their operational envelope. In such scenarios, the quantifications of experimental errors/uncertainties that accompany a (model) validation dataset may be difficult to estimate. This causes complications when the dataset is used to validate a model — if the model predictions and experimental measurements do not agree, it is difficult to judge whether it is due to a shortcoming of the model or due to a mis-specification of the error/uncertainty bounds on the measurements in the experimental dataset. Ref. [1] describes one such case when an experimental dataset of laminar hypersonic flow over a $25^\circ/55^\circ$ double-cone, gathered in the Calspan-University at Buffalo Research Center’s (CUBRC’s) Large Energy National Shock Tunnel (LENS-I) wind tunnel, was used to validate a Navier-Stokes model of the flow. The model predictions failed to match experiments. The authors of Ref. [1] simulated the flow in the LENS-I shock-tunnel (the design was made available to them) and found that the inflow into the test section was in vibrational non-equilibrium (the total enthalpy of the flow was 3.71 MJ/kg). This was not known to the experimentalists, and was not reflected in the experimental dataset. When the correct boundary conditions, i.e., vibrational non-equilibrium, were specified, model predictions had far fewer

problems matching experimental measurements [1].

A similar problem was encountered in Ref. [2] when using double-cone data, collected in CUBRC’s LENS-XX expansion tunnel, to validate SPARC (Sandia Parallel Aerodynamics and Reentry Code), a Navier-Stokes model for non-equilibrium hypersonic flows. The experimental datasets contained pressure and heat-flux measurements over the same $25^\circ/55^\circ$ double-cone, but the flows in question had far higher total enthalpies (5.44 MJ/kg and 21.77 MJ/kg). The uncertainties in the inflow conditions were provided as part of the experimental datasets. The authors created an ensemble of inflow conditions consistent with the specified uncertainties and simulated flows over the double-cone, with the aim of bracketing (bounding) the measurements. They found that the distribution of model predictions did not agree with the distribution of measurements – they sometimes overlapped partially but were frequently disjoint. This implied that SPARC could not simulate some of the physics in the dataset. An obvious check would be to simulate the flow inside LENS-XX and compare the predicted inlet conditions (alternatively, the boundary conditions imposed by SPARC) with the specification that accompanied the validation dataset. However, the design details of the LENS-XX expansion tunnel are not publicly known and, consequently, the simulation cannot be performed. In addition, it is known that the flow inside LENS-XX is far more complex than in LENS-I; see Refs [3, 4] for a description. The authors in Ref. [5] simulated the flow inside LENS-X (the predecessor to LENS-XX) and found sections with steady and unsteady flows.

In order to address the problem of unknown inlet (inflow) conditions in LENS-XX, the authors in Ref. [6] employed an inversion approach. They assumed that the inflow was axisymmetric and then, conditional on pressure and heat-flux measurements on the fore-cone, as well as the total pressure and enthalpy of the flow, inferred the velocity, density and temperature at the inflow of the test section. The inversion was Bayesian and used a Markov chain Monte Carlo (MCMC; Ref. [7]) method to estimate the inflow conditions as a 3-dimensional probability density function (PDF). The PDF captured the uncertainty in the estimates due to the limited nature of the measurements, the measurement errors, and the missing physics in the SPARC model. For Case 1 (the 5.44 MJ/kg flow), they found that the measured and modeled PDFs overlapped partially, and the maximum *a posteriori* (MAP) estimate of the inflow conditions (loosely, the most probable estimate) fell within the experimental uncertainty bounds. For the second case (the 21.77 MJ/kg flow, called “Case 4”), the MAP estimate was clearly outside the experimental uncertainty bounds. However, these inferences were drawn using polynomial surrogates for SPARC’s predictions, a necessity due to the tremendous computational cost of MCMC. The polynomial surrogates are themselves approximations of the SPARC model, and the estimated inflow conditions are affected by the accuracy of the surrogates.

In this paper, we revisit the inference of the inflow conditions for Case 1 and Case 4 using the Navier-Stokes model in SPARC (and *not* any surrogates). The inference is cast as a deterministic inverse problem and solved using an optimization method that “tunes” the inflow conditions to maximize the agreement between the experimental data and model predictions. The flow separates at a point on the fore-cone of the double-cone configuration, and, just as in Ref. [6], we only use measurements that lie upstream of the separation point. A quantification of the uncertainty in

the estimated inflow conditions can be obtained by approximating the PDF (called the “posterior” PDF of the inflow conditions) as a multivariate Gaussian, whose variance-covariance matrix is provided by the inverse of the Hessian of the objective function, evaluated at the *optimal* point (inflow estimate which maximizes the match between measurements and predictions). We also develop a test to check the validity of approximating the posterior PDF as a Gaussian (also called the "Gaussian assumption"), and find some evidence that it is not justified.

This paper has two main contributions. The first is the assessment of the inferences drawn in Ref. [6] regarding the disagreement of inferred inflow conditions with that specified in the experimental dataset. Note that unlike Ref. [6], the native Navier-Stokes model in SPARC (and *not* its surrogate) is used in this study. An additional novelty of the paper is the first use of adjoint-based sensitivity analysis to solve an inverse problem in the reconstruction of a hypersonic flow; to date, adjoint-based inverse solutions have been used extensively only in design optimization and mesh adaptation for hypersonic flows. The second contribution of this paper is the development of an *a posteriori* test for the validity of the Gaussian posterior distribution that is computed via a simple extension of the deterministic inverse solution. In problems where this assumption is valid, our method provides a quick but approximate method to quantify uncertainties in the estimated quantities.

The paper is laid out as follows. In Sec. III, we review background literature on double-cone experiments and modeling efforts, as well as adjoint-based sensitivity and inverse solutions. In Sec. IV we perform a grid-refinement study and establish certain other numerical details. In Sec. V, we formulate the inverse problem, and in Sec. VI, we present our results. We first perform an inversion using a dataset measured in LENS-I which has been successfully modeled, and check whether our inversion yields inflow conditions that agree with those specified in the experimental dataset. Thereafter, we proceed to Case 1 and Case 4. We conclude in Sec. VII.

III. Literature Review

A. The Double-Cone Problem

Double-cone experiments have been conducted in the LENS-I [8, 9] and LENS-XX [10] facilities. The dimensions of the double-cone are in Ref. [11], Figure 1. The hypersonic flow sets up a complex shock pattern (see numerical Schlieren in Fig. 1), a separation zone, and complex shock-shock and shock-boundary layer interaction downstream of the separation zone. The flow is laminar throughout. The fore-cone is a 25° half-angle cone, which abuts a truncated 55° aft-cone. The fore-cone sets up an attached shock, whereas the aft-cone creates a detached shock. The two shocks interact to create a triple point and a transmitted shock that impinges on the aft-cone. The adverse pressure gradient set up by the transmitted shock feeds upstream through the (subsonic) boundary layer and causes the flow over the fore-cone to separate at $x \approx 0.05\text{m}$. The separated flow reattaches on the aft-cone. The separation zone, in turn, causes a separation shock that interacts with the transmitted shock and modifies it. Due to the existence of triple points, the

region near the intersection of the two cones contains contact surfaces (vortex sheets), as well as a separation zone. These are clearly labeled in Fig. 1.

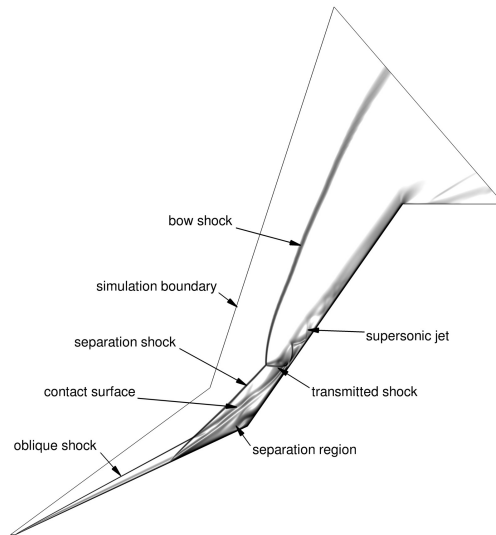


Fig. 1 A numerical Schlieren of the flow over a double-cone, taken from Ref. [6].

Double-cone experiments: Refs. [8, 9] describe double-cone experiments in the LENS-I shock tunnel. When the experiments were modeled using the Navier-Stokes equations, it was found that the model over-predicted heat-flux but did well with predictions of the pressure on the fore-cone, as well as the separation zone. Note that only one of the experiments, called “Run 35” (see specifications in Table 1), was modeled under the assumptions of an axisymmetric flow at the inlet of the shock tunnel that was also in vibrational and reactive equilibrium. However, it was later discovered that the inflow into the test-section of the shock tunnel was in strong vibrational non-equilibrium and was, in addition, frozen [1]. This discovery was made when the authors of Ref. [1] simulated the flow inside the entire tunnel and examined the flow state at the test-section inlet. This necessarily requires knowledge of the geometry of the nozzle, which was made available to the authors. When the inlet flow conditions (henceforth, inflow conditions) were adjusted to reflect vibrational non-equilibrium, Computational Fluid Dynamics (CFD) simulations using the Navier-Stokes model for Run 35 matched experimental data quite well [1]. The LENS-I experiments were conducted in a N_2 environment at very low flow density and total enthalpy, and a perfect-gas model, with no reactions, was quite appropriate.

The double-cone experiments were also conducted in the LENS-XX *expansion* tunnel [10], where temperatures do not exceed about 3000K [3]. Consequently, the thermochemical non-equilibrium issues that affected the LENS-I inflow conditions do not arise. Six double-cone experiments at different flow conditions (total enthalpies and pressures) were conducted. Two of these experiments, called Case 1 and Case 4, are investigated in this paper and are summarized in Table 1. The uncertainty bounds on the measured inflow conditions mentioned in the table were sourced from a private conversation with T. Wadhams, Director, CUBRC Aerosciences Group, CUBRC Headquarters, 4455 Genesee

Street, Buffalo, NY 14225, 2018 and some are also mentioned in Ref. [11]. Attempts to model these six LENS-XX double-cone experiments have met with varying levels of success. They are described below.

Recently, a paper describing the repetition of Cases 1 and 4 in University of Queensland's X3 expansion tunnel was published [12]. The experiment was performed with a scaled-down double-cone test article (less than half the original size). Unlike the LENS-XX experiments, the X3 experiment had four "lines" of sensors along the double-cone, arranged at 90° circumferentially, to detect asymmetries in the flow. As in the case of the LENS-XX experiments, the authors identified, using one pressure and heat flux sensor mounted on the double-cone, a period of "steady flow" sandwiched between periods of transient flow. The "steady flow" periods in the two tunnels are approximately the same (about 0.3 milliseconds; see Refs.[10, 12]). The average heat flux computed using measurements obtained during the "steady flow" period agreed with the LENS-XX results. However, plots of the transient heat flux along the axis of the double-cone show that it is evolving, monotonically decreasing in time over the fore-cone and rising elsewhere. Further, the separation zone continues to shrink. The four lines of sensors reveal that during the "steady" period (as determined using pressure measurements), the separation zone starts off being strongly asymmetric and progressively relaxes into a more symmetric shape, before the second transient process ends the "steady" period. If the X3 experiments are a true reflection of the processes extant in the LENS-XX experiments, then the measurements collected during the "steady" period were neither steady nor axisymmetric, two assumptions made by all modeling efforts to date.

Modeling of double-cone experiments: Ref. [11] provides a good description of modeling studies of the LENS-XX double-cone experiments, including visualizations of the flowfields. The authors showed that the accurate pressure predictions (i.e., they match LENS-XX measurements) on the fore-cone could be achieved with both perfect-gas and non-equilibrium flow models, but the heat-flux predictions were too low, even when accounting for measurement errors (5% for surface pressure and 7% for surface heat-flux on the double-cone; measurement uncertainty estimates from Ref. [13]). The separation zone was predicted to be smaller, primarily because of late separation. Studies have explored the correct thermochemical modeling approach to adopt, often with conflicting results [14, 15]. The high-enthalpy Case 4 shows the largest differences in $p(x)$ and $q(x)$ predictions, depending on the choice of the thermochemical (equilibrium/frozen/non-equilibrium) modeling of the flow, while Case 1 shows the least. Here $p(x)$ and $q(x)$ are the pressure and heat-flux on the double-cone surface, and x is the axial distance from the double-cone nosetip. Aft of the reattachment point, no model does particularly well, with errors increasing with the total enthalpy of the flow. In Ref. [16] the authors examined whether a sophisticated model for the coupling between vibrational and reactive (i.e., dissociation) non-equilibrium [17] could bridge the gap between measurements and predictions, but found that it worked no better than Park's simple model [18]. (The discrepancy in the separation zone size, measured versus predicted, varied between 2% and 20%, depending on the total enthalpy of the flow.) This lack of sensitivity of the results to more sophisticated coupling between vibrational and chemical non-equilibrium has been corroborated by other studies for the same LENS-XX double-cone experiments [14, 19, 20]. The main consequence of changing the thermochemical

model has been to change the size of the separation zone [14, 15, 19, 20]. The authors in Ref. [14] used thermochemical models whose parameter values were computed from more fundamental simulations using forced harmonic oscillators and quasi-classical trajectory (QCT) calculations, and obtained a very good comparison with the separation zone size for Case 4. For the integrated heating (which suppresses local disagreements between predicted heat fluxes and their experimental counterparts), the prediction errors ranged from 5% to 15%, with larger errors occurring at higher total enthalpies. Ref. [15] compared predictions of pressure and heat fluxes on the double-cone to measurements using normalized root-mean-square metrics and found discrepancies of around 20%, clearly outside measurement uncertainty bounds. Researchers have also studied whether catalytic effects could explain the disagreement between measurements and predictions [21] and found them to be minor. LENS-XX experiments have also been the subject of a “blind” modeling comparison, described in Ref. [22]. Here, five different CFD researchers attempted to model the experiments, with very different results, which also did not agree with measurements. The only difference between these studies were the numerical method (e.g., the slope and state limiters) used by them. This numerical issue was explored in Ref. [21], which found that the effect of numerical viscosity was too small to bridge the gap between model predictions and experimental measurements.

Uncertainty quantification and double-cone experiments: The authors of Ref. [2] performed a forward propagation of input uncertainties, as described in Sec. II. The aim was to test whether the distribution of pressure and heat-flux predictions from the ensemble bracket the experimental measurements. They could not, but a $\pm 15\%$ perturbation of the inflow condition at the test-section produced predictions of the surface pressure, heat flux, total enthalpy and Pitot pressure that bounded the experimental data (also called “covering” the experimental data). This raised the possibility that modifying the inflow conditions, outside of the limits placed by the stated experimental uncertainty bounds, could allow CFD simulations to reproduce experimental measurements. This led to the Bayesian inversion studies in Ref. [6] which were described in Sec. II. The authors developed an inference framework, targeted the three experiments listed in Table 1, and inferred the inflow conditions (density, velocity, temperature and vibrational temperature) based on them. The authors found that the MAP estimates for density and velocity were well within the stated experimental error bounds for Run 35, but not for Case 1 and Case 4. They also found that the measurements used in the inference (surface pressure on the double-cone $p(x)$, surface heat-flux $q(x)$, the total enthalpy of the flow h_0 and the Pitot pressure P_{Pitot}) carried very little information on the inflow temperature of the flow which, consequently, could not be estimated. They also investigated whether shortcomings of the thermochemical models could explain the model versus experiment disagreements and found weak evidence for it. However, when the “correct” inflow conditions, as inferred from the fore-cone measurements, were used in CFD predictions, the disagreement between predictions and measurements remained downstream of the separation point, implying that while Bayesian inference found some of the causes of modeling errors, it was far from a comprehensive solution.

The SPARC flow simulator: The studies in this paper are conducted using SPARC, a CFD solver. A description

Table 1 Inflow conditions and their uncertainties (“Error”). Run 35 was conducted in the LENS-I tunnel and is described in Refs. [1, 25]. Cases 1 and 4 were conducted in LENS-XX and are described in Ref. [6, 11]. The double-cone was held that temperature T_{wall} during the experiment.

Test	U_∞ (m/s)	ρ_∞ (g/m ³)	T_∞ (K)	T_{v_∞} (K)	Re (m ⁻¹)	M_∞ (±3%)	h_0 (MJ/kg)	P_{Pitot} (kPa)	T_{wall} (K)	Gas
Error	(±3%)	(±7%)	(±3%)				(±5%)	(±5%)		
Run 35	2545	0.5848	98.27	2562	14.3×10^4	12.59	3.71	3.55	296.11	N ₂
Case 1	3246	0.499	175	175	0.14×10^6	12.2	5.44	5.1	300.00	Air
Case 4	6479	0.964	652	652	0.20×10^6	12.82	21.77	39.5	300.00	Air

of SPARC is in Ref. [6]. Grid resolution [23] and solution verification [24] studies have been performed for SPARC, and it has been used for uncertainty quantification studies, within the context of double-cone experiments [2, 6].

B. Adjoint-Based Inverse Solvers

Adjoint-based sensitivity analysis is used to compute gradients for partial differential equation-constrained optimization problems (strictly, they can also be used for problems governed by algebraic systems), including design and inverse problems. Adjoint sensitivity analysis has been used for solving inverse problems for a wide range of applications [26, 27], but not, to the authors’ knowledge, in hypersonic flows. Instead, adjoint-based sensitivity analysis has been used for design optimization [28], mesh adaptation and uncertainty propagation.

NASA has two CFD codes which support adjoint-based design optimization and mesh adaptation for inviscid and viscous hypersonic flows. CART3D, a Cartesian grid inviscid flow solver, leverages the adjoint for optimization and mesh adaptation on a wide range of geometries and flight conditions [29]. CART3D has been applied to numerous inviscid hypersonic flows including those around reentry capsules [30] and asteroids [31]. Another NASA solver, FUN3D, has an adjoint capability that has been used for shape optimization and mesh adaptation on a wide range of viscous hypersonic flows, including non-equilibrium flows [32–34]. It is a discrete adjoint capability which leverages analytically differentiated Jacobians for verification and validation purposes [33]. The adjoint capability of Stanford University’s SU2 open-source CFD solver has been demonstrated on a range of inviscid and viscous hypersonic flows [35]. Initially, SU2 leveraged a continuous adjoint for mesh adaptation and shape design of reentry vehicles and scramjet inlets [36–38], but more recent work leverages a discrete adjoint capability [39, 40]. Apart from these examples, recent work by Damm et al. demonstrated a discrete adjoint solver for scramjet inlet design [41]. Finally, Lockwood et al. developed and demonstrated a discrete adjoint to solve optimization problems as part of uncertainty propagation studies on various flow field properties in non-equilibrium hypersonic flows around bluff bodies [42–44].

IV. Sensitivity Analysis

In this section, we describe the effect of grid resolution on the pressure $p(x)$ and heat-flux $q(x)$ on the surface of the double-cone, as well as on their sensitivities S_α to perturbations in inlet conditions, $\alpha \in \{\rho_\infty, U_\infty, T_\infty, T_{v_\infty}\}$. We define the sensitivity S_α as

$$S_\alpha = \frac{\alpha}{Z} \frac{\partial Z}{\partial \alpha} \quad (1)$$

where $Z \in \{p(x), q(x)\}$ are two of four Quantities of Interest (QoI) in this paper. Thus, this is Local Sensitivity Analysis, with the gradients in Eq. (1) computed for the nominal test-section inlet conditions tabulated in Table 1. The method for computing these gradients is described in Ref. [45]. The other two QoIs are the total enthalpy h_0 and Pitot pressure P_{Pitot} of the incoming flow, both of which are independent of the mesh resolution. The flow is assumed to be axisymmetric, and we consider three different grids *viz.*, 128×256 (called the ‘‘Coarse’’ mesh), 256×512 (‘‘Medium’’ mesh) and 512×1024 (‘‘Fine’’ mesh; figure in Ref. [45]). These are the same meshes which were used in the grid resolution studies in Ref. [23].

Fig. 2 (top row) plots the surface pressure $p(x)$ and its sensitivities on the left and the surface heat-flux (and its sensitivities) on the right. The computations are performed on the Fine mesh using the inlet conditions corresponding to Case 1. A 5-species 17-reaction chemical mechanism [46] is used, and the model can accommodate flows in vibrational non-equilibrium. The point where the flow separates, at $x = 0.065\text{m}$ is denoted with a vertical line. The QoI is plotted with a solid line, and is represented by the y-axis; the sensitivities are plotted with dashed lines and are represented by the 2nd y-axis. We see that $p(x)$ is constant in the laminar, attached flow region, as are its sensitivities. In contrast, $q(x) \propto x^{-1/2}$, though its sensitivities are also constant upstream of the separation point. Downstream of the separation *i.e.*, of the vertical line, we see significant variations in the sensitivities of $p(x)$ and $q(x)$ primarily due to reattachment of the separated flow and its interactions with complex shock reflections. In the bottom row of the same figure, we plot the same QoIs and their sensitivities, but computed on the Coarse mesh. We see large differences in the sensitivities downstream of the separation point, but not upstream of it. This is not surprising since small changes in mesh resolution can change the location where flow or shocks impinge on a surface. In addition, we see that the sensitivity with respect to temperature, *i.e.*, S_{T_∞} , is very small compared to the sensitivities with respect to ρ_∞ and U_∞ . This is not very surprising either, given that most of the energy in the incoming hypersonic flow is kinetic, not thermal.

Fig. 2 (bottom row) plots the same quantities as the top row, but computed on the Coarse mesh. We see oscillations in the sensitivity of pressure and heat flux even on the fore-cone with the attached flow. Sensitivity calculations generally demand finer meshes for accurate computations. The figure leads us to two conclusions:

- 1) Since the sensitivity of $p(x)$ and $q(x)$ to T_∞ is so small, it will be infeasible to infer it (*i.e.*, T_∞) from surface measurements on the double-cone (an observation also borne out in Ref. [6]). Consequently, we will only estimate (ρ_∞, U_∞) .

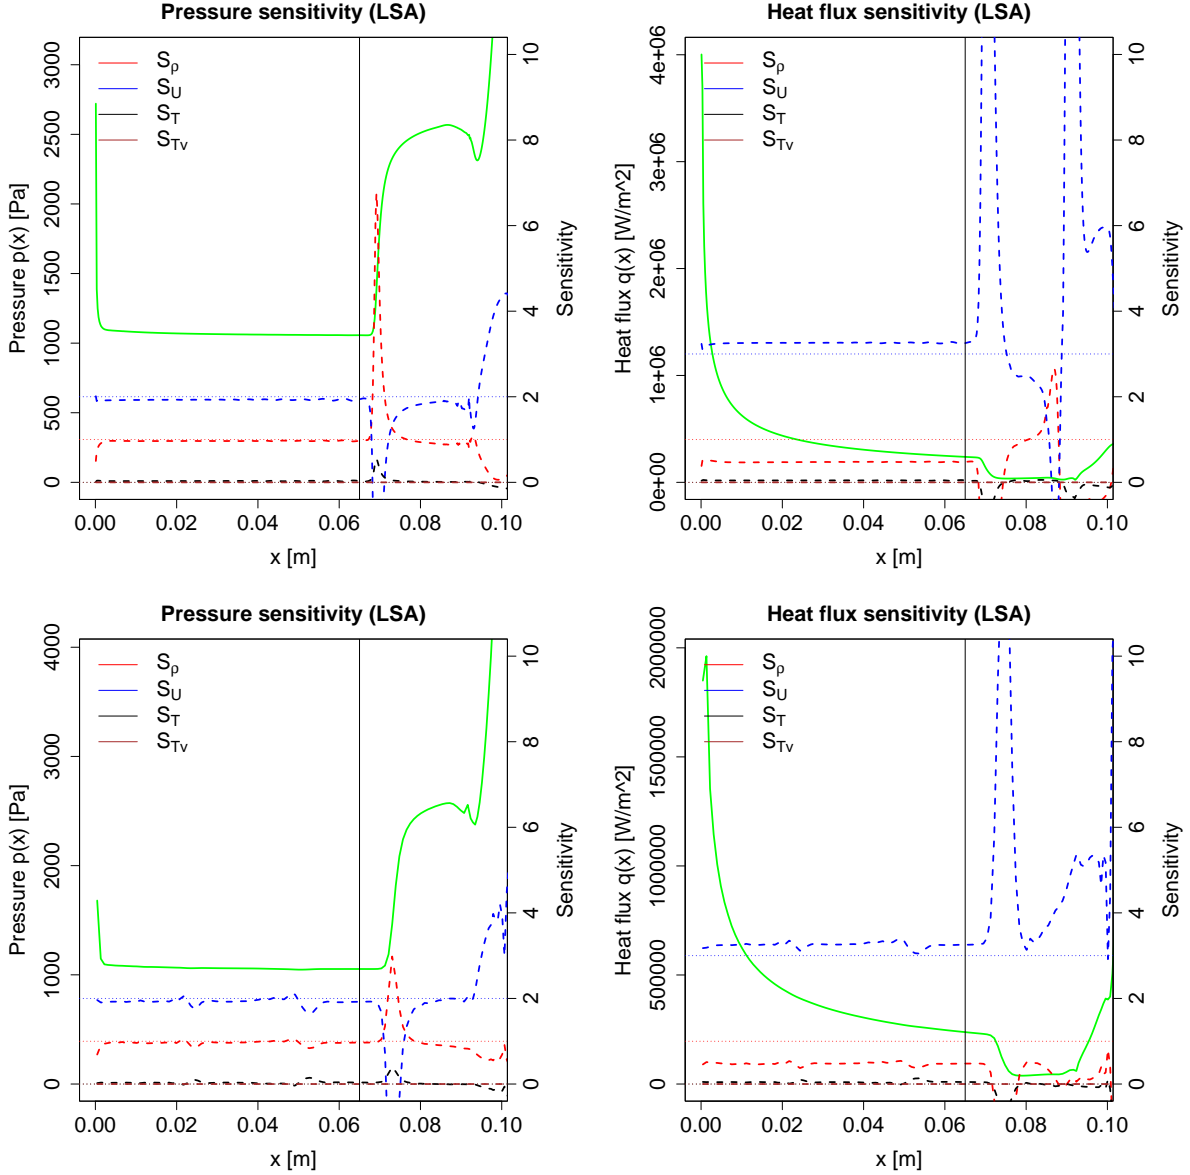


Fig. 2 Top: Predicted pressure $p(x)$ (left) and heat-flux $q(x)$ (right) for Case 1, plotted with a green line. Predictions were obtained on the Fine mesh. The point of flow separation is denoted with a vertical line. On the second y-axis we plot the sensitivity S_α of the relevant quantity of interest w.r.t. $(\rho_\infty, U_\infty, T_\infty, T_{v_\infty})$. Bottom: The same plots but computed on the Coarse mesh.

- 2) Given the large changes in sensitivities downstream of the separation zone due to changes in grid resolution, our inverse problem will use only those model predictions that lie upstream of the separation region, i.e., $x \leq 0.065\text{m}$ for Case 1. Using measurements upstream of the separation point offers the possibility that grid resolution issues will not significantly affect the accuracy of the inverse problem's solution.

V. Formulation of the Inverse Problem

In this section, we pose the inverse problem for the flow density (ρ_∞) and velocity (U_∞) at the test-section inlet, conditional on pressure ($p(x)$) and heat-flux ($q(x)$) measurements, as well as the total enthalpy (h_0) and Pitot pressure (P_{Pitot}) of the incoming flow. Fig. 2 plots the pressure $p(x)$ and heat-flux $q(x)$ on the double-cone surface for Case 1. The point of flow separation is also plotted with a vertical line. We employ only those measurements that are upstream of the separation point, since the flow physics there is well understood. The probes (sensors) that collect measurements in this part of the flow are referred to as the ‘‘calibration’’ probes. Fig. 2 also shows that while the pressures at the calibration probes are constant, the heat-flux $q(x)$ shows a variation in space. In addition, $p(x)$ and $q(x)$ assume very different values, and need to be scaled down to the same order-of-magnitude to be usable in an inverse problem.

For pressure, the scaling s_p is defined as the reciprocal of the mean of all measured pressures $\tilde{p}_i, i = 1 \dots N_p$, where N_p is the number of pressure calibration probes. Henceforth a tilde (“~”) denotes a measured quantity. Since the heat-flux $q(x)$ varies as $x^{-1/2}$ (see Chapter 9 in Ref. [47]), we compute a distance-weighted $Q(x_i; u, y) = q(x_i; u, y) \times \sqrt{x_i}$ which is approximately constant over the attached laminar flow over the fore-cone. The scaling s_Q is defined as the reciprocal of the mean of $\tilde{Q}_i, i = 1 \dots N_q$. N_q is the number of calibration probes for heat-flux measurements. The scalings s_h and s_P are the reciprocals of \tilde{h}_0 and \tilde{P}_{Pitot} . N_q and N_p vary for the three experiments in Table 1. N_p is 5 for Run 35, 3 for Case 1, and 4 for Case 4, while N_q is 9 for Run 35, 15 for Case 1, and 13 for Case 4.

Posing the optimization problem: Since the problems of interest are steady-state flows, we consider the following deterministic optimization problem,

$$\min_y g(u, y) \quad \text{s.t.} \quad f(u, y) = 0, \quad (2)$$

where $u \in \mathbb{R}^n$ is the discrete flow solution consisting of the unknown degrees-of-freedom across the mesh, $y \in \mathbb{R}^m$ is a set of model parameters (for us $y = \{\rho_\infty, U_\infty\}$ with $m = 2$), $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the discretized residual, and $g : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^+$ is a scalar-valued objective function. We formulate the optimization problem based on the cell-centered finite volume approach using block-structured meshes, in which case u contains the value of the degrees-of-freedom at each cell center (density of each chemical species, flow velocity, temperature, and molecular vibration temperature), and f is the finite volume residual (thus n is given by the number of mesh cells times the number of degrees-of-freedom per cell, which varies across the meshes (i.e., different resolutions) and cases considered here). The objective function g quantifies the misfit between simulation quantities-of-interest and experimentally measured values:

$$g(u, y) = \sum_{i=1}^{N_q} [s_q(q(x_i; u) - \tilde{q}_i)]^2 + \sum_{i=1}^{N_p} [s_p(p(\hat{x}_i; u) - \tilde{p}_i)]^2 + [s_h(h_0(y) - \tilde{h}_0)]^2 + [s_P(P_{\text{Pitot}}(y) - \tilde{P}_{\text{Pitot}})]^2. \quad (3)$$

Here $q(x_i; u)$ and $p(\hat{x}_i; u)$ are the (distance-weighted) heat-flux and pressure computed at a set of calibration probe

locations $\{x_i\}$ and $\{\hat{x}_i\}$, respectively (the probe locations are not, in general, the same), $h_0(y)$ is the total enthalpy of the flow, $P_{Pitot}(y)$ is the Pitot pressure of the flow, \tilde{q}_i , \tilde{p}_i , \tilde{h}_0 , and \tilde{P}_{Pitot} are the corresponding measured values, and s_q , s_p , s_h and s_P are scaling values chosen for each problem to equilibrate the contribution of each term to the overall objective function.

Equation (2) falls into the category of PDE-constrained optimization, for which numerous methods have been developed in the literature. In this work, we focused on the so-called reduced space approach that transforms Eq. (2) into an unconstrained optimization problem

$$\min_y h(y), \quad h(y) = g(u(y), y) \quad \text{s.t.} \quad f(u(y), y) = 0 \quad (4)$$

where evaluation of $h(y)$ first involves solving the nonlinear equations $f(u(y)) = 0$ for u , then evaluating the original objective function $g(u(y))$. Under suitable differentiability assumptions for f and g , the corresponding reduced gradient can be computed as

$$\nabla h(y) = - \left(\frac{\partial f}{\partial y} \right)^T w + \left(\frac{\partial g}{\partial y} \right)^T \quad \text{s.t.} \quad \left(\frac{\partial f}{\partial u} \right)^T w = \left(\frac{\partial g}{\partial u} \right)^T \quad (5)$$

which involves solving a single linear system for the adjoint variable $w \in \mathbb{R}^n$. A variety of optimization approaches are available for solving Eq. (4). In this work, we use the truncated conjugate gradient method with trust-region globalization [48, Ch. 4] and BFGS Hessian approximation [48, Ch. 6] as provided by the ROL package [49].

Given a set of inflow conditions y , each step of the optimization method requires computing a new steady-state solution u satisfying $f(u, y) = 0$. Due to numerical difficulties in computing these steady-state solutions, SPARC employs a first-order Backward Euler (BDF1) implicit time integration method, starting from an initial uniform flow, to compute them by solving $M\dot{u} + f(u, y) = 0$, where $M \in \mathbb{R}^{n \times n}$ is the mass matrix, until steady-state is reached with $\dot{u} = 0$. This method is not time-accurate. The mass matrix M is a diagonal matrix whose entries are cell volumes corresponding to each conserved quantity in u . Therefore, M can be scaled and added to the main diagonal of a sparse matrix containing the Jacobian of $f(u, y)$ to form the matrix in the linear system solved as part of the Newton iteration at each implicit time step.

Computing sensitivities/gradients: Once a steady-state solution is computed, the objective function gradient is then given by Eq. (5). We employ Automatic Differentiation (AD) provided by the Sacado package [50–52] to analytically compute $\partial f/\partial y$, $\partial g/\partial y$, $(\partial f/\partial u)v$ and $(\partial g/\partial u)v$ for any vector $v \in \mathbb{R}^n$ without hand-coding. Using this capability, we then employ graph coloring [53–55] provided by the Zoltan package [56] to construct the true Jacobian $\partial f/\partial u$, which is then explicitly transposed to construct the adjoint operator $(\partial f/\partial u)^T$. Since the heat-flux and pressure computations involve a similar stencil as the finite volume residual, the same coloring procedure is used to construct $(\partial g/\partial u)^T$. While numerous solver strategies are available for solving Eq. (5), in this work we apply GMRES [57]

using a block tridiagonal/line implicit solver [58] applied to the transpose of SPARC’s approximate Jacobian as a preconditioner. Due to ill-conditioning of the true adjoint matrix, we furthermore apply multiple steps of iterative refinement to obtain a solution with small residual. Details of the adjoint calculation are in Ref. [45].

Estimation of the Gaussian posterior: Under the assumptions that the heat-flux, pressure, total enthalpy, and Pitot pressure at the calibrated values differ from the experimental measurements by additive, Gaussian noise, and their dependence on the calibrated parameters is approximately linear in the vicinity of the optimal parameter values, the posterior of the calibrated parameters is approximately Gaussian, and solution of Eq. (2) is equivalent to a maximum likelihood estimation procedure for the mean of the posterior. It is straightforward to show the covariance of the posterior is then given by the inverse Hessian of the negative log-likelihood, i.e., the objective function Eq. (3). At the completion of the calibration process, we estimate the Hessian through a first order finite difference approximation, differencing the gradient Eq. (5) with a relative finite difference step size of 10^{-7} , and then compute its inverse through Gaussian elimination. The relative finite difference step size was determined empirically, with successively smaller step sizes until the Hessian became insensitive to it. However, before we use the Gaussian posterior to represent uncertainty in the solution to the inverse problem, we will devise a test to check whether the Gaussian assumption is valid. In addition, our inversion scheme is unconstrained[†], as we suspect that $\{\rho_\infty, U_\infty\}$ may lie outside the bounds of the measurement errors specified by the experimental dataset, but have no principled way of deciding how large the deviation may be. In a Bayesian interpretation, this would be equivalent to specifying a non-informative prior e.g., a uniform distribution with infinite bounds or a Jefferey’s prior on the mean of the Gaussian posterior distribution.

VI. Results

A. Verification of the Inversion Method using Synthetic Data

The first check for our inversion method is to test it in the absence of any model-form error. This implies a synthetic data test i.e., “measurements” of pressure and heat-flux that are generated by running the forward simulation and corrupting it with Gaussian noise. The synthetic measurements are generated by running the forward problem using SPARC with the same inflow conditions as Run 35 (see Table 1). Multiplicative noise, sampled from a uniform distribution i.e., $\epsilon \sim \mathcal{U}(1 - \eta, 1 + \eta)$ is used with SPARC predictions to generate the synthetic measurements. We use $\eta = 5\%$ for pressure and 7% for heat-flux, in line with the measurement errors reported for Run 35 [13]. Inversions are the performed for (ρ_∞, U_∞) and we quantify the uncertainty in the estimates with a Gaussian posterior with a mean provided by the optimal (ρ_∞, U_∞) and a covariance matrix computed as the inverse Hessian. These results are shown in Fig. 3. The Gaussian posterior at the top left shows that the estimated density and velocity are inversely correlated. In addition, the density shows a tremendous amount of estimation uncertainty (under a Gaussian assumption), but the

[†]Even positivity of $\{\rho_\infty, U_\infty\}$ is not explicitly enforced as deviations from their initial guesses are small enough to not require bounds.

velocity is accurately estimated. The other two plots in Fig. 3 show the marginalized posteriors. The true and estimated inflow quantities are indistinguishable. The values of ρ_{norm} and U_{norm} are in Table 2.

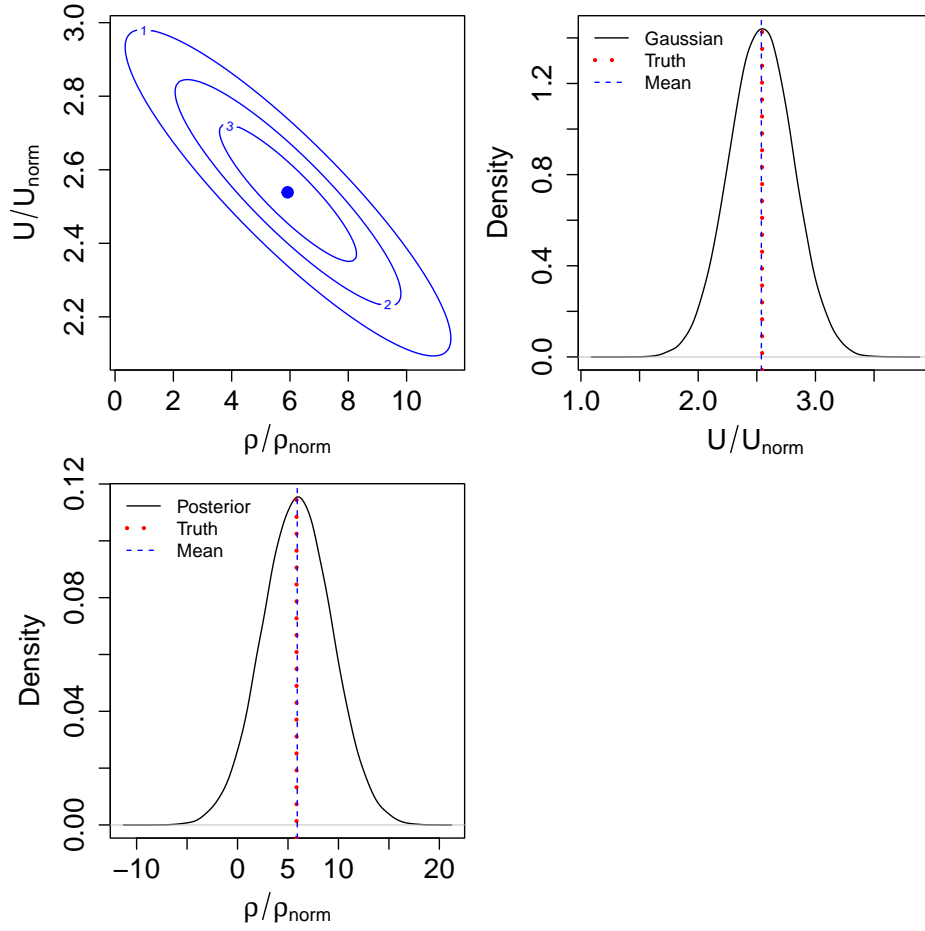


Fig. 3 Gaussian posterior PDFs from the synthetic data inversion. **Top left:** Joint $(\rho_{\infty}/\rho_{norm}, U_{\infty}/U_{norm})$ PDF. The Optimal estimate from the deterministic inversion and the true values are also plotted; they overlap. **Bottom left:** Marginal posterior PDFs for $\rho_{\infty}/\rho_{norm}$, from the Gaussian posterior obtained from the deterministic inversion (“Posterior”). The true value (“Truth”; dotted vertical line), and the optimal estimate (“Mean”; vertical dashed line) from deterministic inversion are also plotted, and are close. **Top right:** Same as the bottom left subfigure, for the normalized inflow velocity. The values of ρ_{norm} and U_{norm} are in Table 2.

Per the discussion in Section IV, the sensitivity of pressure and heat-flux to variation in the inflow temperature is very small and we expect that temperature should be difficult to estimate from the synthetic data. To this end, we perform a three-dimensional inversion for $(\rho_{\infty}, U_{\infty}, T_{\infty})$. The inversion did not converge to a final value, primarily because T_{∞} kept drifting. The evolution of $(\rho_{\infty}, U_{\infty}, T_{\infty})$ as a function of the optimizer iterations is shown in Fig. 4. We see that ρ_{∞} and U_{∞} converge to their true value quite quickly while T_{∞} diverges in the absence of any constraining effect of $p(x)$ and $q(x)$. This implies that if T_{∞} cannot be inferred when there is no model-form error in the inverse problem, there is little point in trying to do so with real data. Therefore, in the rest of the discussion, we will focus on inferring $(\rho_{\infty}, U_{\infty})$ from data.




Fig. 4 The evolution of ρ_∞ (left), U_∞ (middle) and T_∞ (right) (inflow quantities) with optimizer iteration. The horizontal green line is the true value.

Note that in Fig. 3 we have plotted the posterior for the estimated (ρ_∞, U_∞) under a Gaussian assumption. In the figures below, we will omit the Gaussian posterior; this is because the Gaussian assumption is incorrect and leads to posterior distributions that are too wide i.e., the uncertainty is over-estimated. This is seen empirically in Ref. [45] and is captured (somewhat incompletely because of the paucity of measurements) via a test described in Sec. VI.E. We will therefore only use the optimal estimates from the inversion. These are also known as the maximum likelihood estimates of the inverse problem.

B. Validation of the Inversion Method using Run 35

We next demonstrate the utility of our inversion technique (described in Sec. V) to estimate the inlet conditions of the test-section, using pressure and heat-flux measurements on the double-cone. Run 35 is used to test the utility of the inversion technique because (1) the flow has been successfully modeled [1], i.e., the stated experimental conditions can reproduce experimental measurements and (2) the experiment was conducted in an N_2 environment, which, along with its low total enthalpy (see Table 1), ensures that complications due to dissociated and reactive flow do not arise. Note that the flow conditions stated in Table 1 are at variance with the original paper where the experiment was first described [8]. The reason is as follows. The inflow conditions for Run 35 (also called Case D in Ref. [8]) were computed from shock tunnel measurements under the assumption that the flow was in equilibrium i.e., $T_\infty = T_{v_\infty}$. When the authors of Ref. [1] simulated the flow inside the nozzle of the LENS-I shock tunnel (where Run 35 was conducted), they found that the flow entering the test section was in vibrational non-equilibrium i.e., $T_\infty \neq T_{v_\infty}$, and frozen, i.e., the vibrational energy was partitioned out and did not play much of a role in the pressure and heat-flux measurements on the double-cone. The frozen flow inflow conditions are in Ref. [1] and are used in this paper, and the flow model assumes an ideal gas model. In addition, the surface temperature of the double-cone was maintained at 296.11 K, instead of 300 K for Cases 1 and 4.

This problem has also been solved using surrogate models and MCMC in Ref. [6], and we compare the exact

Table 2 Summary of posterior PDFs for ρ_∞/ρ_{norm} and U_∞/U_{norm} for Run 35, Case 1 and Case 4. The experimental specification is also provided. The uncertainty in the experimental measurements is 7% for inflow density and 3% for inflow velocity; the uncertainty bounds are in parentheses. $\rho_{norm} = 1.0 \times 10^{-4} \text{kg/m}^3$ and $U_{norm} = 1.0 \times 10^3 \text{m/s}$. MAP refers to the MAP estimate obtained from the MCMC solution in Ref. [6]. Optimal values are obtained from the deterministic method described in this paper.

Summary	ρ_∞/ρ_{norm}			U_∞/U_{norm}		
	MAP	Optimal	Experiment	MAP	Optimal	Experiment
Run 35	5.83	5.89	5.848 (5.43, 6.26)	2.548	2.506	2.545 (2.47, 2.62)
Case 1	4.897	4.33	4.999 (4.64, 5.34)	3.34	3.54	3.246 (3.15, 3.34)
Case 4	8.608	8.619	9.84 (9.15, 10.52)	7.060	6.95	6.479 (6.23, 6.67)

posterior distribution from the MCMC solution versus the point estimates obtained from deterministic inversion. The deterministic inversion was started with a guessed $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ that was well outside the uncertainty bounds tabulated in Table 1. In Fig. 5 (top left) we plot the joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ posterior PDF from the MCMC solution [6] as contours. The MAP (Maximum A Posteriori) estimate (from the MCMC solution), the Optimal one (from the deterministic inversion) and the experimental values are also plotted in the same figure. In Fig. 5 (bottom left) we plot the marginalized posterior PDF for the normalized inflow density, along with the MAP estimate and its deterministic counterpart, and the experimental value with its measurement error bars. Fig. 5 (top right) contains the same information, but for the normalized inflow velocity. We see that the MAP and the optimal estimates of the inflow density and velocity are close to each other. In addition, they both lie within the support of the marginalized posterior PDFs as well as within the experimental error bounds. These results are summarized in Table 2. *We conclude that the deterministic inversion can infer inflow conditions correctly, i.e., within the experimental error bounds. Further, it is close to the MAP estimate from the MCMC solution and within the support of the posterior distribution.*

C. Case 1

Having tested our inversion capability on low-enthalpy Run 35, which had been previously modeled numerically [1], we proceed to Case 1. Case 1 is of moderate enthalpy (see Table 1), and we model it with a five-species, 17-reaction mechanism [46]. We also assume vibrational non-equilibrium, i.e., the governing equations have two separate temperatures, one each for translational and vibrational energy. The flow at the inlet is assumed to be in vibrational equilibrium ($T_\infty = T_{v_\infty}$). The experiment was conducted in the LENS-XX expansion tunnel. An inversion for the inflow conditions has been performed for this experimental dataset, using Bayesian inference, surrogate models, and MCMC [6]. The study found that the inferred inflow density (MAP value) was within the stated experimental error bounds but the inferred inflow velocity was (barely) outside the measurement error bounds.

In Fig. 6 (top left) we plot the joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ posterior PDF from the MCMC solution [6] as contours. The MAP estimate (from the MCMC solution), the Optimal one (from the deterministic inversion) and the experimental

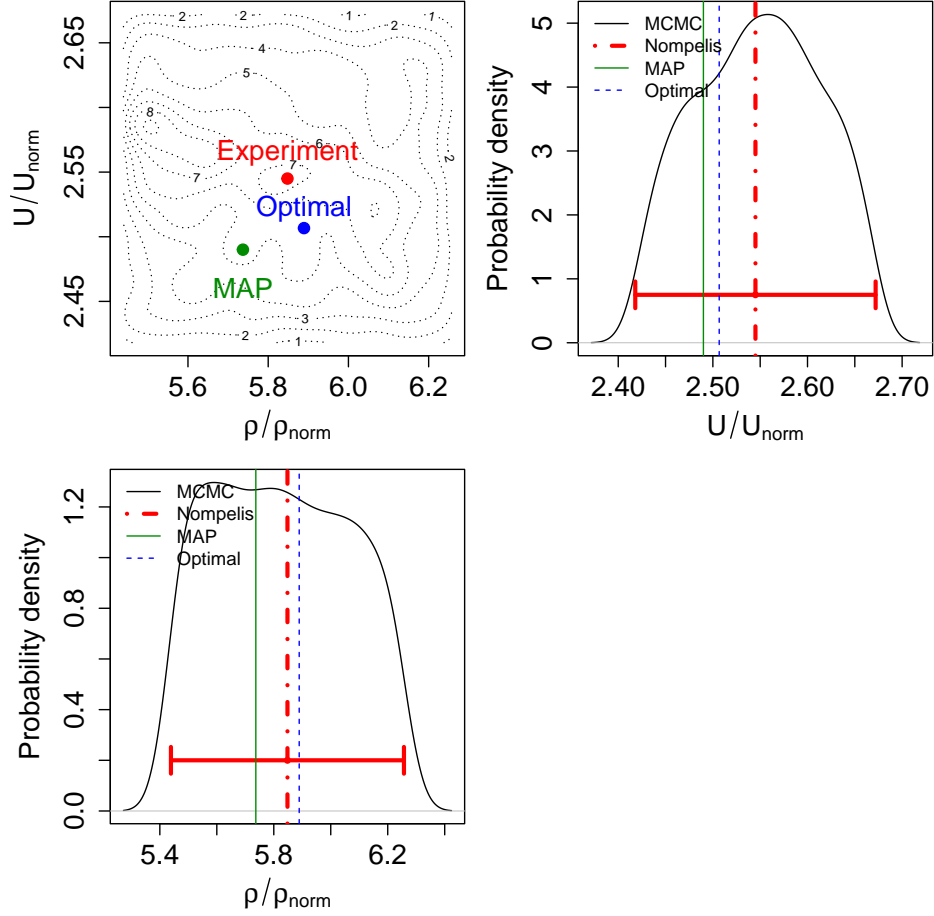


Fig. 5 Posterior PDFs for Run 35. Top left: Contours of the joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ posterior distribution from MCMC solution, plotted with dotted lines. The MAP estimate, the Optimal from the deterministic inversion and the experimental value are also plotted. Bottom left: Marginal posterior PDFs for ρ_∞/ρ_{norm} , from the MCMC solution. The nominal/experimental specification (“Nompelis”; dash-dotted vertical line, with measurement error bars), the MAP estimate from the MCMC solution (“MAP”; solid vertical line) and the Optimal estimate from deterministic inversion (“Optimal”; vertical dashed line) are also plotted. Top right: The same as the bottom left subfigure, but for the normalized velocity.

values are also plotted in the same figure. In Fig. 6 (bottom left) we plot the marginalized posterior PDF for the normalized inflow density, along with the MAP estimate and its deterministic counterpart, and the experimental value with its measurement error bars. Fig. 6 (top right) contains the same information, but for the normalized inflow velocity. We see that the MAP estimate of the normalized inflow velocity and the optimal estimate lie outside the measurement error bounds; further, both the estimates indicate that the inflow velocity is larger than the experimental value. In contrast, the MAP estimate of the inflow density agrees with the experimental data (i.e., lies within the measurement error bounds), whereas the optimal estimate disagrees with it, and is smaller than the experimental value. These results are tabulated in Table 2. Note that the agreement of the MAP estimate of the inflow density with the experimental specification (and the disagreement of the velocity estimate) offers a very weak corroboration of the experimental values.

This is the case for two reasons. First, the MCMC solution relies on approximate surrogate models. Secondly, per the sensitivity analysis in Sec. IV, the measured pressure and heat-flux are far more sensitive to the inflow velocity, and thus the inflow velocity is easier to estimate correctly than the inflow density. Nevertheless, the two approaches - MCMC and the deterministic inversion - do not agree perfectly, and therefore we check the optimal result using a multi-start test.

A common problem with deterministic inversions is that they can often fall into a local minima when starting from an unfortunate guess; this occurs when the objective function is rough. Such a problem does not arise with global methods like MCMC. A local minimum can provide a better agreement with experimental measurements of $p(x)$ and $q(x)$ compared to other (ρ_∞, U_∞) in its vicinity but is not the optimal solution. This problem can be solved by starting multiple inversions from randomly chosen initial guesses of (ρ_∞, U_∞) ; it is hoped that the majority of them will converge to the same global minimum. Consequently, we seed 10 inversions with different guesses of (ρ_∞, U_∞) and have them evolve them to their final converged state. The results are plotted in Fig. 7. We see 10 starting points for the optimization study plotted in green which all converge to the same point (in red). The filled red circle, and the “whiskers” around it, are the inflow conditions provided by the experimental dataset and the measurement errors. We see that all the optimization studies converge to the same minima, indicating that it is the only one in the vicinity. In addition, the converged value is far outside the error bounds provided by the experimental dataset, indicating that the inflow conditions for Case 1 are inconsistent with the measurements of $p(x)$ and $q(x)$. The dotted paths (10 of them) were performed on the “Medium” mesh, because of the computational cost of performing these inversions. The three dashed paths in the figure are “Medium” mesh inversions that were repeated with a “Fine” mesh, as a check. The two sets of paths are identical, reinforcing the conclusion that Case 1’s (experimental) inflow conditions, as stated in Table 1, are potentially not correct.

In Fig. 8 we plot the pressure and heat-flux predictions using the nominal and estimated inflow conditions. The measured values are also plotted. It is clear, and unsurprising, that the estimated inflow conditions accurately predict the pressure and heat-flux at the calibration probes. However, when comparing with all the probes, the nominal inflow conditions provide a better agreement. This is best captured by the root mean square error (RMSE) of the predictions (versus the measured values) provided in Table 3. The predictions at the “calibration” probes show unequivocal improvement when computed with the estimated inflow conditions, while the opposite is true when all the probes are considered (most of the probes lie on the aft-cone where complex shock interactions, flow separation and re-attachment take place). Since the flow over the fore-cone is simple (attached and laminar), this points towards serious model-form error. This is corroborated by Fig. 8 (left) which shows the large discrepancy between predicted and measured pressure, regardless of the inflow conditions used for computing the predictions.

To conclude, both the MCMC and our deterministic inversion suggest that the experimental specification of the inflow velocity may be erroneous, i.e., the flow at the inlet is faster than stated and lies outside the experimentally specified range. The inference of the inflow density is rather unclear, as MCMC and the deterministic inversion provide

opposite suggestions, with our current estimate lying well outside the experimental measurement error bounds. There is significant model-form error and the predictions on the aft-cone show large errors, regardless of the inflow conditions used to compute them.

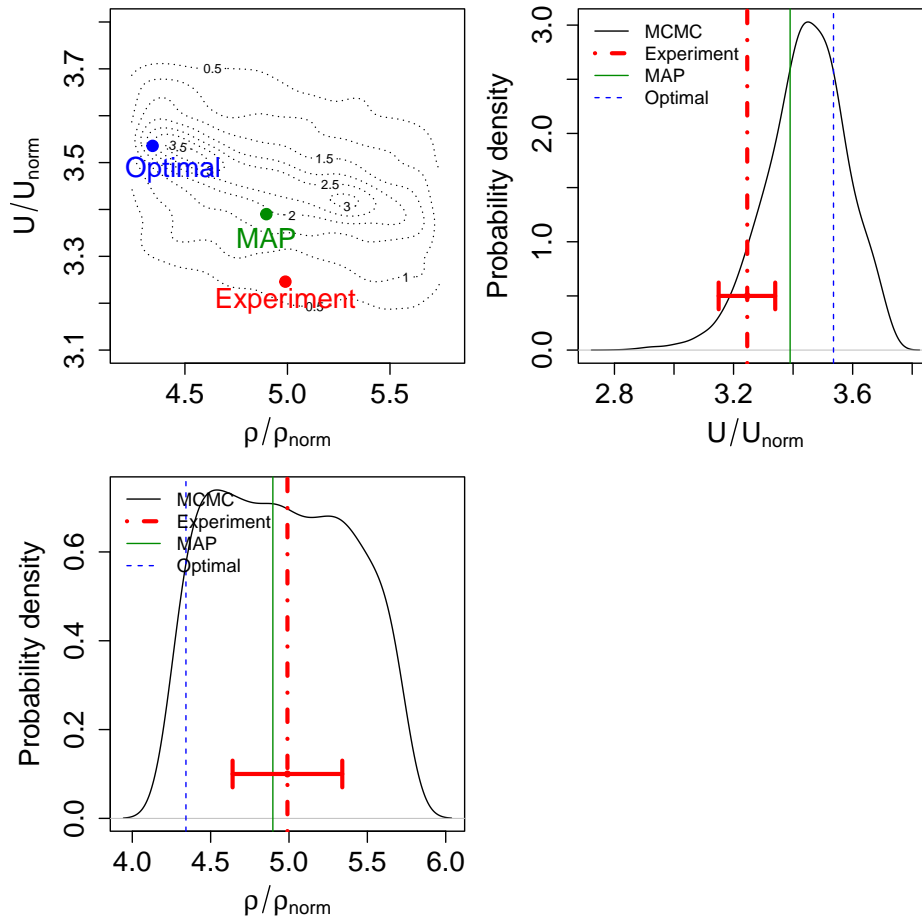


Fig. 6 Posterior PDFs for Case 1. Top left: Contours of the joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ PDF from the MCMC solution, plotted with dotted lines. The MAP estimate, the Optimal estimate from the deterministic inversion and the experimental values are also plotted. Bottom left: Marginal posterior PDFs for ρ_∞/ρ_{norm} , from MCMC. The MAP estimate (“MAP”; solid vertical line), the optimal estimate (“Optimal”; vertical dashed line) from deterministic inversion and the experimental specification (dotted vertical line with measurement error bars) are also plotted. Top right: Same as the bottom left subfigure, but for normalized velocity.

D. Case 4

Finally, we address the high-enthalpy Case 4 (see Table 1 for the flow conditions). The flow is modeled with a 5-species, 17-reaction mechanism [46]. The inflow is assumed to be in equilibrium i.e., $T_\infty = T_{v_\infty}$. The flow has extensive regions of vibrational and chemical non-equilibrium in the vicinity of the double-cone [11]. This experimental dataset has also been studied using Bayesian inference in Ref. [6], which found that the most probable value of both the inflow density and velocity (i.e., the MAP estimate) lay outside the experimental error bounds supplied with the




Fig. 7 Paths through the (ρ_∞, U_∞) -space traced out by the multi-start optimization study. The green circles are the starting points, and the red one the final converged value. Paths plotted with dotted blue lines were inversions performed using a “Medium” mesh (256×512). The ones in dashed red line were repeated with a “Fine” mesh. The filled red circle denotes the “nominal” inflow conditions as provided by the experimentalists (see Table 1). The “whiskers” are the measurement uncertainties provided by the experimental dataset.

dataset. However, since the Bayesian analysis was performed using surrogate models, we repeat it here using exact, if deterministic inversion.

In Fig. 9 (top left) we plot the joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ posterior PDF from the MCMC solution [6] as contours. The MAP estimate (from the MCMC solution), the Optimal one (from the deterministic inversion) and the experimental values are also plotted in the same figure. In Fig. 9 (bottom left) we plot the marginalized posterior PDF for the normalized inflow density, along with the MAP estimate and its deterministic counterpart, and the experimental value with its measurement error bars. Fig. 9 (top right) contains the same information, but for the normalized inflow velocity. We see that for the inflow density, the MAP and the optimal estimates are close, smaller than the experimental specification and far outside the measurement error bounds. Much the same holds true for the inflow velocity, but the estimates are larger than the experimental figures. The results are tabulated in Table 2. *To summarize, neither the*

Table 3 Root-mean-square error (RMSE) for predictions of pressure and heat-flux, using the nominal and estimated inflow conditions. The optimized inflow conditions did not improve predictions at all probes.

Variable	Calibration probes		All probes	
	Nominal	Inferred	Nominal	Inferred
Pressure (N/m^2)	2.6×10^1	1.46×10^1	1.8×10^3	1.9×10^3
Heat flux (W/m^2)	7.6×10^4	1.25×10^4	8.68×10^4	1.5×10^5

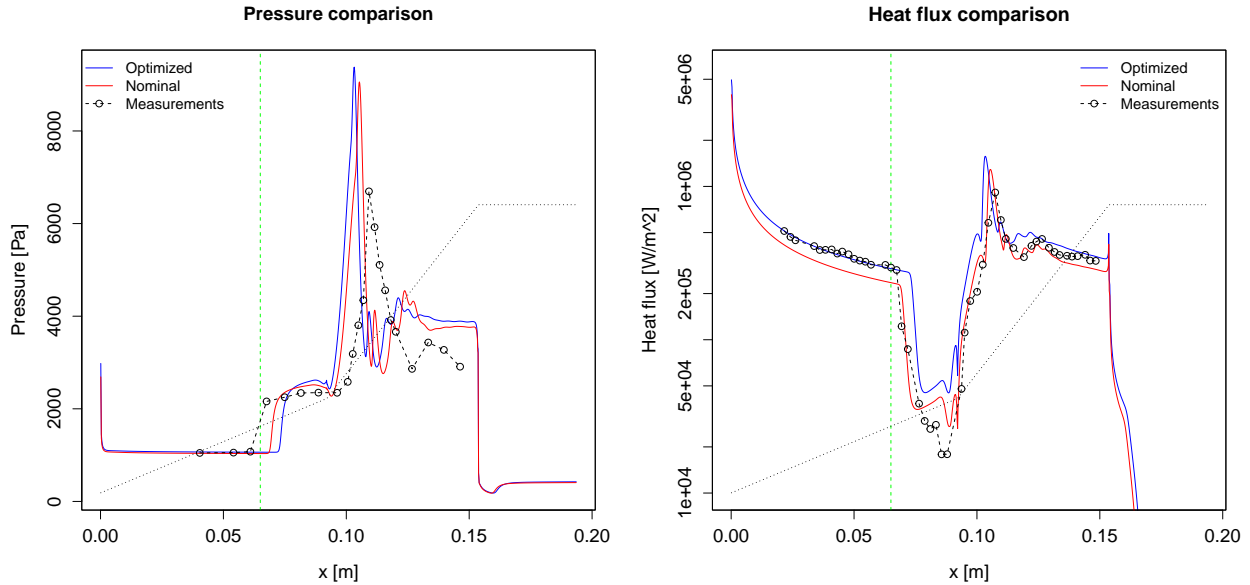


Fig. 8 Predictions of pressure (left) and heat-flux (right) using the optimized (blue) and nominal (red) values of (ρ_∞, U_∞) for Case 1. The measured values at the probes are also plotted. Calibration does not improve the predictive skill of the model for the full flow, though the predictions are far better for the region used to estimate the inflow conditions.

density nor the velocity estimates agree with the experimental data, regardless of whether they are MAP or optimal estimates. This result agrees with the findings in Ref. [6].

E. Discussion

Our investigation into the LENS-XX experiments was motivated by the inability of researchers to reproduce the measurements (see Sec. II and III). In totality, there can be three main sources of prediction errors. The first source is the model itself, specifically the thermochemical models, the transport property models and the models governing the partition of energy into the translational, rotational, electronic and vibrational modes. As reviewed in Sec. III under “Modeling of double-cone experiment”, the examination of alternative models did not improve predictive skill. The second source of errors could be the numerical scheme, e.g., the finite volume discretization, limiters, inviscid reconstruction scheme etc. As reviewed in the same section, these potential sources of errors were investigated and were found to be insignificant. In addition, in the “blind” modeling challenge [22], researchers used different numerical scheme and grids, but to no avail when it came to predictive accuracy, bolstering the contention that the main source of errors was not in the numerics. Grid resolution is unlikely to be a significant source of errors as grid convergence studies are routinely used to drive down discretization errors until they are no longer significant relative to other errors (our own efforts are in Ref.[2]). Monitoring convergence to steady state is automated in most simulators, and it is also unlikely to be a major source of error. The third source of error is the boundary conditions. The wall temperature T_{wall}

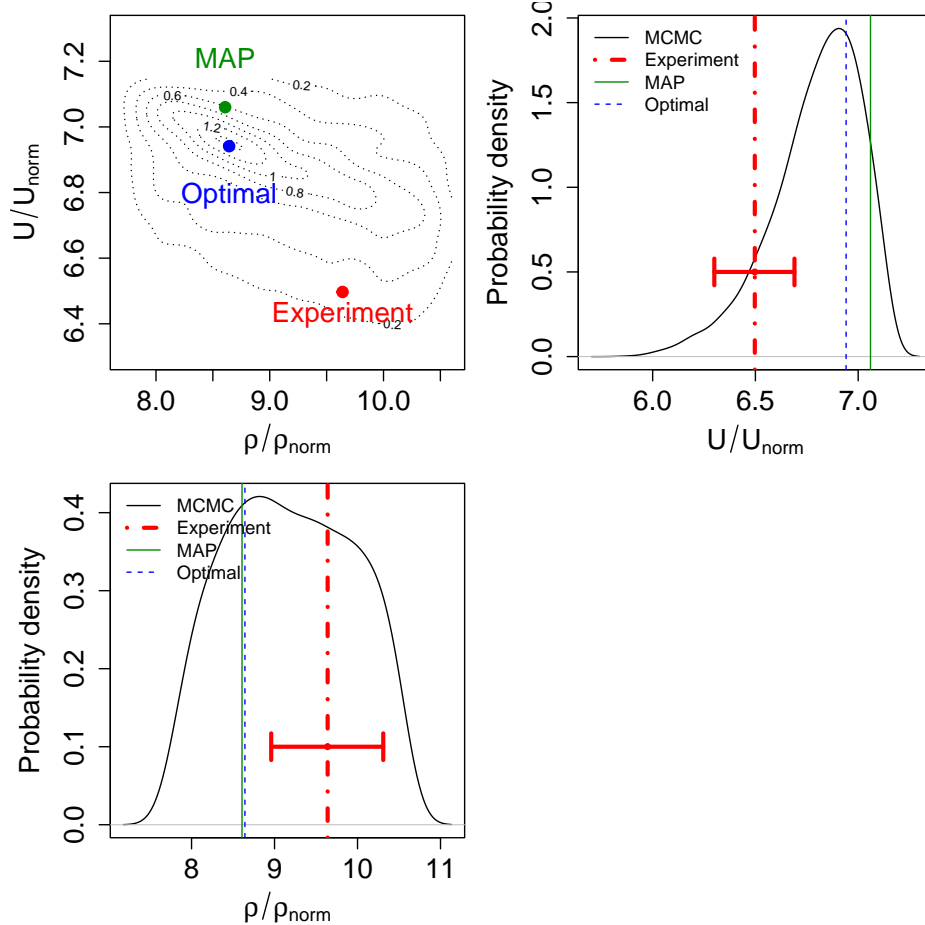


Fig. 9 Posterior PDFs for Case 4. **Top left:** Contours of the joint $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ posterior PDF from the MCMC solution, plotted with dotted lines. The MAP estimate, the Optimal estimate from the deterministic inversion and the experimental value are also plotted. **Bottom left:** Marginal posterior PDFs for ρ_∞/ρ_{norm} , from MCMC. The MAP estimate (“MAP”; solid vertical line), the optimal estimate from deterministic inversion (“Optimal”; vertical dashed line) and the experimental value (dotted vertical line with measurement error bars) are also plotted. **Top right:** Same as the bottom left subfigure, but for the normalized velocity.

of the double-cone is stated by the experimentalists to be constant and its potential variability during an experiment posed no difficulty in previous studies that have reproduced double-cone measurements in the LENS-I shock-tunnel [1]. That leaves the inflow boundary conditions, which are the objects of this investigation.

The results above show that our deterministic inversion method infers inflow conditions that matches the inferences drawn with MCMC in Ref. [6]. For Run 35, which was conducted in the LENS-I shock tunnel and has been modeled successfully, the inferred inflow conditions lie within the experimental error bounds of the inflow specification that accompanied the dataset. For the LENS-XX experiments, which have yet to be modeled with good accuracy, our inversion method leads us to believe that the inflow conditions may have been mis-specified in both experimental datasets. The Bayesian analysis in Ref. [6] reaches much the same conclusion for Case 4, but the opposite one for Case 1. However, this conclusion is dependent on the accuracy of the surrogate models of the Navier-Stokes flow simulator,

which could be the reason for this disagreement. In contrast, our method uses the Navier-Stokes simulator natively and the inflow condition estimates are, therefore, arguably more accurate.

Checking the Gaussian assumption: The usefulness of a Gaussian approximation to the posterior distribution when the uncertainty in a quantity is small is well-known (see Ref. [59], Appendix A) and its applicability can be monitored, without access to the true posterior distribution from MCMC, using the methods in Ref. [60]. Consequently we formulate a test, based on cross-validation arguments, for the appropriateness of a Gaussian assumption for the posterior distribution. The inflow estimates discussed above were drawn from $N_p + N_q + 2$ observations, corresponding to N_p calibration probes for pressure, N_q for the heat-flux, the total enthalpy h_0 and Pitot pressure P_{Pitot} . Consider estimating inflow conditions (ρ^*, U^*) using a subset of the observations (in our case, 75% of the observations), selected randomly. This can be repeated K times, using a different, randomly selected subset of the observations each time, to yield K different sets of inflow estimates, $(\rho^*, U^*)_k, k = 1 \dots K$. The scatter in $(\rho^*, U^*)_k$ reflects the uncertainty in the estimates due to the limited nature of the observations from which they were drawn. In addition, the scatter will be larger (wider) than the true uncertainty distribution e.g., from MCMC, as was done in Ref. [6], as the Bayesian computation would use the full set of $N_p + N_q + 2$ observations. If the Gaussian posterior is far wider than the scatter in $(\rho^*, U^*)_k$, it would be evidence that the Gaussian assumption for the posterior distribution is inappropriate.

Note that for our problem, while we have N_p and N_q calibration probes, they do not carry independent information. The pressure and heat-flux measurements can be reduced to two constants (modulo measurement error) under a self-similar transformation (see Ref. [6]). Thus, we have 4 independent sources of observations, corresponding to pressure, heat-flux, h_0 and P_{Pitot} . We perform a 4-way cross-validation (i.e., $K = 4$) by selecting 3 measurement modalities at a time and computing $(\rho^*, U^*)_k, k = 1 \dots 4$. In Fig. 10, we plot the inflow estimates $(\rho^*, U^*)_k$ as symbols. These are computed for Run 35. The ρ^* estimates are plotted in the left column and the U^* in the right column. The marginalized posterior distributions for ρ_∞ and U_∞ are plotted as box-and-whisker plots, for the MCMC solution from Ref. [6] on top and the Gaussian posterior below. We see that ρ^* and U^* both show scatters that are wider than the IQR (Inter-Quartile Range) for MCMC posterior distribution (top row of Fig. 10). In contrast, the bottom row, which contains the posterior distribution obtained under a Gaussian assumption, shows that ρ_k^* and U_k^* are clustered around the median, and are considerably narrower than the IQR. While four $(\rho^*, U^*)_k$ samples cannot adequately represent a distribution, it is clear that the Gaussian posteriors are far wider than the scatter in $(\rho^*, U^*)_k$, (qualitatively) suggesting that the use of Gaussian posteriors is inappropriate for this problem. Note that the cross-validation test described above does not require one to compute an exact posterior. Given the ease with which the Gaussian posterior can be computed and then checked for the appropriateness of the Gaussian assumption (via cross-validation), it is a convenient tool for quantifying the uncertainty in any estimates obtained using our deterministic inversion method.

Inversions on a coarse mesh: Our inversion approach can be used to check the consistency of an experimental dataset, e.g., to check whether certain measurements are consistent with specifications that accompany a dataset, such as

the boundary conditions. However, these inversions can be computationally expensive, and it is reasonable to wonder if the checks could be done quickly, if approximately, using a coarse mesh. This is feasible, and is described in Ref. [45].

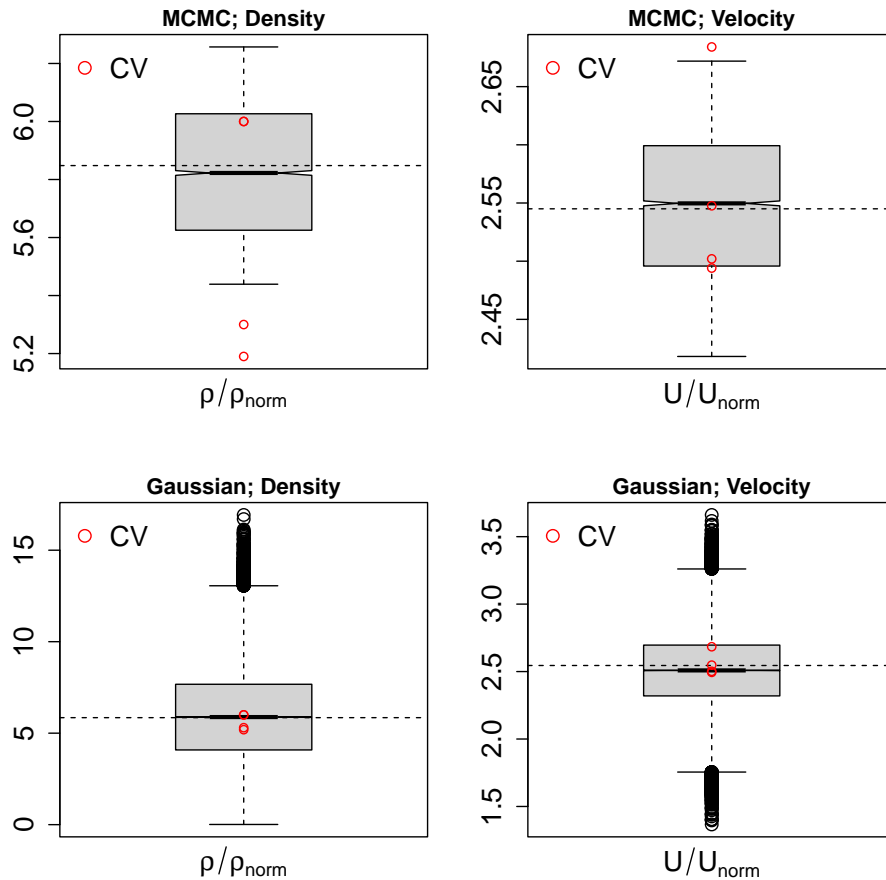


Fig. 10 The “cross-validation” (CV) estimates $(\rho^*, U^*)_k, k = 1 \dots 4$ for Run 35 plotted as symbols, with ρ_∞/ρ_{norm} on the left and U_∞/U_{norm} on the right. The posterior distribution from a Bayesian inference of $(\rho_\infty/\rho_{norm}, U_\infty/U_{norm})$ (from Ref. [6]) are plotted as box-and-whisker plots on the top row. The bottom row has the posterior distributions obtained with a Gaussian assumption. It is clear that the Gaussian posterior is overly wide (compare the vertical axes of the Bayesian and deterministic estimates).

VII. Conclusion

In this paper, we have developed a deterministic inversion method to check the consistency of an experimental dataset. As we begin to validate models against challenging test cases, with experimental data acquired in extreme environments, the possibility of large measurement errors creeping into the model validation dataset becomes uncomfortably high. This is because the experimental apparatus is often operated at the edges of its operational envelope in these experiments, and their behavior under those conditions is not always known. In many cases, it may be possible to partition an experimental dataset into two and infer one from the other. If the inferred quantities lie within the experimental error bounds of the corresponding measurements, the dataset can be deemed consistent. This does require that the model be

an accurate one.

This approach was used in Ref. [6] to analyze the LENS-I (Run 35) and LENS-XX double-cone datasets (Case 1 and Case 4) studied in this paper. They found that Run 35 was consistent, but Cases 1 and 4 were not. However, that study relied on (approximate) surrogate models of the Navier-Stokes flow model that we use. Our studies in Sec. VI, conducted using the deterministic inversion method that uses a Navier-Stokes model natively, largely bears out the older results. If the measurements on the double-cone are correct, then the inlet conditions need to be different, provided the inlet flow is axisymmetric. The measurements used in this study (as well as in Ref. [6]) are collected from the attached laminar-flow region of the fore-cone, where the flow physics are relatively well-known.

The deterministic inference method is formulated as an optimization problem whose objective function is given by the (scaled) sum-of-squares difference between the experimentally measured quantities and those predicted by the CFD simulator. The optimization method is solved using a gradient-based approach, where the gradient of the objective function is computed using adjoint sensitivities facilitated by differentiation of the CFD solver using automatic differentiation. Since the inference method is deterministic, it only directly provides an estimate of the inferred parameters. However, a posterior distribution can be estimated under Gaussian assumptions. We devised a test for the validity of the Gaussian assumption and found that it does not hold for our problem; hence no Gaussian posterior distributions were presented in this paper. They can be found in our conference paper [45].

We find evidence that the inlet conditions for Case 1 could be different from what was stated in the experimental dataset. This conclusion is conditional on the model (Park's, [18]) used for vibrational non-equilibrium, but studies have shown that using a more sophisticated model does not improve predictions materially [14, 16, 19]. Case 4 experiences vibrational and chemical non-equilibrium, and detailed studies have shown that existing thermochemistry models do not reproduce measurements well [11, 14, 15]. The inlet conditions, as inferred by us, are definitely outside the range of stated experimental error. The same result was obtained in Ref. [6]. Note that these conclusions could potentially be dependent on the thermochemical and non-equilibrium models in SPARC; other factors, such as the grid and the numerical method, are not expected to materially affect the conclusions.

This raises the question of what could be the cause of the discrepancy between model predictions and measurements. All CFD studies to date have assumed that the inlet flow is axisymmetric and steady, but the experimental dataset available to us does not allow us to test an alternative hypothesis. More research is definitely warranted, as is more experimental data. Repeated runs, that would allow us to assess the run-to-run variability and measurements on different sides of the double-cone, that would allow us to test the assumption of axisymmetric flow, would most likely top a modeler's wish-list. One interesting study [61] looked at the effect of non-equilibrium flow *at the inlet* and found that it could increase the size of the separation bubble, without materially changing the heat-flux and pressure profiles on the fore-cone. However, the study was performed for an O_2 environment in the LENS-I *shock* tunnel; in contrast, the LENS-XX *expansion* tunnel is designed to suppress non-equilibrium processes upstream of the test-section inlet. In

addition, knowledge of the vibrational and thermochemical non-equilibrium state of the inflow at the test section would be helpful for modeling inflow uncertainties; direct measurements of $T_{v,\infty}$ in a tunnel are feasible [62, 63]. Finally, in light of the findings from the X3 experiments [12], it would be of utmost help if one could ascertain whether assumptions of axisymmetry and steady flow, which underlie all modeling work to date, are warranted for the LENS-XX experiments. Note that assumptions of steady, axisymmetric flow were sufficient for reproducing Run 35 conducted in LENS-I [1].

Acknowledgments

This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://www.energy.gov/downloads/doe-public-access-plan>).

References

- [1] Nompelis, I., Candler, G. V., and Holden, M. S., “Effect of vibrational nonequilibrium on hypersonic double-cone experiments,” *AIAA Journal*, Vol. 41, No. 11, 2003, pp. 2162–2169. <https://doi.org/10.2514/2.6834>.
- [2] Kieweg, S., Ray, J., Weirs, V. G., Carnes, B., Dinzl, D., Freno, B., Howard, M., Rider, W. J., and Smith, T., “Validation Assessment of Hypersonic Double-Cone Flow Simulations using Uncertainty Quantification, Sensitivity Analysis, and Validation Metrics,” *AIAA 2019 Aerosciences Conference, AIAA SciTech 2019 (AIAA-2019-2278)*, 2019. <https://doi.org/10.2514/6.2019-2278>.
- [3] MacLean, M., Dufrene, A., Wadhams, T., and Holden, M., “Numerical and experimental characterization of high enthalpy flow in an expansion tunnel facility,” *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, Aerospace Sciences Meetings, (AIAA-2010-1562)*, 2010. <https://doi.org/10.2514/6.2010-1562>.
- [4] Dufrene, A., MacLean, M., Parker, R., Wadhams, T., and Holden, M., “Characterization of the new LENS expansion tunnel facility,” *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition (AIAA-2010-1564)*, 2010. <https://doi.org/10.2514/6.2010-1564>.
- [5] Nompelis, I., Candler, G., Holden, M., and Wadhams, T., “Numerical simulation of high-enthalpy experiments in the LENS X expansion tube facility,” *42nd AIAA Aerospace Sciences Meeting and Exhibit, Aerospace Sciences Meetings, (AIAA-2004-1000)*, 2004. <https://doi.org/10.2514/6.2004-1000>.
- [6] Ray, J., Kieweg, S., Dinzl, D., Carnes, B., Weirs, V., Freno, B., Howard, M., Smith, T., Nompelis, I., and Candler, G. V., “Estimation of Inflow Uncertainties in Laminar Hypersonic Double-Cone Experiments,” *AIAA Journal*, Vol. 58, No. 10, 2020, pp. 4461–4474. <https://doi.org/10.2514/1.J051609>.

- [7] Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. (eds.), *Markov Chain Monte Carlo in Practice*, Chapman & Hall, 1996.
- [8] Holden, M., “Experimental studies of laminar separated flows induced by shock wave/boundary layer and shock/shock interaction in hypersonic flows for CFD validation,” *38th Aerospace Sciences Meeting and Exhibit, Aerospace Sciences Meetings, (AIAA-2000-0930)*, 2000. <https://doi.org/10.2514/6.2000-930>.
- [9] Harvey, J., Holden, M., and Wadhams, T., “Code validation study laminar shock/boundary layer and shock/shock interactions in hypersonic flow Part B: Comparison with Navier–Stokes and DSMC solutions,” *39th Aerospace Sciences Meeting and Exhibit, Reno, NV, U.S.A (AIAA-2001-1031)*, 2001. <https://doi.org/10.2514/6.2001-1031>.
- [10] MacLean, M., Holden, M. S., and Dufrene, A., “Measurements of real gas effects on regions of laminar shock wave/boundary layer interaction in hypervelocity flows,” Tech. rep., CUBRC Inc., 2014. https://www.cubrc.org/_iassets/docs/laminar-xx-paper.pdf.
- [11] Youssefi, M. R., and Knight, D., “Assessment of CFD capability for hypersonic shock wave laminar boundary layer interactions,” *Aerospace*, Vol. 4, No. 25, 2017. <https://doi.org/10.3390/aerospace4020025>.
- [12] Kennedy, A., Geraets, R. T. P., James, C. M., Thompson, M., Morgan, R. G., Austin, J. M., Zander, F., and McGilvray, M., *Measurements of Hypersonic Double Cone Flows with Shock Wave/Boundary Layer Interactions in the X3 Expansion Tunnel*, ????. <https://doi.org/10.2514/6.2023-0271>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2023-0271>.
- [13] Holden, M., Wadhams, T., Harvey, J., and Chandler, G., “Comparisons between DSMC and Navier-Stokes solutions and measurements in regions of laminar shock wave boundary layer interactions in hypersonic flows,” *40th AIAA Aerospace Sciences Meeting & Exhibit*, 2002, p. 435. AIAA 2002-0435.
- [14] Holloway, M. E., Chaudhry, R. S., and Boyd, I. D., “Assessment of Hypersonic Double-Cone Experiments for Validation of Thermochemistry Models,” *Journal of Spacecraft and Rockets*, Vol. 59, No. 2, 2022, pp. 389–400. <https://doi.org/10.2514/1.A35052>.
- [15] Kianvashrad, N., and Knight, D. D., “Nonequilibrium Effects on Prediction of Aerothermodynamic Loading for a Double Cone,” *AIAA Journal*, Vol. 57, No. 7, 2019, pp. 2946–2963. <https://doi.org/10.2514/1.J057883>.
- [16] Hao, J., Wang, J., and Lee, C., “Numerical simulation of high-enthalpy double-cone flows,” *American Institute of Aeronautics and Astronautics Journal*, Vol. 55, No. 6, 2017, pp. 2471–2475. <https://doi.org/10.2514/1.J055746>.
- [17] Marrone, P. V., and Treanor, C. E., “Chemical relaxation with preferential dissociation from excited vibrational levels,” *Physics of Fluids*, Vol. 6, No. 9, 1963, pp. 1215–1221. <https://doi.org/10.1063/1.1706888>.
- [18] Park, C., “On convergence of computation of chemically reacting flows,” *23rd AIAA Aerospace Sciences Meeting (AIAA-1985-0247)*, 1985. <https://doi.org/10.2514/6.1985-247>.
- [19] Gao, J., Hao, J., Wang, J., and Lee, C., “Effect of Thermochemical Nonequilibrium Modeling on High-Enthalpy Double-Cone Flow,” *Journal of Spacecraft and Rockets*, Vol. 58, No. 4, 2021, pp. 1243–1247. <https://doi.org/10.2514/1.A35002>.

- [20] Zuo, F.-Y., and Hu, S.-L., “Thermochemical non-equilibrium effects on aerothermodynamic prediction of laminar double-cone flow,” *Acta Astronautica*, Vol. 182, 2021, pp. 179–188. <https://doi.org/10.1016/j.actaastro.2021.01.058>, URL <https://www.sciencedirect.com/science/article/pii/S0094576521000692>.
- [21] Nompelis, I., and Candler, G. V., “US3D predictions of double-cone and hollow cylinder and cylinder-flare at high enthalpy,” *44th AIAA Fluid Dynamics Conference (AIAA-2014-3366)*, 2014. <https://doi.org/10.2514/6.2014-3366>.
- [22] Candler, G. V., “Next-generation CFD for hypersonic and aerothermal flows (Invited),” *22nd AIAA Computational Fluid Dynamics Conference, AIAA AVIATION Forum, (AIAA 2015-3048)*, 2015. <https://doi.org/10.2514/6.2015-3048>.
- [23] Carnes, B., Weirs, V. G., and Smith, T., “Code verification and numerical error estimation with application to model validation of laminar, hypersonic flow over a double cone,” *AIAA 2019 Aerosciences Conference, AIAA SciTech 2019 (AIAA-2019-2175)*, 2019. <https://doi.org/10.2514/6.2019-2175>.
- [24] Freno, B. A., Carnes, B. R., and Weirs, V. G., “Code-verification techniques for hypersonic reacting flows in thermochemical nonequilibrium,” *Journal of Computational Physics*, Vol. 425, 2021, p. 109752. <https://doi.org/10.1016/j.jcp.2020.109752>.
- [25] Harvey, J., Holden, M., and Candler, G., “Validation of DSMC/Navier-Stokes Computations for Laminar Shock Wave/Boundary Layer Interactions Part 3,” *36th AIAA Thermophysics Conference*, 2003. <https://doi.org/10.2514/6.2003-3643>.
- [26] Petra, N., Martin, J., Stadler, G., and Ghattas, O., “A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems, Part II: Stochastic Newton MCMC with Application to Ice Sheet Flow Inverse Problems,” *SIAM Journal on Scientific Computing*, Vol. 36, No. 4, 2014, pp. A1525–A1555. <https://doi.org/10.1137/130934805>, URL <https://doi.org/10.1137/130934805>.
- [27] Bui-Thanh, T., Ghattas, O., Martin, J., and Stadler, G., “A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems Part I: The Linearized Case, with Application to Global Seismic Inversion,” *SIAM Journal on Scientific Computing*, Vol. 35, No. 6, 2013, pp. A2494–A2523. <https://doi.org/10.1137/12089586X>, URL <https://doi.org/10.1137/12089586X>.
- [28] Iollo, A., Ferlauto, M., and Zannetti, L., “An Aerodynamic Optimization Method Based on the Inverse Problem Adjoint Equations,” *Journal of Computational Physics*, Vol. 173, No. 1, 2001, pp. 87–115. <https://doi.org/https://doi.org/10.1006/jcph.2001.6845>, URL <https://www.sciencedirect.com/science/article/pii/S0021999101968457>.
- [29] Nemec, M., Aftosmis, M., Murman, S., and Pulliam, T., “Adjoint formulation for an embedded-boundary Cartesian method,” *43rd AIAA Aerospace Sciences Meeting and Exhibit*, 2005, p. 877. <https://doi.org/10.2514/6.2005-877>.
- [30] Nemec, M., and Aftosmis, M., “Aerodynamic shape optimization using a Cartesian adjoint method and CAD geometry,” *24th AIAA Applied Aerodynamics Conference*, 2006, p. 3456. <https://doi.org/10.2514/6.2006-3456>.
- [31] Nemec, M., Aftosmis, M. J., and Brown, P. G., “Numerical prediction of meteoric infrasound signatures,” *Planetary and Space Science*, Vol. 140, 2017, pp. 11–20. <https://doi.org/10.1016/j.pss.2017.03.003>.

- [32] Nielsen, E. J., Lu, J., Park, M. A., and Darmofal, D. L., “An implicit, exact dual adjoint solution method for turbulent flows on unstructured grids,” *Computers & Fluids*, Vol. 33, No. 9, 2004, pp. 1131–1155. <https://doi.org/10.1016/j.compfluid.2003.09.005>.
- [33] Nielsen, E. J., and Kleb, W. L., “Efficient construction of discrete adjoint operators on unstructured grids using complex variables,” *AIAA journal*, Vol. 44, No. 4, 2006, pp. 827–836. <https://doi.org/10.2514/6.2005-324>.
- [34] Thompson, K. B., “Aerothermodynamic Design Sensitivities for a Reacting Gas Flow Solver on an Unstructured Mesh Using a Discrete Adjoint Formulation,” Ph.D. thesis, North Carolina State University, 2017.
- [35] Economon, T. D., Palacios, F., Copeland, S. R., Lukaczyk, T. W., and Alonso, J. J., “SU2: An Open-Source Suite for Multiphysics Simulation and Design,” *AIAA Journal*, Vol. 54, No. 3, 2016, pp. 828–846. <https://doi.org/10.2514/1.J053813>, URL <https://doi.org/10.2514/1.J053813>.
- [36] Copeland, S., Lonkar, A., Palacios, F., and Alonso, J., “Adjoint-Based Goal-Oriented Mesh Adaptation for Nonequilibrium Hypersonic Flows,” *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2013. <https://doi.org/10.2514/6.2013-552>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2013-552>.
- [37] Copeland, S. R., “A Continuous Adjoint Formulation for Hypersonic Flows in Thermochemical Nonequilibrium,” Ph.D. thesis, Stanford University, 2015.
- [38] Kline, H., and Alonso, J., “Adjoint of generalized outflow-based functionals applied to hypersonic inlet design,” *AIAA Journal*, Vol. 55, No. 11, 2017, pp. 3903–3915. <https://doi.org/10.2514/1.J055863>.
- [39] Albring, T. A., Sagebaum, M., and Gauger, N. R., “Efficient aerodynamic design using the discrete adjoint method in SU2,” *17th AIAA/ISSMO multidisciplinary analysis and optimization conference*, 2016, p. 3518. <https://doi.org/10.2514/6.2016-3518>.
- [40] Munguía, B. C., Maier, W. T., Needels, J. T., and Alonso, J. J., “Goal-Oriented Mesh Adaptation for Flows in Thermochemical Nonequilibrium,” *23rd AIAA International Space Planes and Hypersonic Systems and Technologies Conference*, 2020. <https://doi.org/10.2514/6.2020-2400>, URL <https://doi.org/10.2514/6.2020-2400>.
- [41] Damm, K. A., Gollan, R. J., Jacobs, P. A., Smart, M. K., Lee, S., Kim, E., and Kim, C., “Discrete adjoint optimization of a hypersonic inlet,” *AIAA Journal*, Vol. 58, No. 6, 2020, pp. 2621–2634. <https://doi.org/10.2514/1.J058913>.
- [42] Lockwood, B., and Mavriplis, D., “Parameter Sensitivity Analysis for Hypersonic Viscous Flow Using a Discrete Adjoint Approach,” *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2010. <https://doi.org/10.2514/6.2010-447>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2010-447>.
- [43] Lockwood, B., Rumpfkeil, M., Yamazaki, W., and Mavriplis, D., “Uncertainty quantification in viscous hypersonic flows using gradient information and surrogate modeling,” *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2010, p. 885. <https://doi.org/g/10.2514/6.2011-885>.

- [44] Lockwood, B., Anitescu, M., and Mavriplis, D., “Mixed aleatory/epistemic uncertainty quantification for hypersonic flows via gradient-based optimization and surrogate models,” *50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2012, p. 1254. <https://doi.org/10.2514/6.2012-1254>.
- [45] Ray, J., Blonigan, P., Phipps, E. T., and Maupin, K., “An Assessment of the Laminar Hypersonic Double-Cone Experiments in the LENS-XX Tunnel,” *AIAA 2023 Aerosciences Conference, AIAA SciTech 2023 (AIAA-2023-0856)*, 2023. <https://doi.org/10.2514/6.2023-0856>.
- [46] Park, C., “Review of chemical-kinetic problems of future NASA missions I: Earth entries,” *Journal of Thermophysics and Heat Transfer*, Vol. 7, No. 3, 1990, pp. 385–398. <https://doi.org/10.2514/3.431>.
- [47] Schlichting, H., and Gersten, K., *Boundary Layer Theory*, 8th ed., Springer, 2000.
- [48] Nocedal, J., and Wright, S., *Numerical Optimization*, Springer, 2006.
- [49] Kouri, D., Ridzal, D., and von Wincke, G., *The ROL Project Website*, 2020 (accessed May 22, 2020). URL <https://trilinos.github.io/rol.html>.
- [50] Phipps, E. T., and Gay, D. M., “Sacado Automatic Differentiation Package,” <http://trilinos.org/packages/sacado/>, 2018.
- [51] Phipps, E. T., Bartlett, R. A., Gay, D. M., and Hoekstra, R. J., “Large-Scale Transient Sensitivity Analysis of a Radiation-Damaged Bipolar Junction Transistor via Automatic Differentiation,” *Advances in Automatic Differentiation*, edited by C. H. Bischof, H. M. Bücker, P. D. Hovland, U. Naumann, and J. Utke, Springer, 2008, pp. 351–362. https://doi.org/10.1007/978-3-540-68942-3_31.
- [52] Phipps, E. T., and Pawlowski, R. P., “Efficient Expression Templates for Operator Overloading-based Automatic Differentiation,” *Recent Advances in Algorithmic Differentiation*, edited by S. Forth, P. Hovland, E. Phipps, J. Utke, and A. Walther, Springer, 2012.
- [53] Gebremedhin, A. H., Manne, F., and Pothen, A., “What Color Is Your Jacobian? Graph Coloring for Computing Derivatives,” *SIAM Review*, Vol. 47, No. 4, 2005, pp. 629–705. <https://doi.org/10.1137/S0036144504444711>.
- [54] Bozdag, D., Gebremedhin, A. H., Manne, F., Boman, E. G., and Catalyurek, U. V., “A framework for scalable greedy coloring on distributed-memory parallel computers,” *Journal of Parallel and Distributed Computing*, Vol. 68, No. 4, 2008, pp. 515–535. <https://doi.org/10.1016/j.jpdc.2007.08.002>.
- [55] Bozdag, D., Catalyurek, U. V., Gebremedhin, A. H., Manne, F., Boman, E. G., and Ozguner, F., “Distributed-Memory Parallel Algorithms for Distance-2 Coloring and Related Problems in Derivative Computation,” *SIAM Journal on Scientific Computing*, Vol. 32, No. 4, 2010, pp. 2418–2446. <https://doi.org/10.1137/080732158>.
- [56] Boman, E. G., Çatalyürek, Ü. V., Chevalier, C., and Devine, K. D., “The Zoltan and Isorropia Parallel Toolkits for Combinatorial Scientific Computing: Partitioning, Ordering and Coloring,” *Scientific Programming*, Vol. 20, 2012, p. 713587. <https://doi.org/10.3233/SPR-2012-0342>.

- [57] Saad, Y., and Schultz, M. H., "GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems," *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1986, pp. 856–869. <https://doi.org/10.1137/0907058>.
- [58] Wright, M. J., Candler, G. V., and Bose, D., "Data-Parallel Line Relaxation Method for the Navier-Stokes Equations," *AIAA Journal*, Vol. 36, No. 9, 1998, pp. 1603–1609. <https://doi.org/10.2514/2.586>.
- [59] Yuen, K.-V., *Bayesian Methods for Structural Dynamics and Civil Engineering*, John Wiley and Sons (Asia) Pte Ltd, 2010.
- [60] Dovì, V. G., Paladino, O., and Reverberi, A. P., "Some remarks on the use of the inverse Hessian matrix of the likelihood function in the estimation of statistical properties of parameters," *Applied Mathematics Letters*, Vol. 4, No. 1, 1991, pp. 87–90. [https://doi.org/10.1016/0893-9659\(91\)90129-J](https://doi.org/10.1016/0893-9659(91)90129-J).
- [61] Ninni, D., Bonelli, F., Colonna, G., and Pascazio, G., "On the influence of non equilibrium in the free stream conditions of high enthalpy oxygen flows around a double-cone," *Acta Astronautica*, Vol. 201, 2022, pp. 247–258. <https://doi.org/10.1016/j.actaastro.2022.09.017>, URL <https://www.sciencedirect.com/science/article/pii/S0094576522004829>.
- [62] Daniel, K., Jans, E., Lynch, K. P., Swain, W., Downing, C., and Wagner, J. L., *Nitric Oxide Emission Spectroscopy of a Cylinder Stagnation Flow in a Shock Tunnel*, ??? <https://doi.org/10.2514/6.2023-1174>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2023-1174>.
- [63] Jans, E., Lynch, K. P., Daniel, K., Downing, C., Wagner, J. L., Wagnild, R., Maeng, J. B., and Kearney, S. P., *Characterization of Shock Tunnel Free-Stream Nonequilibrium using Nanosecond Pulse-Burst Coherent Anti-Stokes Raman Scattering*, ??? <https://doi.org/10.2514/6.2023-1929>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2023-1929>.