Adjoint-based calibration of nonlinear coupled constitutive relations for argon normal shocks with a differentiable finite-volume solver

Qirong Tu [1], Mengjia Wu [2] and Gaoxiang Xiang [2, *]

[1] School of Aeronautics, Northwestern Polytechnical University, Xi'an 710072, People's Republic of China

[2] School of Mechanics and Transportation Engineering, Northwestern Polytechnical University, Xi'an 710129, People's Republic of China

* Corresponding Author Email: xianggx@nwpu.edu.cn

Abstract. Nonlinear coupled constitutive relations (NCCR) extend the Navier–Stokes–Fourier equations through nonlinear dissipation and stress–stress and heat-flux–stress couplings. We calibrate their three constants for argon normal shocks using an implicit constitutive layer, a differentiable finite-volume solver and a discrete adjoint. A phase condition removes the translational mode of the steady shock. With the viscosity law and Prandtl number fixed, joint calibration at Mach 2.05–8 gives $\boldsymbol{\theta}^* = (0,\ 0.775,\ 2.258)$ and reduces the density–temperature misfit by 86.8%. The Mach-8 heat-flux peak underprediction falls from 21.4% to 6.7%. The enhanced heat-flux–stress coupling persists after recalibration with each training Mach number omitted in turn and with independent DS1V data. Strong shocks drive the dissipation scale to zero, whereas weak shocks constrain it poorly. The upstream temperature precursor remains too short, limiting extrapolation beyond the fitted one-dimensional compression flows.

Keywords: Nonlinear coupled constitutive relations; argon normal shocks; parameter calibration; discrete adjoint; differentiable finite-volume solver.

1. Introduction

Normal-shock structure is a stringent test of continuum descriptions of rarefied gases: density and temperature vary over only a few mean free paths, allowing comparison with measurements, [1] direct simulation Monte Carlo (DSMC) [2] and ab initio simulations. [3] Within the Navier–Stokes–Fourier (NSF) framework, transport corrections such as bulk viscosity address specific nonequilibrium effects, [4] but constitutive descriptions beyond linear stress and heat-flux laws are needed as gradients increase.

Eu's generalized hydrodynamics [5] and Myong's NCCR closure [6,7] retain nonlinear dissipation and coupling between thermodynamic fluxes. Recent work examines shock relaxation,[8] hypersonic configurations [9] and vibrational nonequilibrium in nccrFOAM.[10] Developments in unstructured-grid solution strategies,[11] multiple-solution rectification [12] and three-temperature polyatomic relaxation [13] further emphasize the importance of constitutive assumptions and reliable nonlinear solvers. These studies motivate testing the constants of an established algebraic closure against kinetic reference data.

Data-driven constitutive modeling includes shock-flow learning, [14] field inversion,[15] embedded closure training, [16] and optimization of first- and second-order transport models. [17, 18] Differentiable flow solvers [19, 20] provide a route to sensitivity through the governing equations. Here we identify only the three NCCR constants, preserving the model structure, viscosity law and Prandtl number to isolate nonlinear dissipation and stress coupling. The objective uses density and temperature; stress and heat flux provide conservation-consistency checks. We assess robustness using the Mach-1.55 case excluded from calibration, independent DSMC data and grid refinement.

2. Governing equations and constitutive model

For monatomic gas with gas constant R, $\gamma = 5/3$ and $c_p = \gamma R/(\gamma - 1)$, the steady one-dimensional conservation laws are

$$\frac{\mathrm{d}}{\mathrm{d}x}(\rho u) = 0, \quad \frac{\mathrm{d}}{\mathrm{d}x}(\rho u^2 + p + \Pi_{xx}) = 0, \quad \frac{\mathrm{d}}{\mathrm{d}x}[u(\rho E + p) + \Pi_{xx}u + Q_x] = 0, \tag{1}$$

Here Π_{xx} and Q_x are normal traceless stress and heat flux, $E = c_v T + u^2/2$ and $p = \rho RT$. We prescribe $\mu = \mu_{\mathrm{ref}}(T/T_{\mathrm{ref}})^{\omega}$ with $\omega = 0.72$, $k = \mu c_p/\mathrm{Pr}$ and $\mathrm{Pr} = 2/3$. Length is normalized by $\lambda_\infty = (\mu_1/p_1)\sqrt{\pi RT_1/2}$, and density and temperature by $\hat{\rho} = (\rho - \rho_1)/(\rho_2 - \rho_1)$ and $\hat{T} = (T - T_1)/(T_2 - T_1)$, where subscripts 1 and 2 denote upstream and downstream Rankine–Hugoniot states. Stress and heat flux are normalized by $\rho_1 u_1^2$ and $\rho_1 u_1^3$, respectively.

The dimensionless stress, heat flux and imposed gradients are

$$a = \frac{\Pi_{xx}}{p}, \quad g = \frac{c_Q Q_x}{p\sqrt{T}}, \quad c_Q = \sqrt{\frac{2\mathrm{Pr}}{c_p}}, \qquad A = -\frac{4}{3}\frac{\mu}{p}\frac{\mathrm{d}u}{\mathrm{d}x}, \quad G = -\frac{c_Q k}{p\sqrt{T}}\frac{\mathrm{d}T}{\mathrm{d}x}, \tag{2}$$

The one-dimensional NCCR closure is

$$a\,q(\kappa) = A(1 + \alpha_{\Pi\Pi}a), \qquad g\,q(\kappa) = G(1 + \alpha_{Q\Pi}a), \tag{3}$$

With

$$q(\kappa) = \frac{\sinh\kappa}{\kappa}, \qquad \kappa = c\sqrt{\frac{3}{2}a^2 + g^2} \tag{4}$$

And $\Pi_{yy} = \Pi_{zz} = -\Pi_{xx}/2$. The scale constant controls nonlinear dissipation through $q(\kappa)$, reducing the effective viscosity and conductivity to μ/q and k/q. The theoretical scale is 1.0138 for Maxwell molecules; a viscosity exponent of 0.72 gives approximately 1.02. Both coupling constants are unity in theoretical NCCR. Setting $c = 0$ gives $q \equiv 1$ and

$$a = \frac{A}{1 - \alpha_{\Pi\Pi}A}, \qquad g = G(1 + \alpha_{Q\Pi}a) \tag{5}$$

Both couplings remain active; NSF is recovered only for small A and G. The stress–stress coefficient controls stress amplification in compression, whereas the heat-flux–stress coefficient enhances heat transfer in the stressed gas. We calibrate $\boldsymbol{\theta} = (c, \alpha_{\Pi\Pi}, \alpha_{Q\Pi})$ within $c \in [0,1.08]$, $\alpha_{\Pi\Pi} \in [0,3]$ and $\alpha_{Q\Pi} \in [0,5]$. The theoretical baseline is $\boldsymbol{\theta} = (1.0138,1,1)$.

3. Differentiable solver and discrete adjoint

We use $N = 200$ uniform finite-volume cells over $L = 60\lambda_\infty$, Rankine–Hugoniot ghost states, unlimited MUSCL reconstruction and HLLC fluxes in double-precision PyTorch. The constitutive state $\boldsymbol{y} = (a, g)$ is solved at each face by damped Newton iteration to residual 10^{-14}. Treating the converged closure $\boldsymbol{G}(\boldsymbol{y}^*, \boldsymbol{x}, \boldsymbol{\theta}) = 0$ as an implicit layer, with $\boldsymbol{x} = (A, G)$, gives

$$\frac{\partial \boldsymbol{y}^*}{\partial(\boldsymbol{x}, \boldsymbol{\theta})} = -\left(\frac{\partial \boldsymbol{G}}{\partial \boldsymbol{y}}\right)^{-1}\frac{\partial \boldsymbol{G}}{\partial(\boldsymbol{x}, \boldsymbol{\theta})} \tag{6}$$

Reverse-mode differentiation requires one transposed 2×2 solve per face, with memory independent of the Newton iteration count, following modular implicit differentiation. [21]

3.1. Phase condition and objective

Translation invariance gives the steady residual $\boldsymbol{R}(\boldsymbol{U})$ a near-null mode $\partial \boldsymbol{U}/\partial x$. The unmodified Jacobian has condition number of order 10^{10}, stalls Newton convergence and produces

percent-level gradient errors. Following the freezing approach,[22] we introduce a shift variable s and impose $\varphi(\boldsymbol{U}) = \sum_i (\rho_i - \rho_i^{\mathrm{ref}})\, n_i/(\rho_2 - \rho_1) = 0$, where $n_i \propto \mathrm{d}\rho^{\mathrm{ref}}/\mathrm{d}x$ has unit length and ρ^{ref} is the DSMC density profile. The bordered Newton system is

$$\begin{pmatrix} \partial_{\boldsymbol{U}}\boldsymbol{R} & \partial_s \boldsymbol{R} \\ \partial_{\boldsymbol{U}}\varphi & 0 \end{pmatrix} \begin{pmatrix} \delta \boldsymbol{U} \\ \delta s \end{pmatrix} = -\begin{pmatrix} \boldsymbol{R} \\ \varphi \end{pmatrix} \tag{7}$$

Its condition number is of order 10^3; Newton converges quadratically and the profiles are registered consistently. We minimize the density–temperature discrepancy over the training Mach numbers:

$$J(\boldsymbol{\theta}) = \sum_{M_\infty \in \mathcal{M}} \frac{1}{N} \sum_{i=1}^{N} \left[(\hat{\rho}_i - \hat{\rho}_i^{\mathrm{D}})^2 + (\hat{T}_i - \hat{T}_i^{\mathrm{D}})^2 \right] \tag{8}$$

Stress and heat flux remain predicted quantities. At the converged state $\boldsymbol{U}^*$, the discrete adjoint gives

$$\frac{\mathrm{d}J}{\mathrm{d}\boldsymbol{\theta}} = \frac{\partial J}{\partial \boldsymbol{\theta}} + \hat{\boldsymbol{u}}^{\mathsf{T}} \frac{\partial \boldsymbol{R}}{\partial \boldsymbol{\theta}}, \qquad \left(\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{U}}\right)^{\mathsf{T}} \hat{\boldsymbol{u}} = -\left(\frac{\partial J}{\partial \boldsymbol{U}}\right)^{\mathsf{T}} \tag{9}$$

The bordered operator in equation (7) replaces $\partial \boldsymbol{R}/\,\partial \boldsymbol{U}$. L-BFGS-B typically converges in fewer than 20 evaluations, each using one forward and one adjoint solve per Mach number. Gradients agree with finite differences to 10^{-10} and with unrolled Newton differentiation to 10^{-15}.

3.2. Reference DSMC data

SPARTA [23] supplies one-dimensional argon shocks with a variable-hard-sphere model,[24] $\omega = 0.72$, reference viscosity 2.116×10^{-5} Pa s and upstream temperature 273.15 K. Mach numbers 2.05, 3.2, 4, 6 and 8 are used for training; 1.55 is held out. Figure 1 compares inverse density thickness with experiment [1] and published simulations. [25, 3] Here $\delta = (\rho_2 - \rho_1)/\max|\mathrm{d}\rho/\mathrm{d}x|$ and $\lambda_1 = 1.0186\lambda_\infty$. At Mach 3.2–8, present DSMC shocks are 1–8% thinner than experiment and within −4% to +2% of ab initio results. At Mach 2.05 and below, the experimental discrepancy grows to 9–18%; the ab initio trend and experimental spread near Mach 2 motivate retaining the Mach-2.05 case.

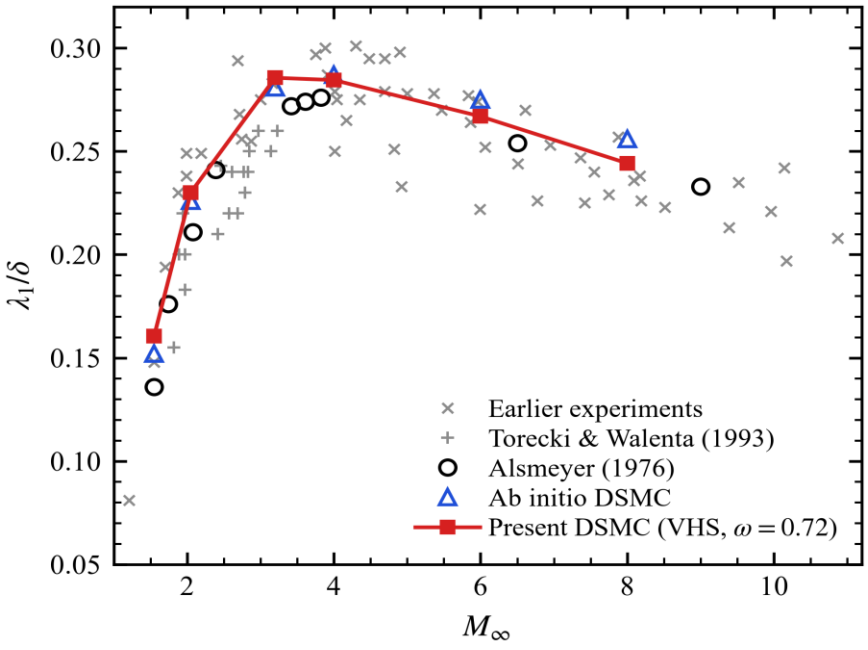


Figure 1. Inverse density thickness λ_1/δ: present DSMC (filled squares), Alsmeyer (circles), earlier experiments (crosses), Torecki and Walenta (plus signs), and ab initio DSMC (triangles).

4. Results and discussion

Joint calibration yields $\boldsymbol{\theta}^* = (0, 0.775, 2.258)$. The objective falls from 1.59×10^{-2} to 2.11×10^{-3} (86.8%), with individual reductions of 6% at Mach 2.05 and 72–91% at Mach 3.2–8. Figures 2–4 show the profile comparisons. Density and temperature errors decrease throughout the training set; temperature errors are approximately halved at Mach 3.2 and above. At the held-out Mach number 1.55, density and temperature errors are no larger than the theoretical baseline (Table1).

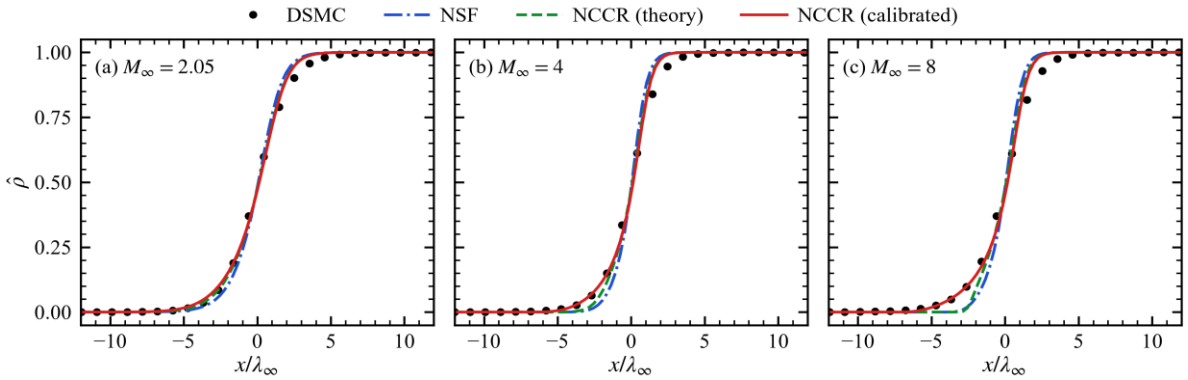


Figure 2. Normalized density $\hat{\rho}$ at Mach (a) 2.05, (b) 4 and (c) 8. Symbols: DSMC; dash-dotted: NSF; dashed: theoretical NCCR; solid: calibrated NCCR.

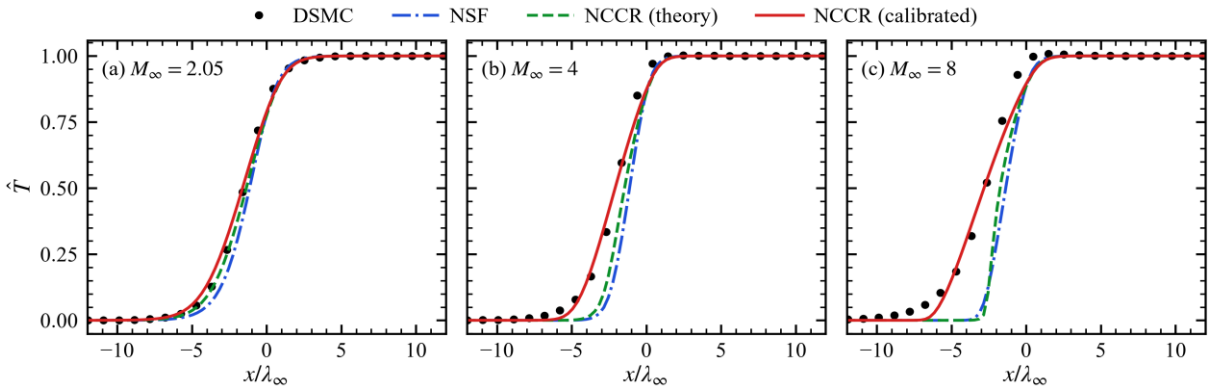


Figure 3. Normalized temperature $\hat{T}$; panels and line conventions as in Figure 2.

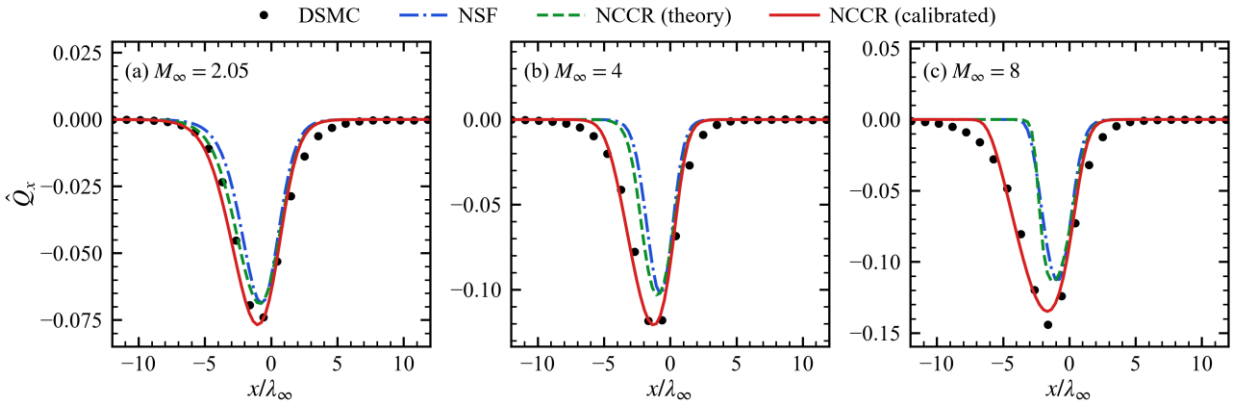


Figure 4. Predicted normalized heat flux $\hat{Q}_x$; panels and line conventions as in Figure 2.

Heat-flux peak errors improve from 9–21% underprediction to within ±7%; at Mach 8, the error changes from −21.4% to −6.7% and whole-domain heat-flux RMS error is one quarter of the NSF value. Stress responds less: its Mach-8 peak is 0.146 in DSMC, 0.185 in theoretical NCCR and 0.172 after calibration. Stress and heat flux are not fitted directly, but the steady conservation equations link them to the fitted density and temperature profiles. Their agreement with DSMC therefore provides a consistency check, rather than a fully independent validation of the model.

4.1. Error metrics and scale identifiability

Table 1. RMS errors relative to DSMC, as percentages of each quantity's DSMC peak, in the shock core $|x| < 2\lambda_\infty$; w.d. denotes whole-domain heat flux. N, T and C denote NSF, theoretical NCCR and calibrated NCCR. Mach 1.55 is held out.

M_∞	$\hat{\rho}$ N	$\hat{\rho}$ T	$\hat{\rho}$ C	$\hat{T}$ N	$\hat{T}$ T	$\hat{T}$ C	$\hat{Q}$ N	$\hat{Q}$ T	$\hat{Q}$ C	$\hat{Q}$ w.d. N	$\hat{Q}$ w.d. T	$\hat{Q}$ w.d. C
1.55	2.2	1.4	1.4	2.8	2.0	1.2	5.6	3.8	2.5	3.1	2.4	3.0
2.05	4.7	2.5	2.3	6.3	3.7	2.0	12.7	8.4	5.1	5.0	2.9	2.9
3.2	7.2	4.0	3.6	13.3	7.8	3.7	20.6	14.0	6.8	8.5	5.7	2.6
4.0	8.3	4.9	4.3	16.0	9.9	4.4	22.9	16.0	7.5	10.1	7.8	2.7
6.0	9.3	5.7	4.8	18.6	12.5	5.8	23.9	18.1	8.4	12.3	11.4	3.1
8.0	9.8	6.1	5.0	18.8	12.8	6.5	22.8	17.8	8.9	13.7	13.4	3.4

The inverse density thickness remains 10–41% above DSMC over Mach 2.05–8 and changes by only 0.8–1.5% after calibration. Temperature–density separation improves more strongly: at Mach 8, an underprediction of $0.92\lambda_\infty$ becomes an overprediction of $0.26\lambda_\infty$.

We varied c and re-optimized both coupling coefficients at each value to minimize the density–temperature fitting error (Figure 5). For Mach numbers up to 3.2, this minimum error changes by only about 10% as c increases from 0.2 to 1.08. Thus, after re-optimizing the couplings, the fitting error is insensitive to c, making it difficult to determine c from these cases alone. For the full training set, increasing c from 0 to the Maxwell-molecule value of 1.0138 raises the minimum error to 1.84 times its value at $c = 0$. Thus, including the higher-Mach-number cases favors $c = 0$ for the overall fit.

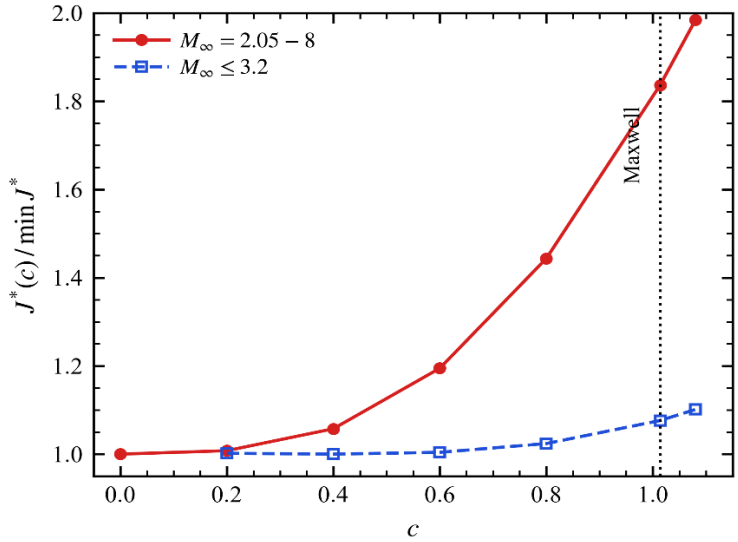


Figure 5. Minimum objective at fixed dissipation scale, normalized by its minimum, for the full Mach-2.05–8 set and the low-Mach-number set (Mach ≤ 3.2). Dotted line: theoretical scale 1.0138. Here $J^*(c) = \min_{\alpha_{\Pi\Pi}, \alpha_{Q\Pi}} J$.

This result can be understood from the dissipation factor q. The effective viscosity and thermal conductivity are μ/q and k/q, so larger q suppresses viscous and heat transport more strongly. In theoretical NCCR, q stays below 1.05 at Mach numbers up to 2.05, reaches 1.27 at Mach 3.2 and exceeds 7 on the upstream side of the Mach-8 shock. Its effect is therefore small at lower Mach numbers and substantial at higher Mach numbers. Together with Figure 5, these results suggest that the theoretical dissipation factor may suppress viscous and heat transport too strongly in the higher-Mach-number cases studied here. Setting $c = 0$ gives $q = 1$, removing this additional suppression.

As c increases, the re-optimized stress–stress coupling coefficient also increases. The heat-flux–stress coupling coefficient remains between 2.26 and 2.97, above its theoretical value of 1. Thus, calibration gives stronger heat-flux–stress coupling than the theoretical model throughout the range of c considered.

4.2. Robustness and limitations

To assess sensitivity to the calibration set, we omitted all data for one training Mach number at a time and recalibrated the three parameters using the remaining four cases. All five recalibrations reached the lower bound c = 0, with stress–stress coupling 0.73–0.80 and heat-flux–stress coupling 2.18–2.36. Independent DS1V data with the same molecular model give (0.764, 2.245), within 1.4% and 0.5% of the SPARTA fit. Refining from 200 to 400 cells changes the couplings by +1.8% and −1.3%; 60 cells instead give (0.18, 0.769, 2.545). A free Prandtl number gives Pr = 0.575 and heat-flux–stress coupling 1.53, reducing the objective by only 1.6%. The upstream temperature precursor remains too short, and one-dimensional compression does not constrain shear or expansion.

5. Conclusion

A phase-fixed implicit adjoint enables accurate calibration of the three NCCR constants. At fixed transport properties, enhanced heat-flux–stress coupling improves shock-core temperature and heat flux across Mach 2.05–8. Strong shocks favor removal of nonlinear dissipation, while weak shocks constrain its scale poorly. The enhanced coupling persists when each training Mach number is omitted in turn and the parameters are recalibrated, and when independent DSMC data are used. These conclusions remain specific to argon normal shocks.

Acknowledgement

This work was supported by the National Natural Science Foundation of China Basic Science Center Program (No. 12388101), the National Natural Science Foundation of China (12202365), the Guangdong Basic and Applied Basic Research Foundation (2022A1515011565, 2023A1515010031), the Natural Science Foundation of Chongqing (2022NSCQ-MSX5709), Northwestern Polytechnical University's Future Industry Proof-of-Concept Project for Aerospace Propulsion (24GNYZ0001-04), and the China Postdoctoral Science Foundation funded project (2021M692633, 2022T150534).

References

[1] Alsmeyer H (1976) Density profiles in argon and nitrogen shock waves measured by the absorption of an electron beam. J. Fluid Mech., 74(3): 497–513. DOI: 10.1017/S0022112076001912

[2] Bird G A (1994) Molecular Gas Dynamics and the Direct Simulation of Gas Flows. Clarendon Press, Oxford. DOI: 10.1093/oso/9780198561958.001.0001

[3] Sharipov F and Dias F C (2017) Ab initio simulation of planar shock waves. Comput. Fluids, 150: 115–122. DOI: 10.1016/j.compfluid.2017.04.002

[4] Tu Q, Xiang G and Zheng X (2025) Bulk Viscosity Effects in Hypersonic Near-Continuum Flow. AIAA J., 63(2): 521–539. DOI: 10.2514/1. J064560

[5] Eu B C (1980) A modified moment method and irreversible thermodynamics. J. Chem. Phys., 73(6): 2958–2969. DOI: 10.1063/1.440469

[6] Myong R S (1999) Thermodynamically consistent hydrodynamic computational models for high-Knudsen-number gas flows. Phys. Fluids, 11(9): 2788–2802. DOI: 10.1063/1.870137

[7] Myong R S (2004) A generalized hydrodynamic computational model for rarefied and microscale diatomic gas flows. J. Comput. Phys., 195(2): 655–676. DOI: 10.1016/j.jcp.2003.10.015

[8] Zeng S, Yang J, Zhao W, Agarwal R K and Chen W (2025) Effect of bulk viscosity and relaxation model on nonequilibrium shock structure of diatomic gases. Phys. Rev. Fluids, 10(6): 063401. DOI: 10.1103/8mtq-n1dn

[9] Jiang Z, Zhao W, Yuan Z, Chen W and Myong R S (2019) Computation of Hypersonic Flows over Flying Configurations Using a Nonlinear Constitutive Model. AIAA J., 57(12): 5252–5268. DOI: 10.2514/1. J057688

[10] Mankodi T K, Ejtehadi O, Chourushi T, Rahimi A and Myong R S (2024) nccrFOAM suite: Nonlinear coupled constitutive relation solver in the OpenFOAM framework for rarefied and microscale gas flows with vibrational non-equilibrium. Comput. Phys. Commun., 296: 109024. DOI: 10.1016/j.cpc.2023.109024

[11] Zeng S, He Z, Zhao W and Chen W (2023) A computational strategy for nonlinear coupled constitutive relations of rarefied nonequilibrium flows on unstructured grids. Aerosp. Sci. Technol., 141: 108477. DOI: 10.1016/j.ast.2023.108477

[12] Cao J, Liu S, Zhong C, Zhuo C and Xu K (2024) Multiple solutions of nonlinear coupled constitutive relation model and its rectification in non-equilibrium flow computation. Comput. Math. Appl., 174: 1–17. DOI: 10.1016/j.camwa.2024.08.017

[13] Srivastava H, Mankodi T K and Myong R S (2026) Non-equilibrium polyatomic gas flows: Nonlinear coupled constitutive relations and topological study of three-temperature relaxation. Phys. Fluids, 38(7): 076132. DOI: 10.1063/5.0340960

[14] Garg G, Mankodi T K, Esmaeilifar E and Myong R S (2024) Neural network-based finite volume method and direct simulation Monte Carlo solutions of non-equilibrium shock flow guided by nonlinear coupled constitutive relations. Phys. Fluids, 36(10): 106113. DOI: 10.1063/5.0223654

[15] Parish E J and Duraisamy K (2016) A paradigm for data-driven predictive modeling using field inversion and machine learning. J. Comput. Phys., 305: 758–774. DOI: 10.1016/j.jcp.2015.11.012

[16] MacArt J F, Sirignano J and Freund J B (2021) Embedded training of neural-network subgrid-scale turbulence models. Phys. Rev. Fluids, 6(5): 050502. DOI: 10.1103/PhysRevFluids.6.050502

[17] Nair A S, Sirignano J, Panesi M and MacArt J F (2023) Deep Learning Closure of the Navier–Stokes Equations for Transition-Continuum Flows. AIAA J., 61(12): 5484–5497. DOI: 10.2514/1. J062935

[18] Kryger M and MacArt J F (2025) Optimization of Second-Order Transport Models for Transition-Continuum Flows. AIAA J., 63(10): 4223–4233. DOI: 10.2514/1. J064990

[19] Bezgin D A, Buhendwa A B and Adams N A (2023) JAX-Fluids: A fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. Comput. Phys. Commun., 282: 108527. DOI: 10.1016/j.cpc.2022.108527

[20] Xiao T (2025) Solving continuum and rarefied flows using differentiable programming. J. Comput. Phys., 539: 114224. DOI: 10.1016/j.jcp.2025.114224

[21] Blondel M, Berthet Q, Cuturi M, Frostig R, Hoyer S, Llinares-López F, Pedregosa F and Vert J-P (2022) Efficient and Modular Implicit Differentiation. Adv. Neural Inf. Process. Syst., 35: 5230–5242. DOI: 10.52202/068431-0378

[22] Beyn W-J and Thümmler V (2004) Freezing Solutions of Equivariant Evolution Equations. SIAM J. Appl. Dyn. Syst., 3(2): 85–116. DOI: 10.1137/030600515

[23] Plimpton S J, Moore S G, Borner A, Stagg A K, Koehler T P, Torczynski J R and Gallis M A (2019) Direct simulation Monte Carlo on petaflop supercomputers and beyond. Phys. Fluids, 31(8): 086101. DOI: 10.1063/1.5108534

[24] Uribe F J and Velasco R M (2018) Shock-wave structure based on the Navier-Stokes-Fourier equations. Phys. Rev. E, 97(4): 043117. DOI: 10.1103/PhysRevE.97.043117

[25] Greenshields C J and Reese J M (2007) The structure of shock waves as a test of Brenner's modifications to the Navier–Stokes equations. J. Fluid Mech., 580: 407–429. DOI: 10.1017/S0022112007005575