

RESEARCH ARTICLE | JANUARY 20 2026

Petrov–Galerkin model reduction for collisional–radiative argon plasma

Ivan Zanardi ; Alessandro Meini ; Alberto Padovan ; Daniel J. Bodony ; Marco Panesi  



Phys. Plasmas 33, 013905 (2026)

<https://doi.org/10.1063/5.0306085>

 CHORUS



View
Online



Export
Citation

Articles You May Be Interested In

The arbitrary Lagrangian–Eulerian coupled finite element-discontinuous Galerkin method for the non-Newtonian two phase fluid structure interaction via the phase field model

Physics of Fluids (November 2025)

Finite volume method for the Black-Scholes equation transformed on finite interval

AIP Conf. Proc. (November 2012)

Generalized similarity solutions for the type D fluid in five-dimensional flat space

J. Math. Phys. (June 1998)

21 February 2026 02:44:35

AIP Advances

Why Publish With Us?



21DAYS
average time
to 1st decision



OVER 4 MILLION
views in the last year



INCLUSIVE
scope

[Learn More](#)

Petrov–Galerkin model reduction for collisional–radiative argon plasma

Cite as: Phys. Plasmas **33**, 013905 (2026); doi: [10.1063/5.0306085](https://doi.org/10.1063/5.0306085)

Submitted: 8 October 2025 · Accepted: 31 December 2025 ·

Published Online: 20 January 2026



View Online



Export Citation



CrossMark

Ivan Zanardi,¹  Alessandro Meini,²  Alberto Padovan,³  Daniel J. Bodony,¹  and Marco Panesi^{2,a)} 

AFFILIATIONS

¹Department of Aerospace Engineering, University of Illinois Urbana-Champaign, Urbana, Illinois 61801, USA

²Department of Mechanical and Aerospace Engineering, University of California Irvine, Irvine, California 92697, USA

³Department of Mechanical and Industrial Engineering, New Jersey Institute of Technology, Newark, New Jersey 07102, USA

^{a)} Author to whom correspondence should be addressed: mpanesi@uci.edu

ABSTRACT

Accurate simulation of nonequilibrium plasmas, essential in hypersonic reentry, fusion energy, and astrophysical flows, relies on state-specific collisional–radiative (CR) kinetic models but often comes with prohibitive computational cost. While traditional approaches reduce this burden through empirical or physics-based simplifications, they frequently compromise accuracy in strongly nonequilibrium regimes. To address these limitations, we develop a Petrov–Galerkin reduced-order model (ROM) for CR argon plasma based on oblique projections that optimally balance the covariance of full-order state trajectories with that of the system’s output sensitivities. This construction ensures that the ROM captures both the dominant energetic modes and the directions most relevant to input–output behavior. After offline training in a zero-dimensional setting using nonlinear forward and adjoint simulations, the ROM is coupled to a finite-volume solver and applied to one- (1D) and two-dimensional (2D) ionizing shock-tube problems. The ROM achieves a $3\times$ reduction in state dimension and more than one order of magnitude savings in floating-point operations, while maintaining errors below 1% for macroscopic quantities. In both 1D and 2D, it robustly reproduces complex unsteady plasma flow features, including periodic fluctuations, electron avalanches, triple points, and cellular ionization patterns. This stands in contrast to standard ROM strategies, which become unstable or inaccurate under these challenging conditions. These results demonstrate the potential of the proposed model reduction strategy to enable high-fidelity simulation of reactive plasma flows at reduced computational cost, offering new capabilities for exploring high-speed fluid systems governed by coupled transport, wave propagation, and kinetic nonequilibrium.

© 2026 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0306085>

I. INTRODUCTION

Nonequilibrium plasmas play a crucial role in a wide range of scientific and engineering applications, including electric propulsion, plasma-assisted combustion, fusion energy, atmospheric reentry, and astrophysical flows.^{1–11} In such regimes, the internal energy distributions of the constituent species, arising from electronic, vibrational, and rotational excitations, can deviate significantly from equilibrium, requiring detailed, state-specific collisional–radiative (CR) kinetic models to faithfully capture the underlying physics.^{12–19} These models, typically formulated as state-to-state (StS) master equations,^{20–29} track the population dynamics of individual internal energy levels through an extensive set of collisional and radiative processes. While offering an unprecedented level of physical accuracy in representing the thermochemical state of the plasma, the associated computational cost is prohibitive, particularly in multidimensional simulations, due to the

large number of degrees of freedom (i.e., the molecules’ and atoms’ energy levels) and the extreme stiffness associated with the governing equations.

Historically, a number of model reduction strategies have been proposed to alleviate this cost. Multi-temperature (MT) models,^{3,30–37} for instance, simplify the problem by assuming thermal equilibrium distributions within each energy mode, introducing additional temperatures to describe vibrational, rotational, and electronic state populations. While computationally attractive, MT models rely on crude, semi-empirical assumptions and are known to break down in strongly nonequilibrium regimes, such as those found in ionizing shock layers or rapid relaxation zones.^{28,29,38} Additional macroscopic reduction strategies have also been proposed, including approaches based on global rate coefficients constructed from detailed StS kinetics. Examples include the formulations of Colonna *et al.*,^{39,40} which

demonstrate how vibrational nonequilibrium can be reduced to macroscopic models through appropriately defined effective rates. Alternatively, coarse-grained (CG) models^{18,41–44} provide a more effective solution compared to the traditional MT models by clustering energy levels into bins and reconstructing level populations using entropy-based closures. These approaches, while grounded in physical principles, require careful tuning and may not capture complex features of the internal state distribution function without significant prior analysis and optimization.^{45–51}

More recently, the emergence of data-driven model reduction techniques has opened new avenues for compressing detailed kinetic systems without relying on physics-based closure assumptions. A simplistic approach involves orthogonally projecting the governing equations onto the span of proper orthogonal decomposition (POD) modes derived from high-fidelity simulation data.^{52–59} Bellemans *et al.*^{60,61} demonstrated the application of such orthogonal projections to reduce the dimensionality of CR models for argon plasmas, achieving good agreement with full-order results in steady 1D shock-tube configurations. However, POD constructs its basis purely from snapshot data, without accounting for sensitivity information encoded in the governing equations. As a result, the method often fails to capture dynamically important directions, particularly in systems exhibiting transient growth and high sensitivity to low-energy disturbances (e.g., low-population states in the distribution function).^{62–65} Consequently, when deployed in predictive simulations, especially under highly nonlinear, unsteady, or extrapolative conditions, POD-based ROMs frequently suffer from instability, degraded accuracy, and poor generalization.

To overcome the limitations of these approaches, Zanardi *et al.*^{66,67} proposed a reduced-order modeling framework based on Petrov–Galerkin projection, in which the full-order dynamics (i.e., the StS master equations) are *obliquely* projected onto a low-dimensional linear subspace. This method builds upon the “Covariance Balancing Reduction using Adjoint Snapshots” (CoBRAS) framework,⁶⁴ which constructs reduced subspaces by balancing the covariance of the full-order state trajectory with that of the system’s output sensitivities (i.e., gradients), similar to the balance truncation technique and its variants.^{68–72} This balance is defined in a statistical mean square sense, ensuring that the reduced space captures both the dominant energetic modes and the directions most influential to the system’s input–output behavior. Therefore, CoBRAS introduces a physically grounded alternative to POD by generating oblique projection operators specifically tailored to minimize output reconstruction error and preserve the system’s most relevant dynamical features. Zanardi *et al.*⁶⁶ demonstrated that CoBRAS-based Petrov–Galerkin ROMs yield stable and accurate reduced models even in highly stiff and nonlinear kinetic systems. A key innovation in their framework is the use of linearized forward and adjoint equations to efficiently approximate the state and gradient covariances, enabling fast offline construction of the reduced basis while preserving fidelity.

In this work, we extend the CoBRAS-based reduced-order modeling framework, previously demonstrated only in zero-dimensional (0D) settings, to simulate multidimensional, unsteady, CR plasma flows, with a focus on ionizing shock dynamics. These flows are particularly demanding due to their strong nonlinearity, multiscale nature,

and tight coupling between thermochemistry and hydrodynamics, conditions under which conventional ROMs often fail due to instability or loss of accuracy. As a first step, we construct a CoBRAS-based Petrov–Galerkin ROM for detailed CR argon kinetics by computing projection spaces from nonlinear 0D forward and adjoint simulations of the full-order system. Second, we integrate the resulting ROM into a finite-volume flow solver and apply it to one- (1D) and two-dimensional (2D) shock-tube configurations characterized by strongly unsteady plasma behavior. Third, we systematically evaluate the performance of the proposed ROM against both the full-order model (FOM) and a standard POD-based ROM, demonstrating substantial improvements in numerical stability, accuracy, and physical fidelity. The CoBRAS approach proves robust even in regimes where traditional ROMs typically fail. This marks a significant step forward in model reduction for reactive plasma flows, opening new avenues for both fundamental investigation and applied analysis of high-speed, nonequilibrium flows in aerospace, astrophysics, and energy science.

The remainder of this paper is organized as follows: Secs. II and III describe the physical modeling of CR argon plasma and the numerical framework used to solve multidimensional plasma flows. Section IV presents the reduced-order modeling methodology and its integration with a finite-volume solver. Section V reports on numerical results obtained in 0D, 1D, and 2D configurations, comparing CoBRAS performance against FOM and POD baselines. Finally, Sec. VI summarizes the key findings and outlines directions for future work.

II. PHYSICAL MODELING

This section outlines the thermochemical and radiative modeling framework adopted in the present study, along with the governing equations describing the plasma flow. The formulation incorporates detailed kinetics and energy exchange processes relevant to weakly ionized argon. For a comprehensive description of the reaction rate coefficients, collisional–radiative source terms, and thermodynamic closures, the reader is referred to the methodology presented by Kapper and Cambier.^{14,15}

A. Collisional–radiative model

To model the kinetics of weakly ionized argon plasmas, we employ a detailed CR framework in which each electronic state is treated as a distinct pseudo-species. This enables a non-Boltzmann description of the atomic state distribution function (ASDF) and offers improved fidelity over traditional MT models. The CR mechanism used in this work was originally developed by Vlček¹² and Bultel¹³ and has been validated against experimental data from the UTIAS (University of Toronto, Institute for Aerospace Studies) shock tube⁷³ by Kapper and Cambier.^{14,15}

The plasma composition includes free electrons (e^-), neutral argon (Ar), and singly ionized argon (Ar^+), while higher-order ionization states (e.g., Ar^{++} and Ar_2^+) are neglected due to their negligible contributions under the conditions considered. The argon species are resolved into their respective electronic states (31 for Ar and 2 for Ar^+) yielding a total of 34 pseudo-species.

The CR model accounts for five types of fundamental processes:^{12–15}

- Excitation and ionization by electron impact:

$$\text{Ar}(i) + e^- \rightleftharpoons \text{Ar}(j) + e^-, \quad (1)$$

$$\text{Ar}(i) + e^- \rightleftharpoons \text{Ar}^+ + e^- + e^-. \quad (2)$$

- Excitation and ionization by heavy-particle impact:

$$\text{Ar}(i) + \text{Ar}(1) \rightleftharpoons \text{Ar}(j) + \text{Ar}(1), \quad (3)$$

$$\text{Ar}(i) + \text{Ar}(1) \rightleftharpoons \text{Ar}^+ + e^- + \text{Ar}(1). \quad (4)$$

- Spontaneous emission/absorption (bound-bound):

$$\text{Ar}(i) + h\nu_{ij} \rightleftharpoons \text{Ar}(j). \quad (5)$$

- Photo-ionization/radiative recombination (free-bound and bound-free):

$$\text{Ar}(i) + h\nu_i \rightleftharpoons \text{Ar}^+ + e^-. \quad (6)$$

- Bremsstrahlung radiation (free-free):

$$e^- + \text{Ar}^+ \rightleftharpoons e^- + \text{Ar}^+ + h\nu. \quad (7)$$

Here, indices $i \in \{1, \dots, 31\}$ and $j > i$ refer to the electronic levels of Ar; the two levels of Ar^+ are omitted for brevity, but their governing equations follow the same formulation as for Ar. Heavy-particle impact collisions between two excited argon states are neglected in the present model, following the formulation of Kapper and Cambier.¹⁴ Under the shock-driven conditions considered here, excited states remain several orders of magnitude less populated than the ground state, making the corresponding heavy particle collision rates negligible compared to ground state dominated processes. This assumption is appropriate for strong ionizing shock flows but may not hold in other plasma regimes. Radiative absorption is modeled using escape factors, Λ , which range from 0 to 1, corresponding to optically thick and optically thin plasma. A value of $\Lambda = 0$ indicates that all emitted photons are reabsorbed locally, while $\Lambda = 1$ implies complete photon escape without reabsorption. In this work, the plasma is assumed to be optically thin for all radiative transitions, with the exception of high-frequency bound-bound emissions to the ground state $\text{Ar}(1)$, which are treated as fully reabsorbed by the gas.^{14,15} Planck's constant and the radiation frequency are denoted by h and ν , respectively.

B. Governing equations

The plasma is described using a single-fluid formulation, in which all species share a common velocity field, and thermal nonequilibrium is captured by distinguishing between electron and heavy-particle temperatures, T_e and T_h , respectively. Charge neutrality is maintained by the short Debye length and the absence of applied electric fields. The assumption of strong collisional coupling justifies a single-fluid velocity. Viscous and diffusive transport effects are neglected. As a result, the dynamics are governed by the two-temperature Euler equations,¹⁵ expressed in conservation form as

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}_\alpha}{\partial x_\alpha} = \mathbf{S} \quad (8)$$

for $\alpha \in \{1, 2, 3\}$, with $\mathbf{x} = [x_1, x_2, x_3] = [x, y, z]$ being the Cartesian coordinates and t denoting time. Einstein summation is implied over

repeated indices. The conservative variables \mathbf{U} , inviscid fluxes \mathbf{F}_α , and source terms \mathbf{S} are defined as

$$\mathbf{U} = \begin{bmatrix} \rho_s^i \\ \rho u_\beta \\ \rho E \\ \rho e_e \end{bmatrix}, \quad (9)$$

$$\mathbf{F}_\alpha = \begin{bmatrix} \rho_s^i u_\alpha \\ \rho u_\alpha u_\beta + p \delta_{\alpha\beta} \\ \rho H u_\alpha \\ \rho e_e u_\alpha \end{bmatrix}, \quad (10)$$

$$\mathbf{S} = \begin{bmatrix} \omega_s^i \\ 0 \\ \Omega^R \\ \Omega_e^C + \Omega_e^R - p_e \partial u_\alpha / \partial x_\alpha \end{bmatrix}, \quad (11)$$

with $\beta \in \{1, 2, 3\}$, $s \in \mathcal{S}$, $i \in \mathcal{S}_s$, and

$$\Omega^R = \Omega_{\text{bb}}^R + \Omega_{\text{ff}}^R, \quad (12)$$

$$\Omega_e^C = \Omega_{\text{el}}^C + \Omega_{\text{in}}^C + \Omega_{\text{ion}}^C, \quad (13)$$

$$\Omega_e^R = \Omega_{\text{bf}}^R + \Omega_{\text{ff}}^R. \quad (14)$$

Here, $\mathcal{S} = \{e^-, \text{Ar}, \text{Ar}^+\}$ denotes the set of chemical species, and \mathcal{S}_s is the set of electronic states for species s . The symbols ρ_s^i and p_s^i represent the partial density and pressure of state i of species s , while ρ and p denote the total plasma density and pressure. u_α corresponds to the α -component of the mass-averaged velocity with $\delta_{\alpha\beta}$ being the Kronecker delta, whereas ρE , ρH , and ρe_e denote the total energy, total enthalpy, and electron internal energy densities, respectively. Mass source terms are denoted by ω_s^i , while energy source terms include radiative contributions from bound-bound (Ω_{bb}^R), bound-free (Ω_{bf}^R), and free-free (Ω_{ff}^R) transitions, as well as collisional energy transfer from elastic (Ω_{el}^C), inelastic (Ω_{in}^C), and ionizing (Ω_{ion}^C) collisions. The bound-free radiative term is omitted from the total energy conservation equation, as its contribution is negligible compared to the free-free and especially the bound-bound terms.¹⁴ The non-conservative term $-p_e \partial u_\alpha / \partial x_\alpha$ is related to the work of the self-consistent plasma electric field, which acts against charge separation, whose expression can be derived from the free-electron momentum equation when neglecting inertial and diffusive terms.⁷⁴

III. COMPUTATIONAL METHOD

The cell-centered finite volume discretization⁷⁵ is employed, with inviscid fluxes computed using the Van Leer flux-splitting scheme.⁷⁶ To achieve second-order spatial accuracy, left and right states at cell interfaces are reconstructed via the MUSCL approach,⁷⁷ applied to primitive variables (partial densities, velocity components, and temperatures) using van Albada's slope limiter.⁷⁸ To mitigate numerical instabilities associated with strong shock fronts, particularly when using reduced-order models, a reconstruction blending technique is implemented. In this approach, the reconstructed state at each cell interface, \mathbf{U}_f , is defined as a weighted combination of the high-order (MUSCL) reconstruction, \mathbf{U}_f^H , and the low-order (first-order upwind) reconstruction, \mathbf{U}_f^L such that

$$\mathbf{U}_f = (1 - \eta) \mathbf{U}_f^H + \eta \mathbf{U}_f^L, \quad (15)$$

where η is a shock-sensing function that transitions smoothly between 1 (upstream and within the shock) and 0 (downstream of the shock) via a Gaussian smoothing profile. Shock locations are identified using the detection algorithm described by Quirk.⁷⁹

Time integration is performed using the operator-splitting technique proposed by Strang.⁸⁰ This method integrates the transport operator, $\mathcal{T}(\mathbf{U}) = \partial \mathbf{F}_\alpha / \partial x_\alpha$, and the reactive operator, $\mathcal{R}(\mathbf{U}) = \mathbf{S}$, sequentially in a symmetric fashion:

$$\partial_t \mathbf{U}^{(1)} = \mathcal{T}(\mathbf{U}^{(1)}), \quad \mathbf{U}^{(1)}(t_n) = \mathbf{U}_n, \quad (16)$$

$$\partial_t \mathbf{U}^{(2)} = \mathcal{R}(\mathbf{U}^{(2)}), \quad \mathbf{U}^{(2)}(t_n) = \mathbf{U}^{(1)}(t_n + \Delta t/2), \quad (17)$$

$$\partial_t \mathbf{U}^{(3)} = \mathcal{T}(\mathbf{U}^{(3)}), \quad \mathbf{U}^{(3)}(t_n + \Delta t/2) = \mathbf{U}^{(2)}(t_n + \Delta t), \quad (18)$$

$$\mathbf{U}_{n+1} = \mathbf{U}^{(3)}(t_n + \Delta t), \quad (19)$$

where Δt is the time step. The splitting formulation is second-order accurate, strongly stable, and symplectic for nonlinear equations. Its convergence and stability properties have been extensively studied for reacting flow simulations.^{81–84} The transport operator is advanced in time using a four-stage, third-order strong-stability-preserving Runge–Kutta (SSP-RK3) scheme,⁸⁵ which allows a Courant–Friedrichs–Lewy (CFL) number of up to 2. The reactive operator is integrated implicitly via a second-order backward differentiation formula (BDF) solver provided by the “Livermore Solver for Ordinary Differential Equations” (LSODE) library.⁸⁶

IV. PROJECTION-BASED MODEL REDUCTION

The numerical solution of the nonlinear governing equations (8) is often computationally expensive, due to the large number of pseudo-species involved and the extremely fast time scales introduced by the CR model, which impose severe restrictions on temporal integration schemes. These challenges are typical when solving thermochemical StS models. To alleviate these issues, ROMs are employed to decrease the number of pseudo-species integrated in time, resulting in a reduction of computational cost compared to the FOM. In this work, ROMs are constructed using Petrov–Galerkin projections, following the procedure outlined by Zanardi *et al.*⁶⁶

A. OD state-space representation

To reduce the number of pseudo-species in the simulation, we adopt a 0D reactor formulation of the governing equations (8), neglecting transport terms and expressing the system in terms of primitive variables, namely, species mass densities and temperatures. Exploiting mass conservation, where $d\rho/dt = 0$, the system can be expressed compactly as

$$\frac{d}{dt} \mathbf{q}(t) = \mathbf{f}(\mathbf{q}(t); \rho), \quad \mathbf{q}(0) = \mathbf{q}_0(\boldsymbol{\mu}), \quad (20)$$

where the thermochemical state vector $\mathbf{q}(t) = [w_s^i, T_h, T_e] \in \mathbb{R}^N$ includes the mass fractions w_s^i associated with each electronic state $i \in \mathcal{S}_s$ of species $s \in \mathcal{S}$, along with the heavy-particle and electron temperatures, T_h and T_e . Here, $N = 2 + \sum_{s \in \mathcal{S}} |\mathcal{S}_s|$ denotes the total number of variables, accounting for the two temperatures and all internal states across all species, while $|\mathcal{S}_s|$ represents the cardinality of the set \mathcal{S}_s . The initial condition $\mathbf{q}(0)$ is prescribed by the function

$\mathbf{q}_0 : \mathbb{R}^{N_0} \rightarrow \mathbb{R}^N$, which depends on a set of parameters $\boldsymbol{\mu} \in \mathbb{R}^{N_0}$ that macroscopically characterizes the reactor’s initial state (see Sec. IV B for practical details).

In addition, we define an output vector $\mathbf{y}(t)$ as

$$\mathbf{y}(t) = \mathbf{C}\mathbf{q}(t), \quad (21)$$

where \mathbf{C} is an output matrix, and $\mathbf{y}(t)$ collects physical quantities of interest. Specifically, $\mathbf{y}(t)$ contains the moments of orders 0 through m of the ASDF for heavy species only. The output dimension is therefore, $\dim(\mathbf{y}) = (m + 1)(|\mathcal{S}| - 1)$. The exclusion of electrons is justified by the quasi-neutrality of the plasma: the zeroth-order moment of Ar^+ , which is included in $\mathbf{y}(t)$, directly corresponds to that of e^- . The j -th moment corresponding to species s is defined as⁶⁶

$$y_{s,j} = (\mathbf{C}\mathbf{q})_{s,j} = \frac{1}{j!} \sum_{i \in \mathcal{S}_s} \frac{1}{M_s} (e_s^i)^j w_s^i, \quad (22)$$

where M_s is the molar mass of species s , and e_s^i denotes the energy of state i (expressed in electron-volts, eV). In this work, temperature-related outputs are not included, as the focus is solely on reducing the number of pseudo-species.

B. Normalization

Model reduction techniques based on singular value decomposition (SVD), such as the one employed in this work, are inherently sensitive to the relative magnitudes of the state variables. This sensitivity arises from the impact of variable scaling on the sample covariance matrix, which the algorithm seeks to diagonalize. In practice, the components of the state vector may span several orders of magnitude, making proper normalization a critical preprocessing step.^{55,87} To address this, we apply a linear transformation of the form

$$\mathbf{q}(t) = \mathbf{D}\tilde{\mathbf{q}}(t) + \bar{\mathbf{q}}, \quad (23)$$

where $\mathbf{D} \in \mathbb{R}^{N \times N}$ is a diagonal matrix of scaling weights and $\bar{\mathbf{q}} \in \mathbb{R}^N$ is a centering vector. We adopt Pareto scaling, in which each variable is centered by its sample mean and scaled by the square root of its standard deviation, both computed from a representative ensemble of 0D trajectories by solving Eq. (20). Unlike standard normalization, which enforces unit variance, Pareto scaling preserves the relative variance structure by assigning each variable a variance equal to its standard deviation.⁸⁷ Substituting Eq. (23) into the governing equations (20) and the output map (21), we obtain

$$\begin{aligned} \frac{d}{dt} \tilde{\mathbf{q}}(t) &= \mathbf{D}^{-1} \mathbf{f}(\mathbf{D}\tilde{\mathbf{q}}(t) + \bar{\mathbf{q}}; \rho), \quad \tilde{\mathbf{q}}(0) = \mathbf{D}^{-1} \mathbf{q}_0(\boldsymbol{\mu}) - \bar{\mathbf{q}}, \\ \mathbf{y}(t) &= \mathbf{C}\mathbf{D}\tilde{\mathbf{q}}(t) + \mathbf{C}\bar{\mathbf{q}}, \end{aligned} \quad (24)$$

which defines the normalized system upon which the reduced-order model is constructed. For notational convenience, we rewrite system (24) in a more compact form as

$$\begin{aligned} \frac{d}{dt} \tilde{\mathbf{q}}(t) &= \tilde{\mathbf{f}}(\tilde{\mathbf{q}}(t); \rho), \quad \tilde{\mathbf{q}}(0) = \tilde{\mathbf{q}}_0(\boldsymbol{\mu}), \\ \mathbf{y}(t) &= \tilde{\mathbf{C}}\tilde{\mathbf{q}}(t) + \bar{\mathbf{y}}, \end{aligned} \quad (25)$$

where the transformed terms $\tilde{\mathbf{f}}$, $\tilde{\mathbf{q}}_0$, $\tilde{\mathbf{C}}$, and $\bar{\mathbf{y}}$ are defined by matching terms between Eqs. (24) and (25).

C. Petrov–Galerkin model reduction

In (linear) Petrov–Galerkin model reduction, we replace the normalized full state $\tilde{\mathbf{q}}(t) \in \mathbb{R}^N$ by its projection onto a lower-dimensional space, $\hat{\mathbf{q}}(t) = \mathbb{P}\tilde{\mathbf{q}}(t) \in \mathbb{R}^N$, where $\mathbb{P} \in \mathbb{R}^{N \times N}$ is a rank- p projector operator (with $p \ll N$) whose range defines the approximation subspace. Inserting this ansatz into (24), it is easy to see that the dynamics of $\hat{\mathbf{q}}$ are governed by

$$\begin{aligned} \frac{d}{dt}\hat{\mathbf{q}}(t) &= \mathbb{P}\tilde{\mathbf{f}}(\mathbb{P}\hat{\mathbf{q}}(t); \rho), \quad \hat{\mathbf{q}}(0) = \mathbb{P}\tilde{\mathbf{q}}_0(\boldsymbol{\mu}), \\ \hat{\mathbf{y}}(t) &= \tilde{\mathbf{C}}\mathbb{P}\hat{\mathbf{q}}(t) + \bar{\mathbf{y}}. \end{aligned} \quad (26)$$

Although $\hat{\mathbf{q}}(t)$ remains N -dimensional, its dynamics evolve within the p -dimensional range of \mathbb{P} . To make this explicit, we can factor the projector as $\mathbb{P} = \Phi(\Psi^\top\Phi)^{-1}\Psi^\top$, where $\Phi, \Psi \in \mathbb{R}^{N \times p}$ are the trial and test basis matrices, respectively. Defining the reduced coordinate vector $\hat{\mathbf{z}}(t) = \Psi^\top\hat{\mathbf{q}}(t)$ and left-multiplying the first equation in (26) by Ψ^\top yields the final p -dimensional Petrov–Galerkin ROM:

$$\begin{aligned} \frac{d}{dt}\hat{\mathbf{z}}(t) &= \Psi^\top\tilde{\mathbf{f}}\left(\Phi(\Psi^\top\Phi)^{-1}\hat{\mathbf{z}}(t); \rho\right), \quad \hat{\mathbf{z}}(0) = \Psi^\top\tilde{\mathbf{q}}_0(\boldsymbol{\mu}), \\ \hat{\mathbf{y}}(t) &= \tilde{\mathbf{C}}\Phi(\Psi^\top\Phi)^{-1}\hat{\mathbf{z}}(t) + \bar{\mathbf{y}}. \end{aligned} \quad (27)$$

The accuracy of the reduced-order model is determined entirely by the choice of trial and test subspaces spanned by Φ and Ψ , respectively. In this work, these subspaces are constructed following the methodology proposed by Zanardi *et al.*,⁶⁶ which builds upon the CoBRAS framework recently introduced by Otto *et al.*⁶⁴ Throughout the paper, the reduced-order state $\hat{\mathbf{z}}(t) \in \mathbb{R}^p$ is taken to be $\hat{\mathbf{z}}(t) = [\hat{\mathbf{z}}_h, w_e, T_h, T_e]$, where $\hat{\mathbf{z}}_h \in \mathbb{R}^r$ contains the latent variables representing the normalized ASDF of the heavy-particle species Ar and Ar⁺. The remaining three variables, electron mass fraction w_e , heavy-particle temperature T_h , and electron temperature T_e , are neither reduced nor normalized, but are instead propagated through the model via identity mapping. The electron mass fraction is retained exactly because it corresponds to a single-state species, and its evolution equation can be used to enforce the quasi-neutrality condition of the plasma, i.e., $dn_e/dt = dn_{\text{Ar}^+}/dt$. The two temperatures are similarly excluded from the reduction, following previous findings⁸⁸ showing that applying model reduction solely to species mass fractions yields better accuracy and stability.

D. Computing the optimal projector

We aim to approximate the full-order input–output map defined by

$$\begin{aligned} F : \mathbb{R}^N &\rightarrow \mathbb{R}^{\dim(\mathbf{y})} : \tilde{\mathbf{q}}(t_0; \rho, \boldsymbol{\mu}) \\ &\mapsto (\mathbf{y}(t_0; \rho, \boldsymbol{\mu}), \mathbf{y}(t_1; \rho, \boldsymbol{\mu}), \dots, \mathbf{y}(t_{L-1}; \rho, \boldsymbol{\mu})), \end{aligned} \quad (28)$$

which maps an initial state $\mathbf{q}(t_0; \rho, \boldsymbol{\mu})$ to a sequence of L output observations along a trajectory of (24), using its Petrov–Galerkin surrogate:

$$\begin{aligned} \hat{F} : \mathbb{R}^N &\rightarrow \mathbb{R}^{\dim(\mathbf{y})} : \mathbb{P}\tilde{\mathbf{q}}(t_0; \rho, \boldsymbol{\mu}) \\ &\mapsto (\hat{\mathbf{y}}(t_0; \rho, \boldsymbol{\mu}), \hat{\mathbf{y}}(t_1; \rho, \boldsymbol{\mu}), \dots, \hat{\mathbf{y}}(t_{L-1}; \rho, \boldsymbol{\mu})). \end{aligned} \quad (29)$$

Here, the state $\tilde{\mathbf{q}}$ is a function of the temporal variable t_0 and of the physical parameters ρ and $\boldsymbol{\mu}$. Rather than directly minimizing the

mean square error $\mathbb{E}[F(\tilde{\mathbf{q}}) - \hat{F}(\mathbb{P}\tilde{\mathbf{q}})]$, we follow Zahm *et al.*,⁸⁹ who show this error is bounded by

$$\text{trace}\left(W_g(I - \mathbb{P})W_s(I - \mathbb{P})^\top\right), \quad (30)$$

where $W_s =: \mathbb{E}[\tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top] \approx XX^\top$ and $W_g =: \mathbb{E}[\nabla F\nabla F^\top] \approx YY^\top$ are the state and gradient covariance matrices associated with the system, respectively. The matrices X and Y correspond to numerical quadrature factors defined in Eqs. (20) and (21) of Rowley’s work.⁷¹ Here, W_s captures the variance of state snapshots, while W_g encodes the sensitivity of the output sequence to the initial state, with $\nabla F = D_{\tilde{\mathbf{q}}}F(\tilde{\mathbf{q}})^\top \in \mathbb{R}^{N \times L\dim(\mathbf{y})}$ and $D_{\tilde{\mathbf{q}}}$ denoting the differential (or Jacobian) operator with respect to $\tilde{\mathbf{q}}$. The optimal \mathbb{P} of Eq. (30) may be written as $\mathbb{P} = \Phi\Psi^\top$, where

$$\Phi = XV_r\Sigma_r^{-1/2}, \quad \Psi = YU_r\Sigma_r^{-1/2}, \quad (31)$$

and U_r, Σ_r and V_r are the first r singular triplets of $Y^\top X$. Rows corresponding to the retained variables (w_e, T_h, T_e) are removed from both X and Y prior to computing the singular value decomposition of $Y^\top X$. This ensures that the resulting projection operator \mathbb{P} is constructed exclusively for the normalized ASDF associated with the heavy-particle species Ar and Ar⁺ (see Sec. IV C).

In contrast to the approach of Zanardi *et al.*,⁶⁶ which relies on linearizing the dynamics to approximate the covariance factors X and Y , we obtain these directly from the full nonlinear system (24) and its associated adjoint (see Proposition 1 in Zanardi *et al.*⁶⁶) as in the original CoBRAS formulation.⁶⁴ Linearization was found to inadequately capture the dominant subspaces for the CR model considered in this work. Accordingly, for each sampled training instance $(t_0, \rho, \boldsymbol{\mu})$, we (i) integrate the nonlinear forward model to generate state snapshots for X and (ii) solve the corresponding nonlinear adjoint to obtain gradient snapshots for Y . While this approach preserves the system’s intrinsic nonlinearity, it significantly increases computational cost due to repeated adjoint integrations. To mitigate this, we adopt the subsampling strategy of Otto *et al.*,⁶⁴ selecting a uniformly sparse set of t_0 values (approximately 10% of the total) along extended trajectories to estimate the gradient covariance with far fewer adjoint integrations per trajectory.

E. Multidimensional reduced governing equations

Once the zero-dimensional ROM is constructed, it can be embedded in a multidimensional flow solver by projecting the full Euler species-continuity equations onto the reduced basis. Recall that only the heavy-particle species mass fractions are subject to dimensionality reduction, while temperatures and the electron mass fraction are retained via identity mapping. Therefore, we restrict our attention to the species-continuity equations from the Euler system (8), which are written as

$$\frac{\partial}{\partial t}\rho\mathbf{w} + \frac{\partial}{\partial x_x}\rho u_x\mathbf{w} = \boldsymbol{\omega}(\mathbf{w}; \theta), \quad (32)$$

where $\boldsymbol{\omega}$ is the source term associated with species mass densities and depends on the thermochemical state vector, expressed here as the mass fraction vector \mathbf{w} and the thermodynamic parameters $\theta = (\rho, T_h, T_e)$. The parameters in θ are obtained from the mass

and energy conservation equations, which are solved concurrently within (8).

Let \mathbf{D}_w and $\bar{\mathbf{w}}$ denote the scaling matrix and centering vector for \mathbf{w} , computed as described in Sec. IV B, and let Φ_w and Ψ_w denote the optimal trial and test basis matrices for the normalized species mass fractions, $\tilde{\mathbf{w}}$, constructed following Sec. IV D. Applying the normalization and projection procedures from Secs. IV B and IV C to Eq. (32), we obtain the reduced species-continuity equations:

$$\frac{\partial}{\partial t} \rho \hat{\mathbf{z}} + \frac{\partial}{\partial x_i} \rho u_i \hat{\mathbf{z}} = \Psi_w^\top \mathbf{D}_w^{-1} \boldsymbol{\omega} (\mathbf{D}_w \Phi_w \hat{\mathbf{z}} + \bar{\mathbf{w}}; \theta), \quad (33)$$

where the reduced coordinate vector is given by $\hat{\mathbf{z}}(t) = [\hat{z}_h, w_e] \in \mathbb{R}^{r+1}$. Equation (33) is then integrated in space and time together with the unmodified momentum and energy equations to simulate the full multidimensional plasma flow. Since the computation of thermodynamic and chemical production terms \mathbf{S} [see (11)] requires the complete set of species mass fractions, we reconstruct $\tilde{\mathbf{w}}$ from the reduced variables $\hat{\mathbf{z}}$ after each time step.⁶¹

V. NUMERICAL RESULTS

In this section, we assess the performance of the proposed CoBRAS-based model reduction strategy for CR argon plasma. Results are compared against both the FOM and the conventional POD-Galerkin approach.^{54,55,60,61} The ROMs are first constructed in a 0D setting, as described in Sec. IV, and subsequently evaluated in 1D and 2D shock-tube simulations, following the configuration described by Kapper and Cambier.¹⁵

The implementation used to build the 0D ROMs is available at <https://github.com/ivanZanardi/romar>. All 1D and 2D simulations were performed using HEGEL,⁹⁰ a parallel, multi-block, structured solver for plasma hydrodynamics, coupled with PLATO,⁹¹ a physico-chemical library for computing thermodynamic properties and evaluating source terms associated with nonequilibrium collisional and radiative processes. The ROMs were integrated into PLATO and seamlessly coupled with HEGEL to enable reduced-order simulations within the plasma flow solver.

A. Physical model validation

We first validate the CR model using a steady, 1D inviscid flow behind a normal shock. The analysis is conducted in the shock reference frame, where the shock front is modeled as a discontinuity located at $x = 0$, across which flow properties exhibit a sharp jump from their freestream values. Given the upstream velocity, pressure, and temperature, post-shock conditions are computed using the Rankine-Hugoniot jump relations, assuming frozen chemistry across the shock. The electron temperature is also held constant through the discontinuity, and any precursor ionization upstream of the shock is neglected. Under these assumptions, the governing equations reduce to a steady, 1D form of the Euler equations (8), which results in a system of ordinary differential equations. These are integrated using the LSODE library⁸⁶ within PLATO.⁹¹

Figure 1 presents a comparison between the simulated results and experimental measurements from shock-tube tests conducted at UTIAS.⁷³ Specifically, we compare total mass density (left panel) and electron number density (center panel) profiles under freestream conditions of $p_\infty = 685.2$ Pa, $T_\infty = 293.6$ K, and $\text{Ma}_\infty = 15.9$ ($u_\infty = 5074.2$ m/s). The close agreement with experimental data demonstrates that the CR model accurately captures the steady-state relaxation behavior behind the shock. We also report representative argon internal population distributions at three locations behind the shock (right panel), together with the corresponding Boltzmann distributions (black solid lines). This comparison highlights the degree of nonequilibrium in the electronic populations and their relaxation toward equilibrium throughout the post-shock region. It also supports the modeling assumption discussed in Sec. IV A that heavy-particle impact collisions between excited argon states can be neglected, since excited state populations remain 6–7 orders of magnitude lower than the ground state under the conditions considered.

The steady solution also enables calibration of the radiation model against experimental data, supporting the assumption that transitions to the ground state are fully reabsorbed. Although unsteady simulations can reproduce the experimentally observed shock-structure fluctuations driven by kinetic-wave interactions,¹⁵ we focus here on the steady case for model validation because it provides a more controlled and well-characterized benchmark.

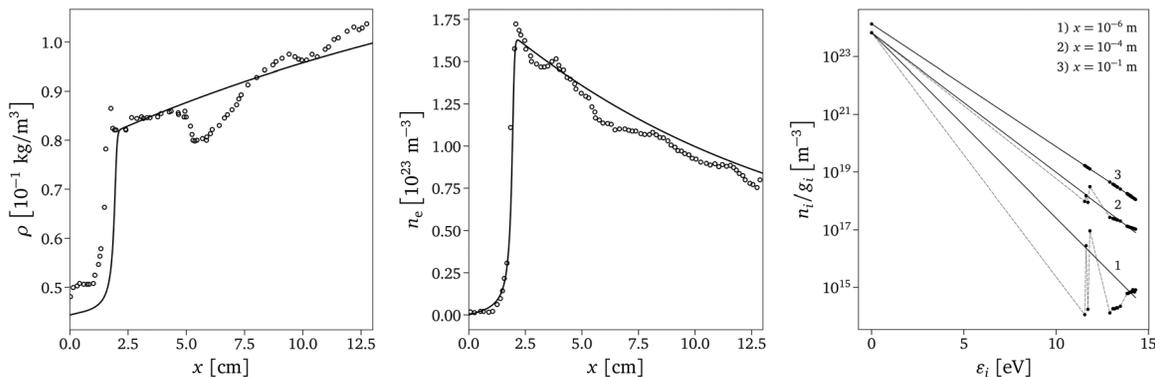


FIG. 1. Steady post-shock relaxation simulation. Comparison of the ionizing shock structure between experimental measurements (circle markers) and numerical solutions (solid line) under the following equilibrium freestream conditions: $p_\infty = 685.2$ Pa, $T_\infty = 293.6$ K, and $\text{Ma}_\infty = 15.9$ ($u_\infty = 5074.2$ m/s). The left and center panels show the total mass density ρ and electron number density n_e , respectively. The right panel presents representative argon ASDFs behind the shock at three different locations, together with their corresponding Boltzmann distributions (black solid lines).

TABLE I. Training dataset. Sampling bounds, distributions, and number of samples for each initial parameter used to generate the training trajectories.

	T_{h_0} (K)	ρ (kg/m ³)
Minimum	20 000	0.005
Maximum	50 000	0.1
Distribution	Uniform	Log-uniform
Samples	40	8

B. Zero-dimensional simulations

To build our ROM and reduce the number of pseudo-species solved, we follow the procedure described in Sec. IV, using the 0D formulation of the governing equations. Training and testing trajectories are generated by varying the total mass density ρ and the parameter vector $\boldsymbol{\mu}$, which includes the initial heavy-particle and electron temperatures (T_{h_0} and T_{e_0}) and the initial species mass fractions w_{s_0} . Specifically, we impose a sudden temperature increase on the argon mixture confined within an ideal, constant-volume reactor. In this setup, T_{h_0} varies across trajectories, while T_{e_0} is fixed at 293.6 K. The density ρ also varies between trajectories but remains constant within each simulation, consistent with the closed, fixed-volume nature of the reactor (i.e., no mass exchange with the surroundings). At the beginning of each trajectory, the initial chemical composition and the ASDF of Ar and Ar⁺ are assumed to be in thermochemical equilibrium, following a Boltzmann distribution at 293.6 K. Consequently, the explored parameters reduce to ρ and $\boldsymbol{\mu} = [T_{h_0}] \in \mathbb{R}^1$. This initialization is consistent with standard CR modeling practice for ionizing shock conditions,¹⁴ where the upstream gas is cold and essentially non-ionized and electron-impact processes rapidly dominate excitation immediately after the shock is traversed. Since our primary interest is predicting the time evolution of macroscopic observables such as mass and energy, we define the output vector $\mathbf{y} \in \mathbb{R}^4$ in (21) to include the zeroth- and first-order moments of the ASDF for both Ar and Ar⁺.

The sampling bounds and distributions used to generate the training dataset are summarized in Table I. Specifically, we generate a uniform grid of 40 values for T_{h_0} and 8 logarithmically spaced values for ρ , yielding a total of 320 distinct training trajectories.

To assess the generalization performance of the ROM, we define three testing datasets, summarized in Table II. The first dataset samples within the training domain and is used to evaluate interpolation accuracy. The second and third datasets represent extrapolation scenarios: the second targets a high-temperature regime ($T_{h_0} > 50\,000$ K), corresponding to more extreme shock conditions, while the third focuses on a low-temperature regime ($T_{h_0} < 20\,000$ K).

TABLE II. Testing datasets. Each testing dataset is defined by the lower and upper bounds of the sampled initial parameters: initial heavy-particle temperature (T_{h_0}) and density (ρ), along with the number of samples and whether the dataset is used for interpolation or extrapolation.

Dataset	T_{h_0} (K)	ρ (kg/m ³)	Samples	Type
1	$[2, 5] \times 10^4$		100	Interpolation
2	$[5, 7] \times 10^4$	[0.005, 0.1]	50	Extrapolation
3	$[1, 2] \times 10^4$		50	Extrapolation

Table III compares the mean relative error of CoBRAS- and POD-based ROMs across the test trajectories of the first dataset. This comparison highlights the effectiveness of the proposed model reduction technique relative to a more conventional approach. Errors are reported for key quantities of interest, including the mixture number density n , the mass fraction vector $\mathbf{w} \in \mathbb{R}^{34}$ (averaged across all pseudo-species), species zeroth- and first-order moments (molar fractions and internal energies, respectively), and the two temperatures, across different values of the reduced dimension r . Here, r denotes the number of latent variables required to represent the normalized ASDF of Ar and Ar⁺ combined (see Sec. IV C). Table III reports the lower error in bold for each quantity, consistently showing that CoBRAS outperforms POD across all reduced dimensions and target variables. In particular, CoBRAS yields significantly more accurate predictions of the electron molar fraction x_e and, by quasi-neutrality, the Ar⁺ molar fraction, which is critical for capturing the plasma ionization level. Furthermore, the CoBRAS model achieves high accuracy in reconstructing the mass fraction vector, despite the fact that the ROM is trained only on macroscopic outputs (chemical species total masses and internal energies). This demonstrates the model’s ability to reconstruct fine-scale features from low-order moment targets, while also reducing adjoint computations by limiting the number of output quantities.

Additionally, the errors produced by the POD Galerkin ROM do not decrease monotonically with increasing reduced dimension, as illustrated by the electron molar fraction x_e . While POD guarantees a monotonic decay of the snapshot reconstruction error, this property does not extend to the error of the projected dynamical system. Because POD modes are selected to maximize snapshot variance rather than sensitivity to specific outputs, increasing the reduced dimension may improve the representation of high variance state components while degrading low variance but dynamically sensitive directions that control ionization. As a result, the inclusion of additional modes can lead to increased error in x_e despite the higher nominal model dimension.

TABLE III. Testing error for dataset 1 of Table II. Mean relative errors (%) for the quantities of interest, computed using CoBRAS and POD ROMs with varying reduced dimensions r (total reduced state dimension $d = r + 3$, including the electron mass fraction and the two temperatures). Results are compared against the FOM solutions. Bold values indicate the lower error between CoBRAS and POD for each entry.

	Testing error (%)—Dataset 1					
	CoBRAS			POD		
r	8	9	10	8	9	10
n	0.013	0.002	0.000	0.032	0.089	0.084
\mathbf{w}	2.005	1.205	1.397	3.638	2.684	2.773
x_e	0.239	0.037	0.009	0.827	2.385	2.098
x_{Ar}	0.021	0.003	0.000	0.101	0.092	0.088
x_{Ar^+}	0.101	0.028	0.014	2.496	0.671	0.574
e_{Ar}	0.270	0.061	0.018	1.003	1.295	1.232
e_{Ar^+}	0.010	0.003	0.003	0.178	0.037	0.028
T_h	0.019	0.004	0.001	0.085	0.077	0.076
T_e	0.129	0.056	0.047	0.241	0.119	0.108

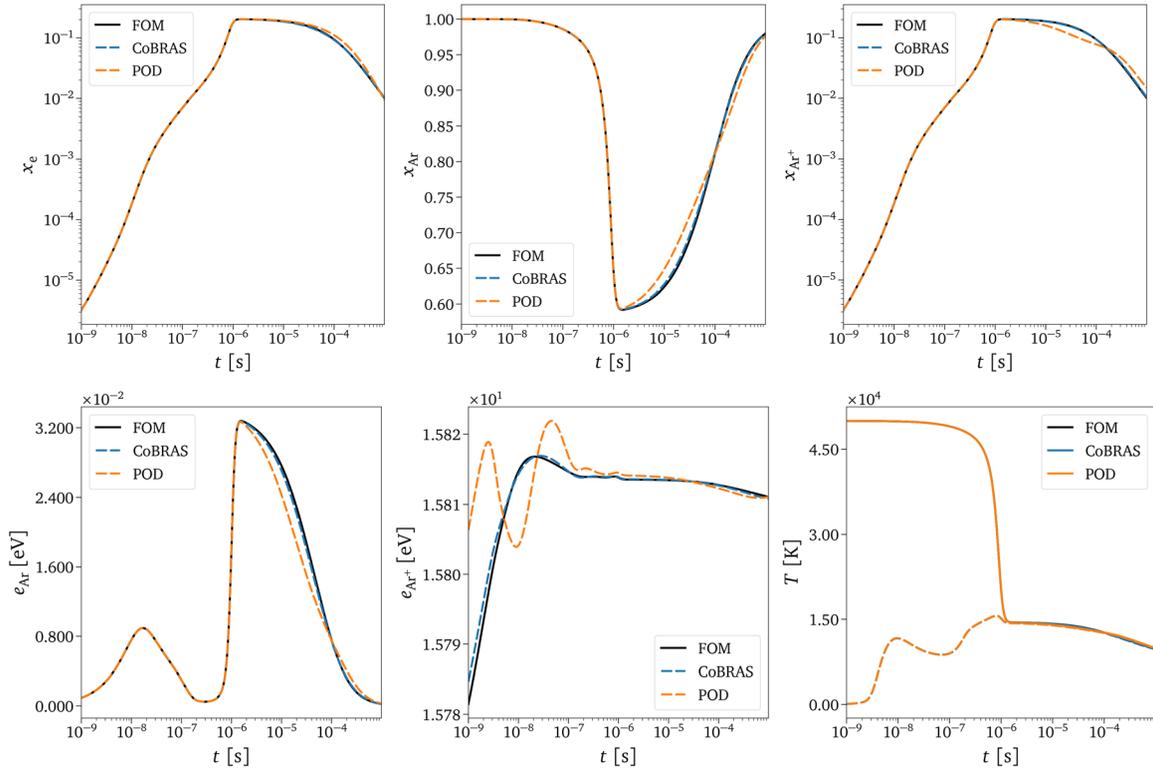


FIG. 2. FOM vs ROMs for 0D simulations: Test case 3. Time evolution of species zeroth-order moments (molar fractions), first-order moments (internal energies), and temperatures, as predicted by the FOM and by CoBRAS and POD reduced-order models with dimension $r = 8$. Results correspond to test case 3 described in Table IV.

Similar error trends are observed in Tables S1 and S2 in the [supplementary material](#), which present results for the two extrapolation datasets. While both ROMs exhibit increased errors in the low-temperature regime, CoBRAS maintains competitive accuracy for $r \geq 9$, even under extrapolative conditions. The performance degradation at lower temperatures likely stems from discrepancies between the subspaces required to accurately capture low-temperature dynamics and those spanned by the training data (see Table I). This is consistent with the physics of the system: the reaction operators governing species evolution are highly temperature-dependent, with rate constants governed by Arrhenius-type laws. Due to their exponential form, $\exp(-1/T)$, these rates exhibit greater sensitivity at lower temperatures, resulting in sharper gradients in the ASDFs and increased difficulty for the reduced basis to generalize accurately.

Figure 2, together with Figs. S1 and S2 ([supplementary material](#)), presents a side-by-side comparison of CoBRAS and POD ROMs predictions with reduced dimension $r = 8$, against FOM results for three test trajectories spanning mild (case 1) to harsh (case 3) conditions, as summarized in Table IV.

In all figures, the top row shows the time evolution of the zeroth-order moments (i.e., molar fractions) of e^- , Ar, and Ar^+ , while the bottom row displays the first-order moments (internal energies) of Ar and Ar^+ , along with the evolution of heavy-particle and electron temperatures. For the first two test cases, both ROMs show good agreement with the FOM, as illustrated in Figs. S1 and S2, except for the Ar^+ internal energy (e_{Ar^+}) evolution, where CoBRAS consistently outperforms POD.

In contrast, under the more demanding conditions of case 3 (Fig. 2), CoBRAS shows superior performance across all quantities, corroborating the error trends reported in Tables III, S1, and S2. Only the temperature profiles appear in good agreement with the FOM for both ROMs. Small deviations arise after the electron peak and within the thermal equilibrium region¹⁴ for the POD ROM, reflecting its reduced accuracy in predicting the evolution of internal energies. These discrepancies are subtle in the plots but are clearly reflected in the quantitative error metrics.

It is also worth noting that, as previously reported by Zanardi *et al.*,⁶⁶ neither CoBRAS nor POD enforces positivity of the mass fractions. In our simulations, this limitation led to small non-physical negative values at early times ($t < 10^{-8}$ s).

1. Computational performance

In this section, we assess computational efficiency by comparing the total floating-point operations (FLOPs) required by the FOM and

TABLE IV. Test cases. Initial parameter values used to generate the three testing trajectories for visual comparison between the ROMs and FOM solutions.

Case	T_{h_0} (K)	ρ (kg/m ³)
1	20 000	0.005
2	35 000	0.01
3	50 000	0.1

TABLE V. FOM vs CoBRAS for 0D simulations: Computational cost. The table reports total FLOPs and FLOP/Byte ratios for key numerical operations, including right-hand side (RHS) evaluation and linear system solve (LSS). Results are shown for the FOM and three CoBRAS ROM configurations.

Model	d	FLOPs		FLOP/Byte	
		RHS - $O(d^2)$	LSS - $O(d^3)$	RHS - $O(1)$	LSS - $O(d)$
FOM	36	2.556×10^3	3.370×10^4	0.234	3.079
	13	3.250×10^2	1.803×10^3	0.208	1.156
CoBRAS	12	2.760×10^2	1.440×10^3	0.205	1.071
	11	2.310×10^2	1.129×10^3	0.202	0.987

the CoBRAS-based ROM as the reduced dimension r varies. In addition, we report the *arithmetic intensity* for key numerical kernels, defined as the ratio of FLOPs to memory access in bytes (FLOP/Byte), following the formulation of Zanardi *et al.*⁶⁶ While actual runtimes are influenced by hardware-specific features such as memory bandwidth, processor speed, and cache hierarchy, FLOP counts and arithmetic intensity offer consistent, architecture-agnostic comparisons.

As shown in Table V, where d is the system dimension ($d = r + 3$ for the ROM), the CoBRAS-based ROM yields significant reductions in computational cost. In particular, the number of FLOPs required for right-hand side (RHS) evaluations, which are dominated by dense matrix–vector multiplications scaling as $O(d^2)$ [see Eqs. (52)–(55) in Kapper and Cambier¹⁴] is reduced by more than an order of magnitude. In our current implementation, these savings are estimated rather than fully realized, as we do not assemble reduced operators offline. Instead, the full-order state is reconstructed at each time step to evaluate the RHS of the FOM, which is subsequently projected to compute the evolution of the reduced coordinates. While this approach avoids the need for precomputing and storing reduced operators, it limits the practical benefits of the theoretical FLOP reductions. Nonetheless, in settings where reduced operators are assembled offline, as demonstrated by Zanardi *et al.*,⁶⁶ the estimates presented here remain directly applicable. The cost of solving the linear system (LSS) associated with the implicit BDF integrator, which scales as $O(d^3)$, is similarly reduced by over a factor of 30. However, the actual efficiency of both RHS and LSS operations is ultimately constrained by their

arithmetic intensity. In the ROM setting, these operations exhibit relatively low FLOP/Byte ratios, suggesting that performance is likely to be memory-bound rather than compute-bound. Consequently, the theoretical FLOP reductions may not translate linearly into runtime gains due to limitations in memory access speed.

Two additional performance metrics are examined to assess numerical stiffness and solver sensitivity. The first is the fastest time-scale, approximated as the inverse of the smallest eigenvalue of the Jacobian, $\tau = 1/|\lambda_{\min}|$. The second is the condition number κ of the matrix \mathbf{P} defined in Eq. (A7) in the Appendix (with $\beta_0 = 1$, i.e., backward Euler method) used in the linear solve of the BDF scheme. Both metrics are computed as averaged values over 100 test trajectories with fixed density $\rho = 10^{-2}$ kg/m³, while varying the initial heavy-particle temperature T_{h0} (within the training regime). As illustrated in Fig. 3 (left panel), the CoBRAS-based ROM exhibits reduced stiffness, as indicated by larger timescales, which enables bigger time steps during integration and contributes to further computational savings. Similarly, the right panel of Fig. 3 shows that the ROM achieves a lower condition number, indicating improved numerical conditioning and a more stable and efficient linear solve. Together, these results highlight that CoBRAS not only reduces dimensional complexity but also improves the numerical properties of the system.

We emphasize that the computational savings achieved by the proposed ROM are not limited to the 0D configuration. In multidimensional simulations, the reduced number of state variables translates directly into fewer scalar quantities to be transported, leading to additional performance gains. Specifically, when using explicit time integration for the transport operator, a linear speedup proportional to the reduction in state dimension can be expected. For steady-state simulations employing implicit schemes, the computational advantage becomes even more pronounced, potentially scaling cubically, due to the reduced cost of solving smaller linear systems.

C. One-dimensional simulations

We evaluate the performance of CoBRAS in a canonical 1D shock-tube configuration, where the coupling between transport and collisional–radiative kinetics plays a critical role in determining the plasma dynamics. This test case serves as a benchmark for assessing the ROM’s ability to accurately capture the highly unsteady and nonlinear behavior

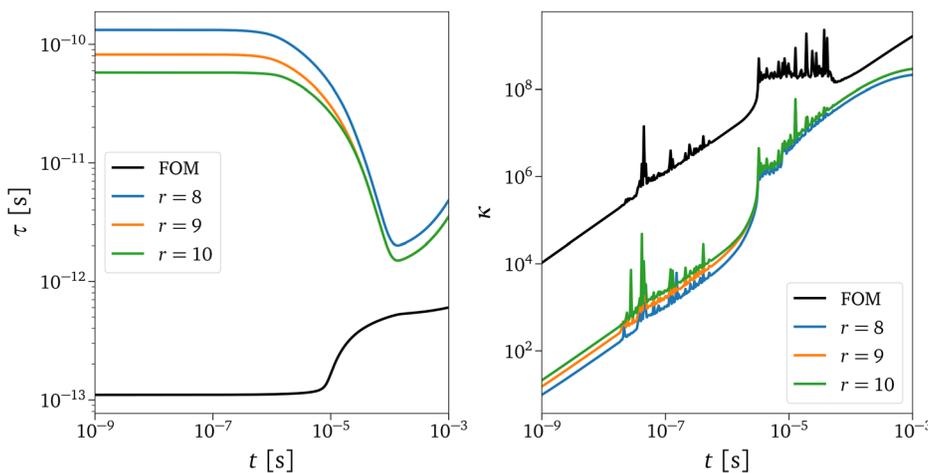


FIG. 3. FOM vs CoBRAS for 0D simulations: Integration scheme parameters. Comparison between the FOM and CoBRAS ROM across different reduced dimensions r , showing the smallest time-scale (left panel) and condition number (right panel) of the implicit integration scheme. These quantities are computed using the Jacobian of the respective system.

of the flow. The numerical discretization and time integration schemes are detailed in Sec. III, while the governing equations for the reduced latent variables are derived in Sec. IV E. The ROM used here employs a latent dimension of $r = 9$, which, when combined with the retained components (electron mass fraction and two temperatures), results in a reduced state vector of size $d = r + 3 = 12$. This represents a $3\times$ reduction compared to the full thermochemical state vector of 36 variables.

The simulation setup follows the shock-reflection experiments of Kapper and Cambier.¹⁵ In this configuration, a supersonic stream of argon gas impinges upon an ideally reflecting wall, generating a strong shock that propagates upstream against the inflow. The selected free-stream conditions are $p_\infty = 685.2$ Pa, $T_\infty = 293.6$ K, $u_\infty = 4\,535$ m/s,

chosen to reproduce a Mach 15.9 shock in the laboratory frame, consistent with UTIAS shock-tube measurements.⁷³ Although the initial conditions correspond to a stronger (Mach 19) shock, the reflected shock rapidly weakens as thermal energy is converted into electronic excitation and ionization, yielding the lower effective Mach number observed experimentally. The computational domain spans $x \in [0, 0.2]$ m, discretized with a spatial resolution of $\Delta x = 10^{-4}$ m. Time integration is carried out over $t \in [0, 2.5 \times 10^{-4}]$ s with a time step of $\Delta t = 5 \times 10^{-8}$ s. Boundary conditions consist of a supersonic inflow on the left and a perfectly reflecting wall on the right.

Figure 4 presents space–time diagrams (space horizontal, time vertical) for three key quantities: pressure p , heavy-particle temperature T_h ,

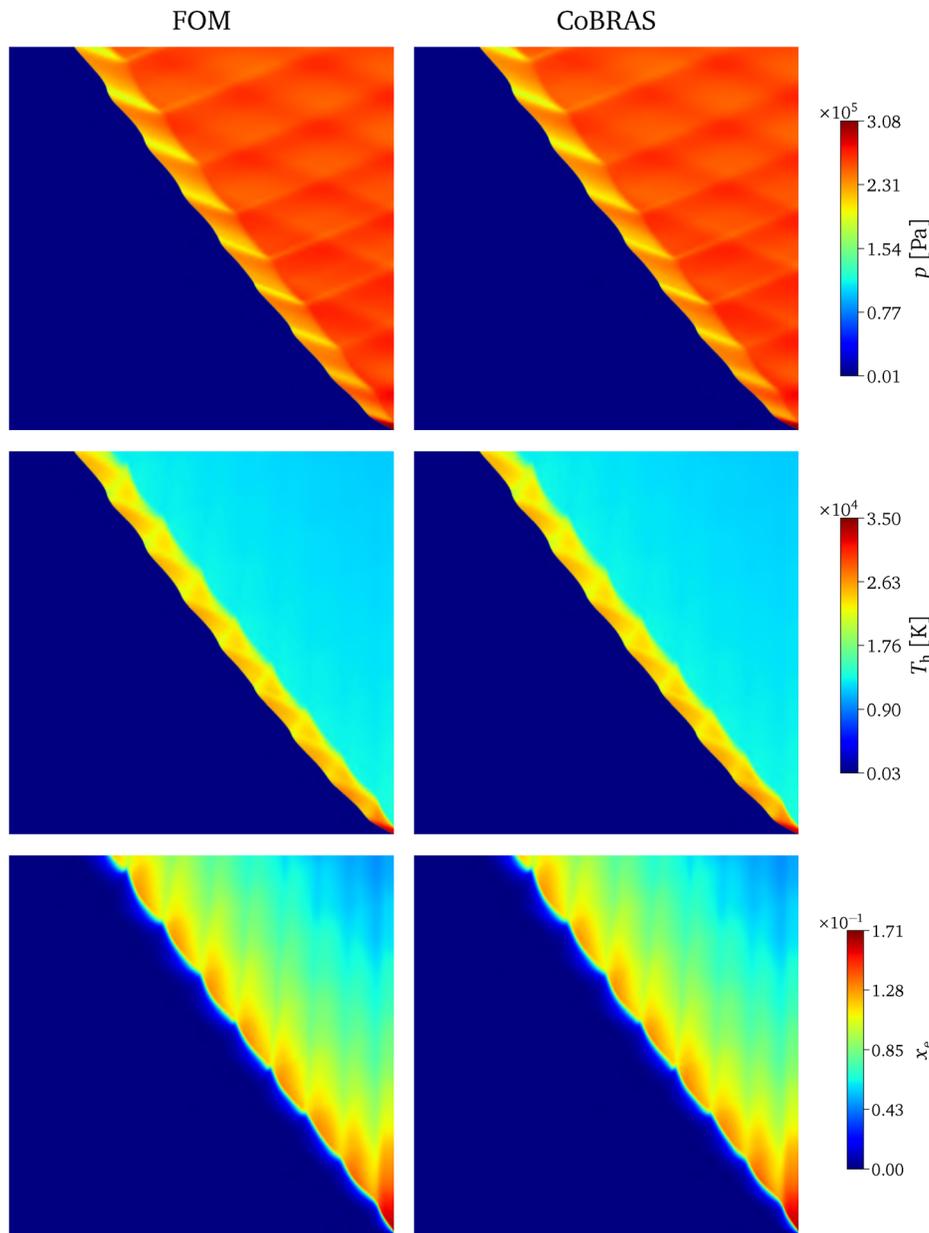


FIG. 4. FOM vs CoBRAS for 1D simulations: Space–time diagrams. Comparison of FOM and CoBRAS ROM ($r = 9$) for pressure p , heavy-particle temperature T_h , and electron molar fraction x_e in space–time diagrams (space horizontal, time vertical) obtained from the 1D simulation.

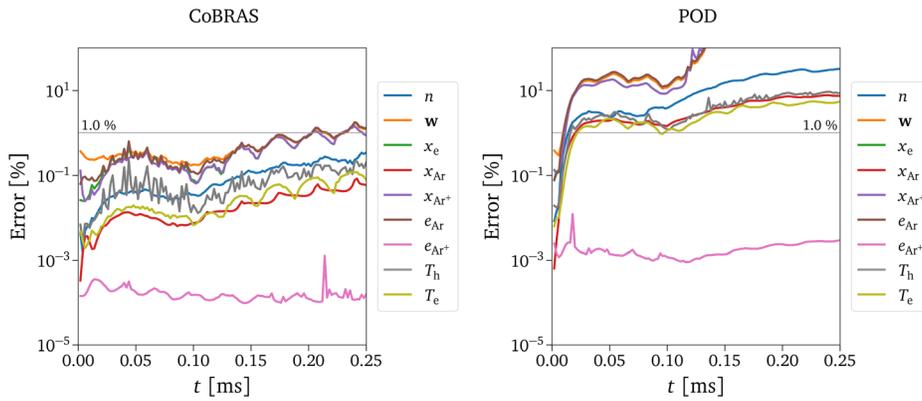


FIG. 5. Error evolution of ROMs in 1D simulations. Space-averaged relative errors (%) for CoBRAS (left panel) and POD (right panel) ROMs with $r = 9$, computed for key quantities of interest.

and electron molar fraction x_e . For each variable, the left column shows the FOM solution and the right column shows the CoBRAS prediction. Both models capture the characteristic undamped, periodic fluctuations observed experimentally in such flows.¹⁵ The shock front is visible as the leftmost propagating feature, followed by the electron avalanche, defined by the first peak in electron density. These structures travel from the bottom-right to the top-left of each panel, with the induction region lying between them. The dominant physical feature of this flow is the presence of periodic disturbances within the induction zone. The electron avalanche generates a pressure wave that travels back to the shock. Upon interacting with the shock front, this

wave alters the local conditions, strengthening the shock and generating a new entropy wave that travels downstream. This entropy wave accelerates excitation and ionization processes, triggering the formation of a subsequent electron avalanche. The new avalanche, in turn, emits another pressure wave, thus repeating the cycle of fluctuations. These mechanisms were analyzed in detail by Kapper and Cambier.¹⁵ The ROM accurately captures this cyclic behavior, preserving both the phase and amplitude of the fluctuations and faithfully reproducing the underlying physical processes.

A quantitative comparison is shown in Fig. 5, which reports the evolution of space-averaged relative errors for a set of thermochemical

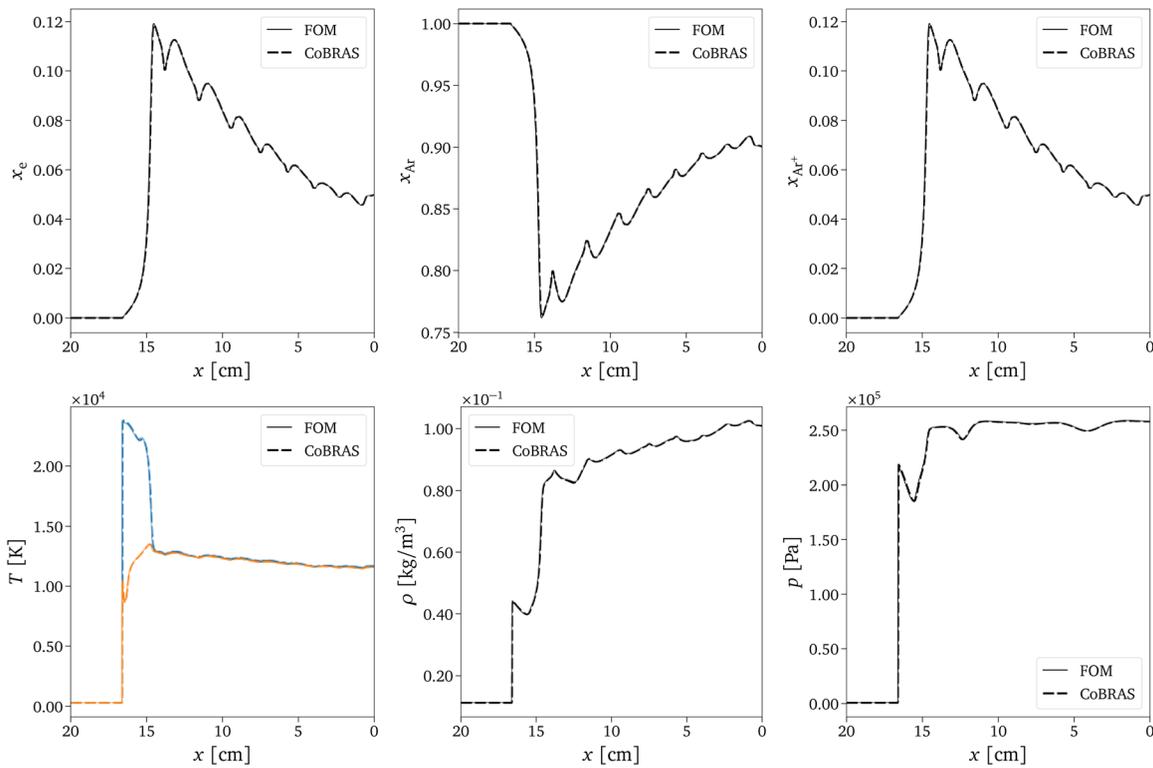


FIG. 6. Final-time comparison of FOM and CoBRAS for 1D simulations. Comparison of FOM and CoBRAS ($r = 9$) at the final time $t = 2.5 \times 10^{-4}$ s for the 1D simulation presented in Fig. 4. Top row: species molar fractions of e^- , Ar, and Ar^+ . Bottom row: heavy-particle temperature T_h (blue), electron temperature T_e (orange), density ρ , and pressure p .

variables: the mixture number density n , the mass fraction vector $\mathbf{w} \in \mathbb{R}^{34}$ (averaged across all pseudo-species), zeroth- and first-order moments of ASDFs, and the heavy-particle and electron temperatures. To ensure the error metric reflects meaningful dynamics, points within the free-stream region, where values remain constant, are excluded from the computation, preventing artificial deflation of the error. The CoBRAS-based ROM consistently maintains relative errors below 1% throughout the simulation, demonstrating excellent agreement with the FOM. In contrast, a standard POD-based ROM exhibits instability, with errors growing rapidly after approximately $t = 0.1$ ms. This highlights the robustness of the CoBRAS projection framework in stiff, nonequilibrium kinetic regimes, where conventional dimensionality reduction techniques may fail.

To further evaluate the ROM's spatial accuracy, Fig. 6 compares the final-time ($t = 2.5 \times 10^{-4}$ s) snapshots of macroscopic fields. The top row shows molar fractions of e^- , Ar, and Ar^+ , while the bottom row displays heavy-particle temperature T_h , electron temperature T_e , density ρ , and pressure p . The CoBRAS-based ROM closely matches the FOM across all variables. In particular, the sharp transitions and oscillatory features in the molar fractions and temperature fields are accurately retained, confirming the ROM's ability to resolve fine-scale plasma dynamics without sacrificing stability or fidelity.

In addition to local field comparisons, we also track global quantities that are relevant to experimental observables. Figure 7 shows the evolution of the instantaneous shock Mach number and the induction length, defined by the distance between the shock front and the first peak electron density. Both quantities exhibit undamped, nonlinear oscillations, consistent with prior observations. The oscillation period is approximately $30.57 \mu\text{s}$ for the FOM and $30.71 \mu\text{s}$ for the ROM, indicating strong agreement. The induction length undergoes discontinuous shifts, further highlighting the nonlinear character of the system. The CoBRAS model captures these behaviors with high fidelity, further validating its predictive capability.

Additional test cases are provided in the supplementary material for freestream velocities of $u_\infty = 4000$ m/s and $u_\infty = 5000$ m/s. Figures S3 and S5, as well as Figs. S6 and S8, show the corresponding space-time evolutions and final-time comparisons for the lower and higher velocity cases, respectively. In both cases, CoBRAS continues to match the FOM closely across all major variables. Moreover, Figs. S4 and S7 confirm that CoBRAS consistently outperforms POD in terms of accuracy, reinforcing its effectiveness across a wide range of conditions.

D. Two-dimensional simulations

Building on the insights gained from the 1D simulations, particularly the identification of oscillation mechanisms and the resulting wave patterns, we extend our analysis to a more complex 2D shock-reflection problem. In the 2D setting, disturbances are no longer confined to the longitudinal direction, allowing for the development of transverse instabilities and rich wave interactions. This presents a rigorous test for the CoBRAS-based ROM and its ability to capture multi-dimensional plasma flow features.

The computational domain spans $(x, y) \in [0, 0.36] \times [0, 0.18]$ m and is discretized using uniform grid spacing $\Delta x = \Delta y = 3 \times 10^{-4}$ m. Time integration proceeds over the interval $t \in [0, 5] \times 10^{-4}$ s with a time step of $\Delta t = 10^{-7}$ s. As in the 1D case, a high-speed argon flow impinges on a perfectly reflecting wall, initiating a strong shock wave that propagates upstream. The freestream conditions are identical to those used in Sec. VC: $p_\infty = 685.2$ Pa, $T_\infty = 293.6$ K, $u_\infty = 4535$ m/s. To promote the growth of transverse instabilities and facilitate the emergence of multidimensional structures, we introduce a weak Gaussian pressure perturbation near the right wall. The perturbation has an amplitude three times the freestream pressure and density, a spatial extent of 3×10^{-3} m in the x -direction, and is centered at 40% of the shock-tube height in the y -direction (see Fig. S9 in the supplementary material). This perturbation accelerates the development of instabilities while remaining weak enough to avoid overwhelming the flow physics.

Figure 8 illustrates the evolution of the electron molar fraction x_e at three representative times, $t_i = [1, 3, 5] \times 10^{-4}$ s, comparing the FOM (left) with CoBRAS (right). Initially disordered transverse waves emerge from the perturbation. Over time, these waves self-organize into a repeating structure consisting of incident and reflected shocks, Mach stems, and triple points. These triple points form a regular cellular pattern in both the longitudinal and transverse directions, indicating a resonant wave interaction. The vorticity generated at these intersections gives rise to ionization cells, which are highly structured regions of increased temperature, electron density, and vorticity, resembling detonation cells in chemical reactive flows. These localized peaks radiate pressure waves outward, reinforcing the oscillatory motion and maintaining the flow instability.¹⁵ The CoBRAS-based ROM accurately captures these dynamics. It reproduces not only the global structure and timing of wave interactions, but also the fine-scale features such as vorticity generation and ionization cell formation. Additional field

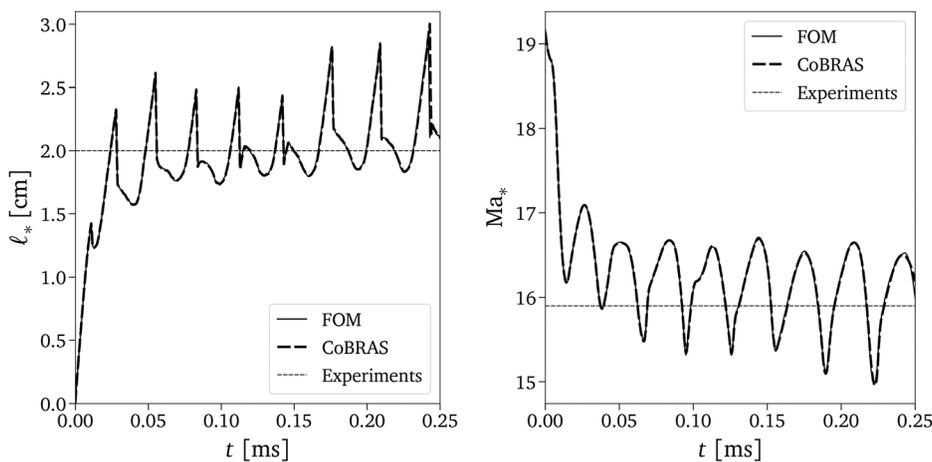


FIG. 7. Comparison against average experimental values for 1D simulations. The FOM and CoBRAS ROM ($r = 9$) are compared against average experimental values for the induction length (left panel) and the shock Mach number (right panel).

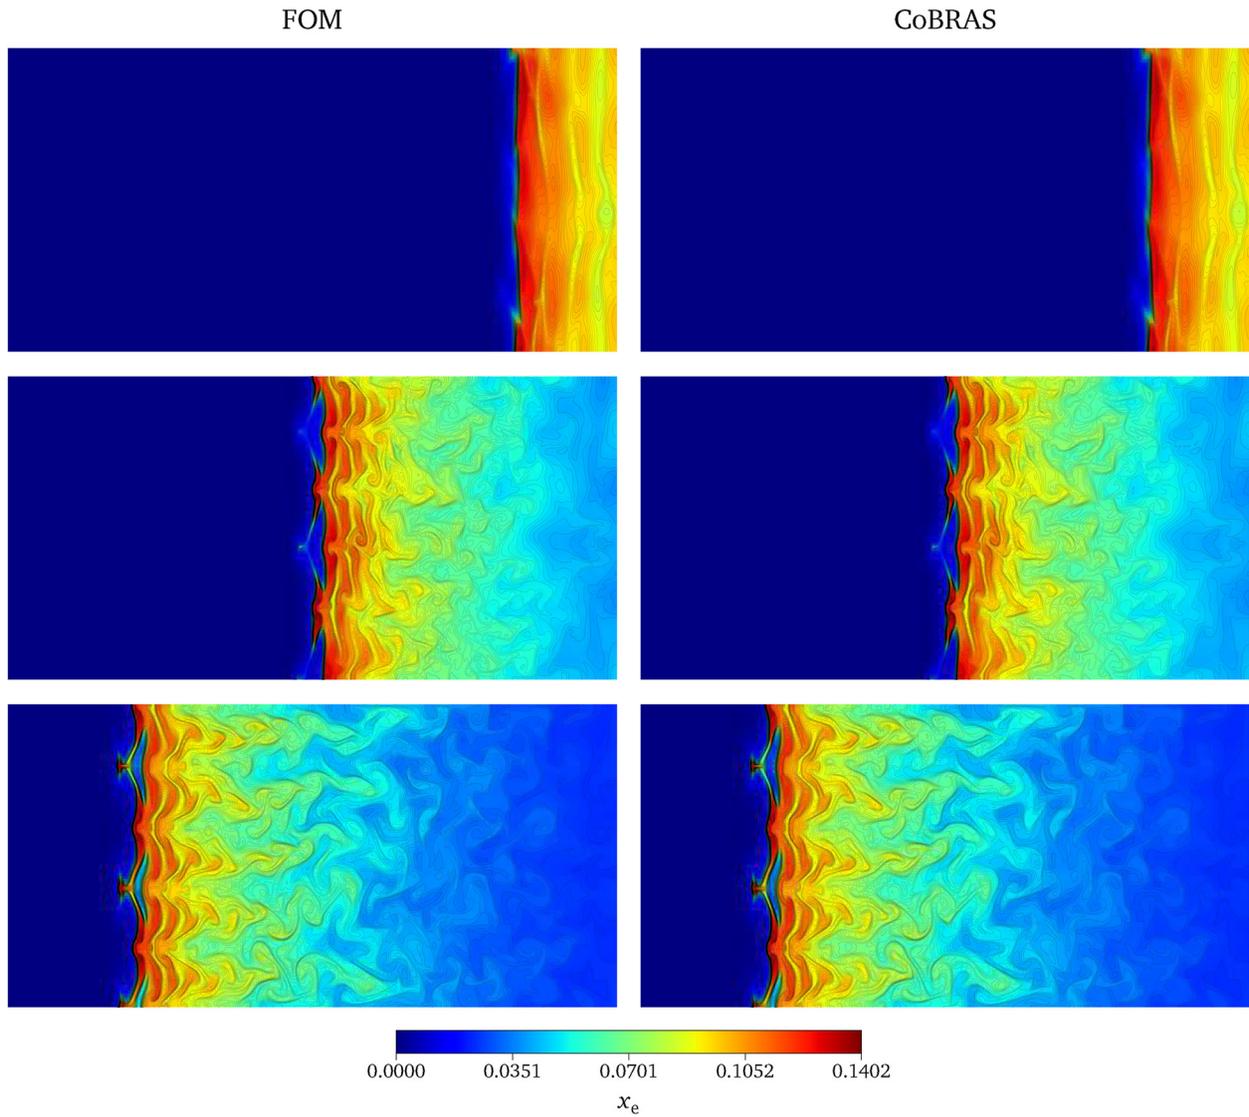


FIG. 8. FOM vs CoBRAS for 2D simulations: Electron molar fraction snapshots. Comparison of FOM and CoBRAS ROM ($r = 9$) at three time instants: $t_i = [1, 3, 5] \times 10^{-4}$ s for the electron molar fraction, x_e .

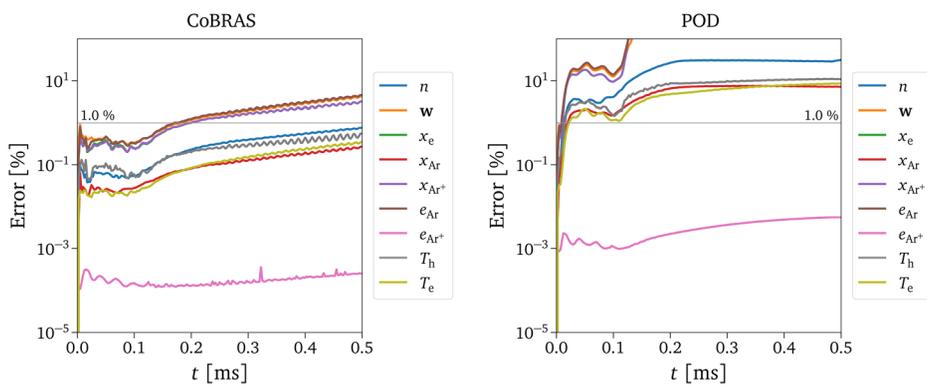


FIG. 9. Error evolution of ROMs in 2D simulations. Space-averaged relative errors (%) for CoBRAS (left panel) and POD (right panel) ROMs with $r = 9$, computed for key quantities of interest.

comparisons are provided in the [supplementary material](#) for mass density ρ (Fig. S10), pressure p (Fig. S11), heavy-particle temperature T_h (Fig. S12), and electron temperature T_e (Fig. S13), all of which show excellent agreement between the FOM and CoBRAS.

As in the 1D case, a quantitative assessment is presented in Fig. 9, which shows the space-averaged relative error for the same set of thermochemical variables. The CoBRAS-based ROM achieves consistently low errors throughout the simulation, with relative errors remaining below 10% and falling below 1% for most quantities. In contrast, POD becomes unstable, with rapidly growing errors beyond $t = 0.15$ ms. This comparison highlights the robustness and accuracy of the CoBRAS framework in handling highly unsteady, nonlinear, and multidimensional plasma flows, regimes in which conventional projection-based ROMs such as POD often fail due to the lack of physically consistent basis selection.

VI. CONCLUSIONS

In this work, we developed and assessed a Petrov–Galerkin reduced-order model (ROM) for collisional–radiative (CR) argon plasma kinetics. The framework builds upon the recent work of Zanardi *et al.*,⁶⁶ who applied the “Covariance Balancing Reduction using Adjoint Snapshots” (CoBRAS) method⁶⁴ to detailed state-to-state kinetics in thermochemical nonequilibrium. ROM construction was performed in a zero-dimensional (0D) setting by computing a linear projection operator through the balancing of state and gradient covariances associated with the full-order nonlinear system. The resulting ROM was then coupled with a finite-volume solver to simulate multidimensional, unsteady, ionizing shock flows in both one- (1D) and two-dimensional (2D) domains. We benchmarked the CoBRAS-based ROM against full-order model (FOM) simulations and compared its performance with standard proper orthogonal decomposition (POD) methods across a wide range of conditions.

The 0D results demonstrated that CoBRAS delivers substantial reductions in computational complexity, exceeding an order of magnitude in FLOP savings, while maintaining high accuracy. Relative errors remained consistently below 1% for macroscopic quantities and around 2% for the fine-scale atomic state distribution functions (ASDFs). In the 1D simulations, the ROM accurately captured periodic shock-induced oscillations and spatially localized features such as electron avalanches and induction zones, achieving a $3\times$ reduction in state dimension with space-averaged errors consistently below 1%. The method remained robust in highly unsteady regimes where POD-based ROMs exhibited instability and rapid error growth. In the more challenging 2D simulations, CoBRAS maintained its accuracy, successfully reproducing complex multidimensional structures, including shock–shock interactions, triple points, and ionization-driven cellular patterns. These results confirm that the CoBRAS projection framework enables efficient model reduction while preserving the essential physical mechanisms required for predictive, high-fidelity plasma simulations.

While the proposed ROM framework already demonstrates substantial gains in efficiency and accuracy, several avenues for improvement remain. One limitation is the lack of a built-in positivity constraint on the ASDF, which can result in small, nonphysical negative populations during early transients ($t < 10^{-8}$ s), as previously observed by Zanardi *et al.*⁶⁶ Future developments could incorporate positivity-preserving projection techniques or constrained optimization formulations to enhance robustness and physical consistency.

Moreover, integrating the ROM into uncertainty quantification pipelines and Bayesian inference frameworks would enable efficient exploration of large parametric spaces, especially in extrapolative regimes.

More broadly, this work contributes to the growing need for physically grounded, numerically stable reduced-order models capable of capturing multiscale, nonlinear interactions in reactive, radiating flows. Such flows arise in a wide range of fluid mechanics contexts, from plasma–shock interactions in hypersonic boundary layers to radiative magnetohydrodynamics in astrophysical jets, where high-fidelity full-order models remain prohibitively expensive. The proposed methodology offers a systematic approach to accelerating these simulations without sacrificing physical fidelity, enabling new opportunities in both fundamental fluid dynamics research and high-impact engineering applications.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for additional results for zero-, one-, and two-dimensional simulations.

ACKNOWLEDGMENTS

I. Zanardi, A. Meini, and M. Panesi acknowledge support from the Vannevar Bush Faculty Fellowship OUSD(RE), under Grant No. N00014-21-1-295. A. Padovan and D. J. Bodony acknowledge support from (a) the National Science Foundation under Grant No. 2139536, issued to the University of Illinois Urbana-Champaign by the Texas Advanced Computing Center under subaward No. UTAUS-SUB00000545 with Dr. Daniel Stanzione as the PI and (b) from the Center for Hypersonics and Entry Systems Studies at the University of Illinois Urbana-Champaign. A. Padovan gratefully acknowledges startup support from the Newark School of Engineering at New Jersey Institute of Technology.

The views and conclusions expressed in this work are those of the authors and do not necessarily reflect the official policies or endorsements, either expressed or implied, of the United States Government.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Ivan Zanardi: Conceptualization (lead); Data curation (lead); Investigation (lead); Methodology (lead); Software (lead); Validation (equal); Visualization (lead); Writing – original draft (lead); Writing – review & editing (equal). **Alessandro Meini:** Data curation (supporting); Validation (equal); Writing – original draft (supporting); Writing – review & editing (equal). **Alberto Padovan:** Conceptualization (supporting); Methodology (supporting); Writing – original draft (supporting); Writing – review & editing (equal). **Daniel J. Bodony:** Supervision (supporting); Writing – original draft (supporting); Writing – review & editing (equal). **Marco Panesi:** Funding acquisition (lead); Project administration (lead); Supervision (lead); Writing – original draft (supporting); Writing – review & editing (equal).

DATA AVAILABILITY

Code to reproduce the zero-dimensional results presented in this publication is available at <https://github.com/ivanZanardi/romar>. Code and data for the one- and two-dimensional simulations are available upon reasonable request.

APPENDIX: CONDITION NUMBER OF THE BDF SCHEME

We consider a general system of ordinary differential equations (ODEs) of the form:

$$\frac{d}{dt} \mathbf{q}(t) = \mathbf{f}(t, \mathbf{q}(t)), \quad \mathbf{q}(0) = \mathbf{q}_0, \quad (A1)$$

where $\mathbf{q}(t) \in \mathbb{R}^N$ is the state vector and $\mathbf{f} : \mathbb{R} \times \mathbb{R}^N \rightarrow \mathbb{R}^N$ is a smooth nonlinear function defining the system dynamics. To numerically solve Eq. (A1), we employ linear multistep methods.⁸⁶ At time step n , the general k -step formulation is given by

$$\mathbf{q}_n = \sum_{j=1}^{k_1} \alpha_j \mathbf{q}_{n-j} + h_n \sum_{j=0}^{k_2} \beta_j \mathbf{f}(t_{n-j}, \mathbf{q}_{n-j}), \quad (A2)$$

where h_n is the time step size at iteration n , and α_j, β_j are method-specific coefficients. Choosing $k_1 = p$ and $k_2 = 0$ yields the p -th order backward differentiation formula (BDF), which takes the form:

$$\mathbf{q}_n = \sum_{j=1}^p \alpha_j \mathbf{q}_{n-j} + h_n \beta_0 \mathbf{f}(t_n, \mathbf{q}_n). \quad (A3)$$

This is an implicit method, requiring the solution of a nonlinear system at each time step. We solve this system using the Newton–Raphson (NR) iteration. At iteration m of time step n , the NR update solves the linear system:

$$\begin{aligned} \mathbf{P}(\mathbf{q}_n^{[m+1]} - \mathbf{q}_n^{[m]}) &= -\mathbf{r}(\mathbf{q}_n^{[m]}) \\ &= \sum_{j=1}^p \alpha_j \mathbf{q}_{n-j} + h_n \beta_0 \mathbf{f}(\mathbf{q}_n^{[m]}) - \mathbf{q}_n^{[m]}, \end{aligned} \quad (A4)$$

where $\mathbf{q}_n^{[m]}$ is the m -th NR iterate, and $\mathbf{r}(\mathbf{q}_n^{[m]})$ is the residual of the BDF equation evaluated at $\mathbf{q}_n^{[m]}$. The Jacobian matrix \mathbf{P} of the residual with respect to \mathbf{q} is given by

$$\mathbf{P} = \frac{\partial \mathbf{r}}{\partial \mathbf{q}} = \mathbf{I} - h_n \beta_0 \mathbf{J}, \quad (A5)$$

where \mathbf{I} is the identity matrix, and $\mathbf{J} = \partial \mathbf{f} / \partial \mathbf{q}$ is the Jacobian of the system dynamics with respect to the state vector. The condition number of the matrix \mathbf{P} is defined as

$$\kappa(\mathbf{P}) = \frac{|\lambda_{\max}|}{|\lambda_{\min}|}, \quad (A6)$$

where λ_{\max} and λ_{\min} denote the eigenvalues of \mathbf{P} with the largest and smallest magnitudes, respectively. Since the time step h_n varies over the course of the integration and only approximate condition number estimates are needed, it is convenient to define a continuous, time-dependent version of \mathbf{P} as

$$\mathbf{P}(t, \mathbf{q}) = \mathbf{I} - t \beta_0 \mathbf{J}(t, \mathbf{q}), \quad (A7)$$

which mirrors the structure of Eq. (A5) by interpreting t as a characteristic time scale of the integration step.

REFERENCES

- ¹D. M. Goebel and I. Katz, *Fundamentals of Electric Propulsion: Ion and Hall Thrusters* (John Wiley and Sons, 2008), pp. 1–507.
- ²S. V. Pancheshnyi, D. A. Lacoste, A. Bourdon, and C. O. Laux, “Ignition of propane-air mixtures by a repetitively pulsed nanosecond discharge,” *IEEE Trans. Plasma Sci.* **34**, 2478–2487 (2006).
- ³C. Park, *Nonequilibrium Hypersonic Aerothermodynamics* (Wiley, New York, 1990).
- ⁴B. Helber, A. Turchi, J. B. Scoggins, A. Hubin, and T. E. Magin, “Experimental investigation of ablation and pyrolysis processes of carbon-phenolic ablators in atmospheric entry plasmas,” *Int. J. Heat Mass Transfer* **100**, 810–824 (2016).
- ⁵J. P. H. Goedbloed and S. Poedts, *Principles of Magnetohydrodynamics* (Cambridge University Press, 2004).
- ⁶I. Hubeny and D. Mihalas, *Theory of Stellar Atmospheres. An Introduction to Astrophysical Non-Equilibrium Quantitative Spectroscopic Analysis* (McGraw-Hill, 2015).
- ⁷M. Carlsson and R. F. Stein, “Dynamic hydrogen ionization,” *Astrophys. J.* **572**, 626–635 (2002).
- ⁸A. Munafo, N. N. Mansour, and M. Panesi, “A reduced-order NLTE kinetic model for radiating plasmas of outer envelopes of stellar atmospheres,” *Astrophys. J.* **838**, 126 (2017).
- ⁹S. Kumar, A. Munafo, S. M. Jo, and M. Panesi, “Investigation of non-equilibrium phenomena in nitrogen RF inductively coupled plasma discharges: A state-to-state approach,” *J. Phys. D: Appl. Phys.* **58**, 025204 (2025).
- ¹⁰D. K. Morozov, E. O. Baronova, and I. Y. Senichenkov, “Impurity radiation from a tokamak plasma,” *Plasma Phys. Rep.* **33**, 906–922 (2007).
- ¹¹A. Kallenbach, M. Bernert, R. Dux, L. Casali, T. Eich, L. Giannone, A. Herrmann, R. McDermott, A. Mlynek, H. W. Müller, F. Reimold, J. Schweinzer, M. Sertoli, G. Tardini, W. Treutterer, E. Viezzer, R. Wenninger, and M. Wischmeier, “Impurity seeding for tokamak power exhaust: From present devices via ITER to DEMO,” *Plasma Phys. Controlled Fusion* **55**, 124041 (2013).
- ¹²J. Vlcek, “A collisional-radiative model applicable to argon discharges over a wide range of conditions. I. Formulation and basic data,” *J. Phys. D: Appl. Phys.* **22**, 623 (1989).
- ¹³A. Bultel, B. van Ootegem, A. Bourdon, and P. Vervisch, “Influence of Ar_2^+ in an argon collisional-radiative model,” *Phys. Rev. E* **65**, 046406 (2002).
- ¹⁴M. G. Kapper and J. L. Cambier, “Ionizing shocks in argon. Part I: Collisional-radiative model and steady-state structure,” *J. Appl. Phys.* **109**, 113308 (2011).
- ¹⁵M. G. Kapper and J. L. Cambier, “Ionizing shocks in argon. Part II: Transient and multi-dimensional effects,” *J. Appl. Phys.* **109**, 113309 (2011).
- ¹⁶A. Bultel, B. G. Chéron, A. Bourdon, O. Motapon, and I. F. Schneider, “Collisional-radiative model in air for earth re-entry problems,” *Phys. Plasmas* **13**, 043502 (2006).
- ¹⁷A. Bultel and J. Annaloro, “Elaboration of collisional-radiative models for flows related to planetary entries into the Earth and Mars atmospheres,” *Plasma Sources Sci. Technol.* **22**, 025008 (2013).
- ¹⁸M. Panesi and A. Lani, “Collisional radiative coarse-grain model for ionization in air,” *Phys. Fluids* **25**, 057101 (2013).
- ¹⁹M. Capitelli, G. Colonna, G. D’Ammando, K. Hassouni, A. Laricchiuta, and L. D. Pietanza, “Coupling of plasma chemistry, vibrational kinetics, collisional-radiative models and electron energy distribution function under non-equilibrium conditions,” *Plasma Processes Polym.* **14**, 1600109 (2017).
- ²⁰E. A. Nagnibeda and E. Kustova, *Non-Equilibrium Reacting Gas Flows, Heat and Mass Transfer* (Springer, Berlin, Heidelberg, 2009).
- ²¹M. Capitelli, R. Celiberto, G. Colonna, F. Esposito, C. Gorse, K. Hassouni, A. Laricchiuta, and S. Longo, *Fundamental Aspects of Plasma Chemical Physics* (Springer, New York, 2016).
- ²²G. D. Billing and E. R. Fisher, “VV and VT rate coefficients in N_2 by a quantum-classical model,” *Chem. Phys.* **43**, 395–401 (1979).
- ²³F. Esposito, I. Armenise, and M. Capitelli, “N- N_2 state to state vibrational-relaxation and dissociation rates based on quasiclassical calculations,” *Chem. Phys.* **331**, 1–8 (2006).

- ²⁴M. Capitelli, I. Armenise, D. Bruno, M. Cacciatore, R. Celiberto, G. Colonna, O. De Pascale, P. Diomede, F. Esposito, C. Gorse, K. Hassouni, A. Laricchiuta, S. Longo, D. Pagano, D. Pietanza, and M. Rutigliano, "Non-equilibrium plasma kinetics: A state-to-state approach," *Plasma Sources Sci. Technol.* **16**, S30 (2007).
- ²⁵M. Panesi, T. E. Magin, A. Bourdon, A. Bultel, and O. Chazot, "Electronic excitation of atoms and molecules for the FIRE II flight experiment," *J. Thermophys. Heat Transfer* **25**, 361–374 (2011).
- ²⁶M. Capitelli, I. Armenise, and C. Gorse, "State-to-state approach in the kinetics of air components under re-entry conditions," *J. Thermophys. Heat Transfer* **11**, 570–578 (1997).
- ²⁷C. O. Laux, L. Pierrot, and R. J. Gessman, "State-to-state modeling of a recombining nitrogen plasma experiment," *Chem. Phys.* **398**, 46–55 (2012).
- ²⁸M. Panesi, R. L. Jaffe, D. W. Schwenke, and T. E. Magin, "Rovibrational internal energy transfer and dissociation of $N_2(^1\Sigma_g^+)-N(^4S_u)$ system in hypersonic flows," *J. Chem. Phys.* **138**, 044312 (2013).
- ²⁹M. Panesi, A. Munafò, T. E. Magin, and R. L. Jaffe, "Nonequilibrium shock-heated nitrogen flows using a rovibrational state-to-state method," *Phys. Rev. E* **90**, 013009 (2014).
- ³⁰P. Hammerling, J. D. Teare, and B. Kivel, "Theory of radiation from luminous shock waves in nitrogen," *Phys. Fluids* **2**, 422–426 (1959).
- ³¹S. P. Heims, "Moment equations for vibrational relaxation coupled with dissociation," *J. Chem. Phys.* **38**, 603–606 (1963).
- ³²P. V. Marrone and C. E. Treanor, "Chemical relaxation with preferential dissociation from excited vibrational levels," *Phys. Fluids* **6**, 1215–1221 (1963).
- ³³C. E. Treanor and P. V. Marrone, "Effect of dissociation on the rate of vibrational relaxation," *Phys. Fluids* **5**, 1022–1026 (1962).
- ³⁴C. Park, "Assessment of two-temperature kinetic model for ionizing air," *J. Thermophys. Heat Transfer* **3**, 233–244 (1989).
- ³⁵C. Park, "Review of chemical-kinetic problems of future NASA missions. I—Earth entries," *J. Thermophys. Heat Transfer* **7**, 385–398 (1993).
- ³⁶C. Park, J. T. Howe, R. L. Jaffe, and G. V. Candler, "Review of chemical-kinetic problems of future NASA missions. II—Mars entries," *J. Thermophys. Heat Transfer* **8**, 9–23 (1994).
- ³⁷C. Park, R. L. Jaffe, and H. Partridge, "Chemical-kinetic parameters of hyperbolic earth entry," *J. Thermophys. Heat Transfer* **15**, 76–90 (2001).
- ³⁸K. L. Heritier, R. L. Jaffe, V. Laporta, and M. Panesi, "Energy transfer models in nitrogen plasmas: Analysis of $N_2(^1\Sigma_g^+)-N(^4S_u)-e^-$ interaction," *J. Chem. Phys.* **141**, 184302 (2014).
- ³⁹G. Colonna, I. Armenise, D. Bruno, and M. Capitelli, "Reduction of state-to-state kinetics to macroscopic models in hypersonic flows," *J. Thermophys. Heat Transfer* **20**, 477–486 (2006).
- ⁴⁰G. Colonna, L. D. Pietanza, and M. Capitelli, "Reduced two-level approach for air kinetics in recombination regime," *AIP Conf. Proc.* **1333**, 1365–1370 (2011).
- ⁴¹T. E. Magin, M. Panesi, A. Bourdon, R. L. Jaffe, and D. W. Schwenke, "Coarse-grain model for internal energy excitation and dissociation of molecular nitrogen," *Chem. Phys.* **398**, 90–95 (2012).
- ⁴²H. P. Le, A. R. Karagozian, and J. L. Cambier, "Complexity reduction of collisional-radiative kinetics for atomic plasma," *Phys. Plasmas* **20**, 123304 (2013).
- ⁴³A. Munafò, M. Panesi, and T. E. Magin, "Boltzmann rovibrational collisional coarse-grained model for internal energy excitation and dissociation in hypersonic flows," *Phys. Rev. E* **89**, 023001 (2014).
- ⁴⁴Y. Liu, M. Panesi, A. Sahai, and M. Vinokur, "General multi-group macroscopic modeling for thermo-chemical non-equilibrium gas mixtures," *J. Chem. Phys.* **142**, 134109 (2015).
- ⁴⁵A. Sahai, B. Lopez, C. O. Johnston, and M. Panesi, "Adaptive coarse graining method for energy transfer and dissociation kinetics of polyatomic species," *J. Chem. Phys.* **147**, 054107 (2017).
- ⁴⁶S. Venturi, M. Sharma Priyadarshini, B. Lopez, and M. Panesi, "Data-inspired and physics-driven model reduction for dissociation: Application to the O_2+O system," *J. Phys. Chem. A* **124**, 8359–8372 (2020).
- ⁴⁷M. Sharma Priyadarshini, Y. Liu, and M. Panesi, "Coarse-grained modeling of thermochemical nonequilibrium using the multigroup maximum entropy quadratic formulation," *Phys. Rev. E* **101**, 013307 (2020).
- ⁴⁸R. J. Abrantes, Sousa, D. Bilyeu, and R. Martin, "Complexity-reduction using automatic level grouping for atomic collisional-radiative models," *J. Comput. Phys.* **407**, 109213 (2020).
- ⁴⁹I. Zanardi, S. Venturi, and M. Panesi, "Adaptive physics-informed neural operator for coarse-grained non-equilibrium flows," *Sci. Rep.* **13**, 1–22 (2023).
- ⁵⁰C. Jacobsen, I. Zanardi, S. Bhola, K. Duraisamy, and M. Panesi, "Information theoretic clustering for coarse-grained modeling of non-equilibrium gas dynamics," *J. Comput. Phys.* **507**, 112977 (2024).
- ⁵¹Y. Huang, Q. Hong, S. Gu, X. Wang, and Q. Sun, "Coarse-grained modeling of high-enthalpy air flows based on the updated vibrational state-to-state kinetics," *Phys. Fluids* **36**, 106133 (2024).
- ⁵²C. W. Rowley, T. Colonius, and R. M. Murray, "Model reduction for compressible flows using POD and Galerkin projection," *Physica D* **189**, 115–129 (2004).
- ⁵³M. F. Barone, I. Kalashnikova, D. J. Segalman, and H. K. Thornquist, "Stable Galerkin reduced order models for linearized compressible flow," *J. Comput. Phys.* **228**, 1932–1946 (2009).
- ⁵⁴I. C. Sutherland and A. Parente, "Combustion modeling using principal component analysis," *Proc. Combust. Inst.* **32**, 1563–1570 (2009).
- ⁵⁵A. Parente and J. C. Sutherland, "Principal component analysis of turbulent combustion data: Data pre-processing and manifold sensitivity," *Combust. Flame* **160**, 340–350 (2013).
- ⁵⁶B. J. Isaac, A. Coussement, O. Gicquel, P. J. Smith, and A. Parente, "Reduced-order PCA models for chemical reacting flows," *Combust. Flame* **161**, 2785–2800 (2014).
- ⁵⁷M. R. Malik, B. J. Isaac, A. Coussement, P. J. Smith, and A. Parente, "Principal component analysis coupled with nonlinear regression for chemistry reduction," *Combust. Flame* **187**, 30–41 (2018).
- ⁵⁸A. Bellemans, G. Aversano, A. Coussement, and A. Parente, "Feature extraction and reduced-order modelling of nitrogen plasma models using principal component analysis," *Computers Chem. Eng.* **115**, 504–514 (2018).
- ⁵⁹A. Bellemans, A. Parente, and T. Magin, "Principal component analysis acceleration of rovibrational coarse-grain models for internal energy excitation and dissociation," *J. Chem. Phys.* **148**, 164107 (2018).
- ⁶⁰A. Bellemans, A. Munafò, T. E. Magin, G. Degrez, and A. Parente, "Reduction of a collisional-radiative mechanism for argon plasma based on principal component analysis," *Phys. Plasmas* **22**, 062108 (2015).
- ⁶¹A. Bellemans, T. Magin, A. Coussement, and A. Parente, "Reduced-order kinetic plasma models using principal component analysis: Model formulation and manifold sensitivity," *Phys. Rev. Fluids* **2**, 073201 (2017).
- ⁶²C. W. Rowley and S. T. Dawson, "Model reduction for flow analysis and control," *Annu. Rev. Fluid Mech.* **49**, 387–417 (2017).
- ⁶³S. E. Otto, A. Padovan, and C. W. Rowley, "Optimizing oblique projections for nonlinear systems using trajectories," *SIAM J. Sci. Comput.* **44**, A1681–A1702 (2022).
- ⁶⁴S. E. Otto, A. Padovan, and C. W. Rowley, "Model reduction for nonlinear systems by balanced truncation of state and gradient covariance," *SIAM J. Sci. Comput.* **45**, A2325–A2355 (2023).
- ⁶⁵A. Padovan, B. Vollmer, and D. J. Bodony, "Data-driven model reduction via non-intrusive optimization of projection operators and reduced-order dynamics," *SIAM J. Appl. Dyn. Syst.* **23**, 3052–3076 (2024).
- ⁶⁶I. Zanardi, A. Padovan, D. J. Bodony, and M. Panesi, "Petrov-Galerkin model reduction for thermochemical nonequilibrium gas mixtures," *J. Comput. Phys.* **533**, 113999 (2025).
- ⁶⁷I. Zanardi, A. Padovan, D. J. Bodony, and M. Panesi, "Petrov-Galerkin model reduction for thermochemical nonequilibrium gas mixtures: Application to the N_2+N system," in *AIAA SciTech Forum and Exposition* (American Institute of Aeronautics and Astronautics Inc., 2025).
- ⁶⁸B. C. Moore, "Principal component analysis in linear systems: controllability, observability, and model reduction," *IEEE Trans. Autom. Control* **26**, 17–32 (1981).
- ⁶⁹G. E. Dullerud and F. Paganini, *A Course in Robust Control Theory*, Vol. 36 (Springer, New York, 2000).
- ⁷⁰A. Varga, "Balanced truncation model reduction of periodic systems," in *Proceedings of the IEEE Conference on Decision and Control* (IEEE, 2000), pp. 2379–2384.
- ⁷¹C. W. Rowley, "Model reduction for fluids using balanced proper orthogonal decomposition," *Int. J. Bifurcation Chaos* **15**, 997–1013 (2005).
- ⁷²A. Padovan and C. W. Rowley, "Continuous-time balanced truncation for time-periodic fluid flows using frequential Gramians," *J. Comput. Phys.* **496**, 112597 (2024).
- ⁷³I. I. Glass, "Over forty years of continuous research at UTIAS on nonstationary flows and shock waves," *Shock Waves* **1**, 75–86 (1991).

- ⁷⁴Y. B. Zel'dovich, Y. P. Raizer, W. D. Hayes, and R. F. Probst, *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena* (Academic Press, 1967).
- ⁷⁵C. Hirsch, *Numerical Computation of Internal and External Flows*, 1st ed. (Elsevier, 2007).
- ⁷⁶B. van Leer, "Flux-vector splitting for the Euler equations," *Lect. Notes Phys.* **170**, 507–512 (1982).
- ⁷⁷B. van Leer, "Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method," *J. Comput. Phys.* **32**, 101–136 (1979).
- ⁷⁸G. D. van Albada, B. van Leer, and W. W. Roberts, "A comparative study of computational methods in cosmic gas dynamics," in *Upwind and High-Resolution Schemes* (Springer, Berlin, Heidelberg, 1997), pp. 95–103.
- ⁷⁹J. J. Quirk, "A contribution to the great Riemann solver debate," *Int. J. Numer. Methods Fluids* **18**, 555–574 (1994).
- ⁸⁰G. Strang, "On the construction and comparison of difference schemes," *SIAM J. Numer. Anal.* **5**, 506–517 (1968).
- ⁸¹O. M. Knio, H. N. Najm, and P. S. Wyckoff, "A semi-implicit numerical scheme for reacting flow," *J. Comput. Phys.* **154**, 428–467 (1999).
- ⁸²M. A. Singer, S. B. Pope, and H. N. Najm, "Operator-splitting with ISAT to model reacting flow with detailed chemistry," *Combust. Theor. Model.* **10**, 199–217 (2006).
- ⁸³Z. Ren, C. Xu, T. Lu, and M. A. Singer, "Dynamic adaptive chemistry with operator splitting schemes for reactive flow simulations," *J. Comput. Phys.* **263**, 19–36 (2014).
- ⁸⁴H. Wu, P. C. Ma, and M. Ihme, "Efficient time-stepping techniques for simulating turbulent reactive flows with stiff chemistry," *Comput. Phys. Commun.* **243**, 81–96 (2019).
- ⁸⁵D. R. Durran, *Numerical Methods for Fluid Dynamics*, 2nd ed., Vol. 32 (Springer, New York, 2010).
- ⁸⁶K. Radhakrishnan and A. C. Hindmarsh, "Description and use of LSODE, the Livermore solver for ordinary differential equations," Tech. Rep. 113855 (1993).
- ⁸⁷I. Noda, "Scaling techniques to enhance two-dimensional correlation spectra," *J. Mol. Struct.* **883–884**, 216–227 (2008).
- ⁸⁸B. J. Isaac, "Reduced-order modeling for reacting flows based on principal component analysis," Ph.D. thesis (University of Utah, Salt Lake City, 2014).
- ⁸⁹O. Zahm, P. G. Constantine, C. Prieur, and Y. M. Marzouk, "Gradient-based dimension reduction of multivariate vector-valued functions," *SIAM J. Sci. Comput.* **42**, A534–A558 (2020).
- ⁹⁰A. Munafò, S. Kumar, S. M. Jo, and M. Panesi, "HEGEL: A high-fidelity flexible software for hypersonics and plasma simulations," in *AIAA SciTech Forum and Exposition* (American Institute of Aeronautics and Astronautics Inc., 2024).
- ⁹¹A. Munafò and M. Panesi, "Plato: A high-fidelity library for multicomponent gases and plasmas," *J. Thermophys. Heat Transfer* **39**, 850–821 (2025).