

# Deep Learning Closure of the Navier–Stokes Equations for Transition-Continuum Flows

Ashish S. Nair\*

*University of Notre Dame, Notre Dame, Indiana 46556, USA*

Justin Sirignano†

*University of Oxford, Oxford OX2 6GG, UK*

Marco Panesi‡

*University of Illinois at Urbana–Champaign, Urbana, Illinois 61801, USA*

Jonathan F. MacArt§

*University of Notre Dame, Notre Dame, Indiana 46556, USA*

The predictive accuracy of the Navier–Stokes equations is known to degrade at the limits of the continuum assumption, thereby necessitating expensive and approximate solutions to the Boltzmann equation. While tractable in one spatial dimension, their high dimensionality increases the computational challenge of multi-dimensional, multi-physical Boltzmann calculations. It is therefore desirable to augment the Navier–Stokes equations for accuracy under these conditions. We present an application of a deep learning method to extend the validity of the Navier–Stokes equations to transition-continuum flows. The technique encodes the missing physics via a neural network, which is trained to reduce the error between the Navier–Stokes and Boltzmann solutions. While standard DL methods can be considered *ad hoc* due to the absence of underlying physical laws, at least in the sense that the systems are not governed by known partial differential equations, the DL framework leverages the *a priori* known Boltzmann physics while ensuring that the trained model is consistent with the Navier–Stokes equations. The online training procedure solves adjoint equations, constructed using algorithmic differentiation, which efficiently provide the gradient of the loss function with respect to the learnable parameters. The model is trained and applied to predict stationary, one-dimensional shock thickness in low-pressure argon.

## I. Nomenclature

$\mathbf{U}$	=	dependent variables
$\hat{\mathbf{u}}$	=	adjoint variables
$\theta$	=	model parameters
$f(\mathbf{U}_x, \theta)$	=	neural network closure model

## II. Introduction

COMPUTATIONALLY efficient and accurate models for flows in translational nonequilibrium are essential for designing, operating, and developing hypersonic flight vehicles and their advanced propulsion systems. The Knudsen number ( $\text{Kn} = \lambda/L$ , where  $\lambda$  is the molecular mean free path and  $L$  is a characteristic length scale) indicates the flow regime, spanning continuum to free-molecule flow. Increasing  $\text{Kn}$  indicates an increasing departure from local thermodynamic

---

Aspects of this work were presented as Paper 2023-1796 at the 2023 AIAA SciTech Forum, National Harbor, MD, January 23–27, 2023, and Paper 2022-1703 at the 2022 AIAA SciTech Forum, San Diego, CA, January 3–7, 2022.

\*Graduate Student, Department of Aerospace and Mechanical Engineering

†Associate Professor, Mathematical Institute

‡Professor, Department of Aerospace Engineering; Director, Center for Hypersonic and Entry Systems Studies. Associate Fellow AIAA.

§Assistant Professor, Department of Aerospace and Mechanical Engineering; jmacart@nd.edu. Member AIAA.

equilibrium, which affects the applicability of various flow models [1]. For  $\text{Kn} \leq 10^{-3}$ , the continuum Navier–Stokes equations are an appropriate fluid model. However, existing computationally efficient hydrodynamic models, perturbative in nature [1], become unreliable as  $\text{Kn}$  rises due to either low densities or regions of high local gradients (as in hypersonic flows), which cause deviations from the linear Newtonian and Fourier laws. The high-Knudsen-number regime can be further divided into the transition-continuum regime ( $0.01 \leq \text{Kn} \leq 1$ ) and the free-molecule flow regime ( $\text{Kn} \geq 1$ ).

From a theoretical standpoint, the kinetic description of interacting particles via the solution of the Boltzmann equation allows for the most accurate description of nonequilibrium hypersonic flows. Because of the difficulties associated with the solution of the Boltzmann equation, approximate physics-based methods—at different spatial and time scales—have been developed over the years by the scientific community. This section aims to provide a brief overview and a critical analysis of the approaches currently used in the literature to alleviate the computational challenges of solving the Boltzmann equation.

These methods can be classified into two broad categories: Direct Simulation Monte Carlo (DSMC) and extended-hydrodynamics methods. DSMC involves the simulation of collisions between many “computational molecules” that statistically describe the physical dynamics and interactions among fluid molecules [2, 3]. While popular, DSMC remains challenging for mixed transition-continuum and continuum flows, though modern DSMC can solve complicated, multiphysical flows given sufficient computing power [4, 5]. The computational cost of DSMC can be alleviated by extended-hydrodynamics methods, which use the fluid equations in conjunction with higher-order constitutive models for the viscous stress and heat flux terms to extend their validity beyond the continuum regime. Many of these methods use the Chapman–Enskog expansion, which uses  $\text{Kn}$ -series approximations (of varying order) to the Boltzmann equation. Retaining the zero-order terms gives the Euler equations, first-order terms the Navier–Stokes equations, and second-order terms the Burnett equations [6]. The Navier–Stokes equations are valid in the continuum limit ( $\text{Kn} \rightarrow 0$ ) of small mean free paths [6, 7]. However, large perturbations of the velocity distribution from equilibrium render the continuum assumption invalid. Nonperturbative moment methods [8] mitigate the problem of the breakdown away from moderate nonequilibrium by predicting values for momentum and energy fluxes that are consistent with the nonnegativity of the particle density. These methods, however, are storage-intensive and computationally impractical due to the large number of independent variables required to evaluate integral terms. Several variants of the Burnett equations have been proposed including the BGK–Burnett (Bhatnagar–Gross–Krook) equations [9], Woods equations [10], Grad’s 13-moment equations [11, 12], regularized Burnett [13], and regularized Grad’s 13-moment equations [14]. A major drawback of most is their mathematical ill-posedness in their original formulations, which leads to physical and numerical instabilities. Thermodynamic consistency is also not an intrinsic property of these equations; this limits their general applicability, despite their computational cost being comparable to the Navier–Stokes equations. The advantages of DSMC’s detailed kinetic information and the continuum-like methods’ computational tractability can be obtained by hybrid schemes [15–24], which solve the Navier–Stokes equations in continuum regions of the flow and employ the DSMC technique elsewhere. Although these methods are attractive for multiscale flows, identifying the continuum and DSMC (molecular) regions is often done from the continuum perspective, numerically interfacing these regions is challenging, and the speedup obtained is often insufficient to render it an attractive alternative to pure DSMC. Deterministic numerical methods that directly solve the Boltzmann equation [25] have been tested for simple flow configurations, though their extension to practical problems is yet to be explored.

This paper describes the construction of a physics-based reduced-order model, able to capture the fundamental physical processes that occur in hypersonic flows, yet computationally efficient, by leveraging fundamental physics, computational science, and applied mathematics. The proposed model compensates for the deficiencies of the Navier–Stokes model in the transition-continuum regime by introducing additional terms into the continuum equations. By estimating the parameters of these terms (we use neural networks) while converging the Navier–Stokes solutions to Boltzmann solutions, the model attempts to encode the unrepresented physics that are missing from the continuum description of the gas. We also propose algebraic constraints to the model outputs to ensure consistency with the second law of thermodynamics.

Applications of deep learning (DL) for closure modeling are widespread in incompressible fluid mechanics [26–30] but are much less common for compressible flows and hypersonics. Most previously employed machine learning methods for flows attempt to estimate unclosed terms (e.g., the Reynolds stress or the subgrid stress) by minimizing the direct mismatch between these terms and “trusted” data, which are often obtained from highly resolved direct numerical simulations (DNS) [31–33]. This approach, which we call *a priori* training, decouples the parameter estimation (optimization) step from the solution of the governing partial differential equations (PDEs), which is mathematically inconsistent due to the noncommutativity of nonlinear operations [34]. In our view, this approach has two primary deficiencies for compressible flows. First, the nonlinear coupling of dependent variables inherent to compressible flows

intensifies the inconsistency of *a priori* training. Second, and more importantly, when the discrepancy is due to a fundamental deficiency in the PDE model—for example, for transition-continuum flows, approximations to the viscous stress and heat flux from kinetic theory—the functional form of the discrepancy is often unknown; thus, even defining the loss function for the direct mismatch is challenging.

Our deep learning approach is conceptually simple yet avoids these challenges. Rather than decouple the optimization step from the PDE solution, we use an *embedded optimization* method that substitutes the untrained model into the PDE and estimates its parameters *in situ* [28, 34, 35]. The model thus respects its effects on the nonlinearly coupled PDE solution. Additionally, this approach can directly compute the loss function from the dependent variables (mass, momentum, and energy) or derived quantities (e.g., temperature) rather than from higher-order unclosed terms appearing in the PDEs. This allows us to obtain target data from Boltzmann solutions and, in theory, enables straightforward incorporation of experimental data. We present an embedded optimization approach that optimizes over the PDE solution by solving the adjoint Navier–Stokes equations, which are obtained using algorithmic differentiation (AD) of the forward equations and are solved efficiently using graphics processing unit (GPU) acceleration. Similar methods have been successfully applied to turbulence modeling for incompressible large-eddy simulation (LES) [34, 36, 37], Reynolds-averaged Navier–Stokes (RANS) simulation [35, 38] and sensitivity analysis of chaotic dynamical systems [39, 40]. Given the code-intrusive nature of adjoint-based optimization, many efforts to avoid this difficulty have used less-complete couplings of the optimization step and the PDE solution [41–43], though these methods typically require trusted data for the exact term being modeled, which is not always possible.

This work presents an example application to the prediction of the shock structure in rarefied argon flows, compares the augmented Navier–Stokes predictions to the unmodified predictions, assesses the Mach number interpolation and extrapolation capability of the DL model, and provides a physical interpretation of the modifications to the continuum transport terms. The governing equations and transport models from kinetic theory are briefly reviewed in Section III. The embedded deep learning framework and constraints to enforce the second law of thermodynamics are introduced in Section IV. An application to transition-continuum argon shock predictions is presented in Section V. A summary and discussion of future work are provided in Section VI.

### III. Continuum Equations and Closure Models

Let  $g(t, \mathbf{x}, \mathbf{v})$  be the solution to the Boltzmann equation (the distribution function) for  $t \in [0, \infty)$  and  $\mathbf{x}, \mathbf{v} \in \mathbb{R}^3$ , and define the mass density  $\rho(t, \mathbf{x}) = \langle m, g(t, \mathbf{x}, \cdot) \rangle$  and fluid-dynamic momentum  $\rho \mathbf{u}(t, \mathbf{x}) = \langle m g(t, \mathbf{x}, \cdot), \mathbf{v} \rangle$ , where  $m$  is the mass of a particle, and integrals of the distribution function over the velocity space are obtained from  $\langle h(\cdot), g(t, \mathbf{x}, \cdot) \rangle = \int_{\mathbb{R}^3} h(\mathbf{v}) g(t, \mathbf{x}, \mathbf{v}) d\mathbf{v}$ . In this manner, all continuum quantities of interest  $\mathbf{U} = \{\rho, \mathbf{u}, p, T, \dots\}$  may be obtained, where  $p$  is the pressure, and  $T$  is the temperature, among other possibilities.

The compressible Navier–Stokes equations for a single-component gas, as derived from the Boltzmann equation, are

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) &= 0 \\ \frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot \left[ \rho \mathbf{u} \otimes \mathbf{u} - \tau(\mathbf{u}, \mathbf{v}, g) \right] &= 0 \\ \frac{\partial \rho E}{\partial t} + \nabla \cdot \left[ \rho E \mathbf{u} - \tau(\mathbf{u}, \mathbf{v}, g) : \mathbf{u} + \mathbf{q}(\mathbf{u}, \mathbf{v}, g) \right] &= 0, \end{aligned} \quad (1)$$

where  $E = e + \mathbf{u}^T \mathbf{u} / 2$  is the total energy and  $e$  is the internal energy. The thermal momentum flux

$$\tau(\mathbf{u}, \mathbf{v}, g) = \left\langle m \mathbf{v} \otimes \mathbf{v} - \rho \mathbf{u} \otimes \mathbf{u}, g(t, \mathbf{x}, \cdot) \right\rangle \quad (2)$$

and the heat flux  $\mathbf{q}(\mathbf{u}, \mathbf{v}, g)$  require integration over the Boltzmann phase space; thus, they are unclosed under the continuum variables. Closure of (2) is commonly achieved by assuming proportionality to the strain-rate tensor and defining the pressure, yielding  $\tau = \sigma - p \mathbf{I}$ , where  $\sigma = \mu(\nabla \mathbf{u} + \nabla \mathbf{u}^T) - \frac{2}{3} \mu(\nabla \cdot \mathbf{u}) \mathbf{I}$ ,  $\mu$  is the viscosity, and the pressure satisfies the ideal-gas equation of state  $p = (\gamma - 1)(\rho E - \rho \mathbf{u}^T \mathbf{u} / 2)$ , where  $\gamma$  is the (constant) ratio of specific heats. Similarly, the heat flux is typically closed using  $\mathbf{q} = -\lambda \nabla T$ , where  $\lambda$  is the thermal conductivity,  $T = p / (\rho R)$ , and  $R$  is the specific gas constant. However, these closures—and the use of the Navier–Stokes equations more generally—require the continuum approximation to hold, which can inaccurately represent the physics of nonequilibrium flows in the transition-continuum regime.

### A. Neural-Network Closures

We introduce three approaches for the closure of (1) using neural networks. First, the **transport coefficients**  $\mu$  and  $\lambda$  may be **corrected** using neural networks (Approach A),

$$\begin{aligned}\mu(\mathbf{U}; \theta) &= \mu_0(1 + f_1(\mathbf{U}_x; \theta)), \\ \lambda(\mathbf{U}; \theta) &= \lambda_0(1 + f_2(\mathbf{U}_x; \theta)),\end{aligned}\tag{3}$$

where  $f_1(\mathbf{U}_x; \theta)$  and  $f_2(\mathbf{U}_x; \theta)$  are the outputs of a neural network  $f(\mathbf{U}_x; \theta)$ , which takes the point-wise derivatives of  $\mathbf{U}$  as its input and has parameters  $\theta$ . The resulting viscous stress and heat flux are

$$\begin{aligned}\sigma(\mathbf{U}; \theta) &= \mu(\mathbf{U}; \theta) \left( (\nabla \mathbf{u} + (\nabla \mathbf{u})^\top) - \frac{2}{3}(\nabla \cdot \mathbf{u})\mathbf{I} \right) \quad \text{and} \\ \mathbf{q}(\mathbf{U}; \theta) &= -\lambda(\mathbf{U}; \theta)\nabla T.\end{aligned}\tag{4}$$

Second, the continuum viscous-stress and heat-flux models may be **augmented** by linearly superimposing neural networks (Approach B). This results in augmented viscous-stress and heat-flux terms,

$$\begin{aligned}\sigma(\mathbf{U}; \theta) &= \mu \left( (\nabla \mathbf{u} + (\nabla \mathbf{u})^\top) - \frac{2}{3}(\nabla \cdot \mathbf{u})\mathbf{I} \right) - f_1(\mathbf{U}_x; \theta) \\ \mathbf{q}(\mathbf{U}; \theta) &= -\lambda\nabla T - f_2(\mathbf{U}_x; \theta),\end{aligned}\tag{5}$$

where  $f_1(\mathbf{U}_x; \theta)$  and  $f_2(\mathbf{U}_x; \theta)$  are again neural-network outputs, though they need not be the same networks as in (3). Finally, the continuum models  $\sigma(\mathbf{U})$  and  $\mathbf{q}(\mathbf{U})$  may be **directly modeled** using neural networks (Approach C),

$$\sigma(\mathbf{U}; \theta) = f_1(\mathbf{U}_x; \theta) \quad \text{and} \quad \mathbf{q}(\mathbf{U}; \theta) = f_2(\mathbf{U}_x; \theta).\tag{6}$$

We pursue only Approaches A and B here. These have more limited potential than Approach C, being constrained by the algebraic models for  $\sigma$  and  $\mathbf{q}$ , but are more straightforward to implement and analyze. Rather than directly model the transport terms, our goal is to select  $\theta$  such that the new models account only for *deviations* of the classical relations from kinetic theory. The neural network  $f(\mathbf{U}_x; \theta)$  structure and optimization are discussed further in Section IV.

An important uncertainty in the physical modeling of the shock structure, especially at high Mach numbers, is the temperature dependence of the transport coefficients  $\mu$  and  $\lambda$ . We use a power-law dependence of the baseline continuum transport coefficients on temperature,

$$\mu(T) = \mu_0 \left( \frac{T}{T_0} \right)^s \quad \text{and} \quad \lambda(T) = \lambda_0 \left( \frac{T}{T_0} \right)^s,\tag{7}$$

where  $T_0 = 273.15$  K, is typically used for argon. Several previous efforts have aimed to adjust the value of  $s$  to reproduce experimental shock profiles [3, 44–46]. Following these, we use  $s = 0.74$  to match the viscosity law exponent used in our DSMC simulations.

### B. Application to 1D Shocks

We consider an application to a 1D stationary shock in argon at  $p_\infty = 6.667$  Pa and  $T_\infty = 300$  K. This is a common shock-structure problem for extended hydrodynamics models [45–51] that has been extensively studied experimentally [44, 52–55]. The flow has a transition-continuum Knudsen number  $\text{Kn} = 0.2 \sim 0.3$  within the shock, which Navier–Stokes equations do not accurately model. A monatomic gas such as argon is generally chosen to study the shock structure due to the lack of vibrational or rotational nonequilibrium. We assume a constant ratio of specific heats  $\gamma = 5/3$  with gas constant  $R = 208.12$  kJ/kg · K. The 1D stationary solution of (1) satisfies

$$\frac{\partial \mathbf{U}}{\partial t} = -\frac{\partial \mathbf{F}_c}{\partial x} + \frac{\partial \mathbf{F}_d}{\partial x} = 0,\tag{8}$$

where

$$\mathbf{U} = \begin{bmatrix} \rho \\ \rho u \\ \rho E \end{bmatrix}, \quad \mathbf{F}_c = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho u H \end{bmatrix}, \quad \mathbf{F}_d = \begin{bmatrix} 0 \\ \sigma(\mathbf{U}; \theta) \\ \sigma(\mathbf{U}; \theta)u - q(\mathbf{U}; \theta) \end{bmatrix},$$

where  $H = e + p/\rho + 1/2u^2$  is the total enthalpy, and  $\sigma(\mathbf{U}; \theta)$  and  $q(\mathbf{U}; \theta)$  are the 1D versions of the viscous-stress tensor and heat-flux vector with embedded neural networks.  $\mathbf{U}$  satisfies Dirichlet (supersonic freestream) boundary conditions at  $x = -\infty$  and subsonic nonreflective characteristic boundary conditions at  $x = +\infty$  [56].

The 1D equations are solved on a domain  $L_x \in [-20 \text{ mm}, 10 \text{ mm}]$  with the shock positioned at  $x = 0$ . Inviscid shock solutions provide the initial conditions and Dirichlet boundary conditions. The equations are discretized over  $n_x = 256$  uniformly spaced mesh cells with the Euler fluxes discretized using a characteristic-based scheme [57]; all other terms are discretized using second-order centered differences. The equations are converged to their steady-state solution using a damped Newton method; details are provided in Section III.C

Target data are obtained from DSMC solutions of the Boltzmann equation at freestream Mach numbers  $M_\infty = 2, 3, 5, 6, 8, 9, \text{ and } 10$ . The DSMC simulations used a VHS collision model and were initialized using the same pre- and post-shock conditions as the Navier–Stokes simulations. The viscosity law exponent  $s = 0.74$  was found to best match experimental data [44]. For further details about the DSMC simulations, readers are directed to [58, 59]. We denote the target data as  $\bar{\rho}$ ,  $\bar{u}$  and  $\bar{T}$ .

### C. Damped Newton Method

Fully converging the forward solution is necessary to ensure well-posedness of the embedded optimization problem. This is achieved using a damped Newton method, which is able to converge the solution to near machine precision (using 64-bit floating-point precision). By comparison, pseudo-time-stepping methods require taking prohibitively many steps to achieve comparable convergence. For the present shock cases with moderately high Mach numbers (e.g.,  $M = 5$ ) converging the relative loss to  $O(10^{-2})$  requires approximately 1200 pseudo-time iterations but only approximately 20 damped Newton iterations. Even with the higher cost of the Newton solver, the lower number of iterations reduces the training time by approximately a factor of ten.

Consider the steady solution of (8),  $\partial(\mathbf{F}_d - \mathbf{F}_c)/\partial x = \partial\mathbf{F}_x/\partial x = 0$ . The  $n^{\text{th}}$  Newton iteration solves  $\Delta\mathbf{U}$  from the linearized system

$$\left. \frac{d\mathbf{F}_x}{d\mathbf{U}} \right| \Delta\mathbf{U} = -\mathbf{F}_x(\mathbf{U}^n) \quad (9)$$

and advances the solution as  $\mathbf{U}^{n+1} = \mathbf{U}^n + \Delta\mathbf{U}$  until  $\Delta\mathbf{U}$  falls below a specified tolerance. Given the low dimensionality of the 1D problem, we solve (9) using Gauss–Jordan elimination, though 2D and 3D problems would require more sophisticated sparse or iterative solution methods. Since the embedded optimization procedure currently uses the Jacobian  $\partial\mathbf{F}_x/\partial\mathbf{U}$ , the Jacobian can be reused for the Newton solver, though both could be posed in a Jacobian-free manner.

A practical difficulty with Newton iteration is the lack of physical constraints on  $\Delta\mathbf{U}$ . Large values of  $\Delta\mathbf{U}$  obtained by solving (9) can result in nonphysical values of  $\mathbf{U}$  (e.g., negative density or pressure). We avoid these problems by bracketing  $\Delta\mathbf{U}$  such that it does not exceed a specified fraction of the local solution. The bracketing operation is

$$\Delta\mathbf{U}_b = \min \left( f_b \|\mathbf{U}_\infty\|_2 \times \frac{\Delta\mathbf{U}}{\|\Delta\mathbf{U}\|_2}, \Delta\mathbf{U} \right), \quad (10)$$

where  $f_b = 0.1$  is the bracketing factor,  $\|\mathbf{U}_\infty\|_2$  is the  $L_2$  norm of the freestream  $\mathbf{U}$ , and  $\Delta\mathbf{U}/\|\Delta\mathbf{U}\|_2$  is the unit vector in the direction of  $\Delta\mathbf{U}$ . To further ensure physical solutions,  $\Delta\mathbf{U}_b$  is recursively damped,  $\Delta\mathbf{U}_d = f_d \Delta\mathbf{U}_b$ , where  $f_d = 0.2$  is a damping factor, until  $\mathbf{U}^{n+1}$  is physically consistent (e.g., the density and pressure are positive). The solution is then advanced as

$$\mathbf{U}^{n+1} = \mathbf{U}^n + \Delta\mathbf{U}_d. \quad (11)$$

Using the Newton method along with bracketing and damping stabilizes the solution before taking an optimization step to update the neural network parameters.

## IV. Machine Learning Closure

This section is organized as follows: IV.A outlines the architecture of the embedded neural network and its inputs, IV.B details the standard *a priori* optimization approach and the adjoint-based optimization approach used in this paper, and IV.C presents a brief derivation of the entropy constraints introduced to enforce the second law of thermodynamics.

## A. Model Architecture

The inputs to the embedded neural network are the first-order spatial derivatives of the computed scalar flow quantities and velocity at a mesh point  $x_i$ :  $\hat{z}_i = \{\rho_x(x_i), p_x(x_i), T_x(x_i), u_x(x_i)\}$ . These inputs are invariant to Galilean transformations (constant relative motion), being spatial derivatives of scalar and vector quantities, which also guarantees that the network output quantities are invariant to Galilean transformations [31, 60]. We do not consider rotational invariance, which is a separate issue that has been assessed elsewhere [31]. The model inputs are expanded to include data at the two neighboring mesh points:  $z_i = \{\hat{z}_{i-1}, \hat{z}_i, \hat{z}_{i+1}\}$ . This expanded set of inputs serves to enable the model to compute higher derivatives, though it introduces a degree of mesh dependence, which could reduce generalizability to out-of-sample meshes. Addressing the issues of rotational invariance and mesh dependence is the subject of ongoing applications to two- and three-dimensional flows.

The neural network has a four-layer structure with one gate layer. For inputs  $z \in \mathbb{R}^{18 \times n_x}$ , the model architecture is

$$\begin{aligned} H^1 &= \sigma_1(W^1 z + b^1) \\ H^2 &= \sigma_1(W^2 H^1 + b^2) \\ H^3 &= G \odot H^2 \quad \text{with } G = \sigma_2(W^3 z + b^3) \\ \hat{f}(z; \theta) &= W^4 H^3 + b^4 \end{aligned}$$

with output dimensions  $\hat{f}(z; \theta) \in \mathbb{R}^{2 \times n_x}$ . In these,

$$\sigma_1(x) = \begin{cases} \alpha(\exp(x) - 1), & x \leq 0 \\ x, & x > 0 \end{cases}$$

is the exponential linear unit (ELU) activation function with hyperparameter  $\alpha = 1$ ,  $\sigma_2(x) = 1/(1 + e^{-x})$  is the sigmoid activation function,  $\odot$  signifies element-wise multiplication, and the parameters are  $\theta = \{W^1, W^2, W^3, W^4, b^1, b^2, b^3, b^4\}$ . We use  $N_H = 1200$  hidden units per layer, which balances training accuracy and computational efficiency.

For Approach B (augmented continuum viscous-stress and heat-flux models; (5)), the output layer is chosen to enforce the entropy conditions introduced in Section IV.C. For Approach A (corrected transport coefficients; (3)), the raw neural network outputs  $\hat{f}(z; \theta)$  are passed through a final ELU layer,

$$f(z; \theta) = \sigma_1(\hat{f}(z; \theta)) + 0.1,$$

which is a simpler constraint to ensure positivity of  $\mu_0(1 + f_1(z; \theta))$  and  $\lambda_0(1 + f_2(z; \theta))$  for all  $z$  and  $\theta$ .

## B. Online Optimization Approach

Optimization of standard *a priori* DL closures is constrained to use objective (loss) functions that are explicit functions of the neural network. For example, if one wished to model the viscosity  $\mu$  with a neural network  $f(z; \theta)$ , then it would be sufficient to minimize the objective function

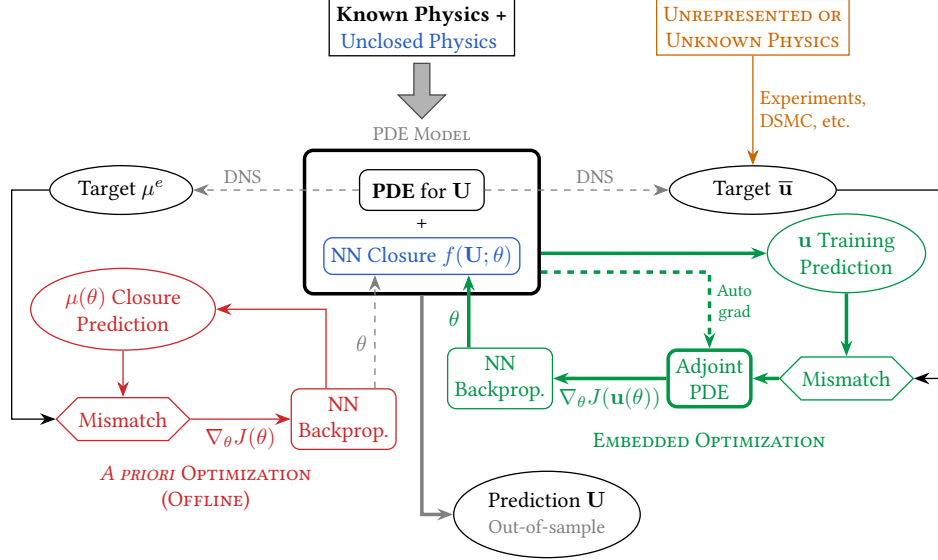
$$J(\theta) = \frac{1}{2} \int_{-\infty}^{\infty} \left( f(z(x); \theta) - \mu^e \right)^2 dx, \quad (12)$$

where  $\mu^e$  is trusted data (usually obtained from DNS), and  $z(x)$  are the input variables. To select  $\theta$  using gradient-descent optimization, one needs to compute the gradient

$$\nabla_{\theta} J = \int_{-\infty}^{\infty} \left( f(z; \theta) - \mu^e \right) \nabla_{\theta} f(z(x); \theta) dx,$$

where  $\nabla_{\theta} f(z(x); \theta)$  may be evaluated using AD over the neural network, which is standard functionality in DL software libraries. However, optimization using (12) is problematic when data for the exact closure term ( $\mu^e$ ) are not available—for example, when using experimental or higher-order simulation data (e.g., DSMC)—or when the target quantity is not known or accurately estimable from the available data. Furthermore, this optimization takes place offline, without solving the governing PDEs, so at best is only indirectly PDE-constrained by the target data.

The embedded DL approach addresses these shortcomings by introducing an online optimization algorithm [38, 61], which enables arbitrarily defined objective functions. For the application to nonequilibrium flows, we construct an



**Fig. 1** Two methods to optimize a NN closure  $f(\mathbf{u}; \theta)$  in a PDE model to represent unclosed physics. *A priori* optimization selects the model parameters  $\theta$  offline (without solving the PDE model). Embedded optimization instead minimizes the error with respect to the PDE solution  $\mathbf{U}$  (or derived quantities  $\mathbf{u}$ ) by solving adjoint PDEs, which enables efficient computation of the gradient  $\nabla_{\theta} J(\mathbf{u}(\theta))$ .

objective function using the mismatch between the computed primitive variables  $\mathbf{u}_p = [\rho, u, T]^{\top}$  and target data obtained from DSMC solutions of the distribution function. Using an objective function constructed of primitive variables  $\mathbf{u}_p$  rather than conserved variables  $\mathbf{U}$  resulted in better convergence and lower losses. In general, the target data may be known from DNS, experiments, and/or higher-order simulations. The loss function in terms of the primitive variables is

$$J(\mathbf{u}_p(\theta)) = \frac{1}{2} \int_{-\infty}^{\infty} \left[ \frac{1}{\rho_{\infty}^2} (\rho(t, x; \theta) - \bar{\rho})^2 + \frac{1}{u_{\infty}^2} (u(t, x; \theta) - \bar{u})^2 + \frac{1}{T_{\infty}^2} (T(t, x; \theta) - \bar{T})^2 \right] dx, \quad (13)$$

where  $\rho_{\infty}$ ,  $u_{\infty}$ , and  $T_{\infty}$  are the reference freestream values used for normalization, and where we emphasize the dependence of the flow quantities on  $\theta$ . The objective function depends on  $\theta$  implicitly through  $\mathbf{u}_p$ .

Figure 1 illustrates the conceptual differences between *a priori* and embedded optimization. Notably, *a priori* optimization requires high-fidelity (e.g., DNS) target data for the closure term, which may not be available or accurately known. (This would be the case, for example, for DNS solutions of the Navier–Stokes equations in the transition-continuum regime.) *A priori* optimization cannot target the flow variables (or derived quantities), thus it cannot straightforwardly incorporate experimental and/or higher-order simulation data. In the same vein, the embedded optimization approach does not necessarily need DNS data.

We now consider the optimization of the discretized versions of (8) and (13); all variables in the remainder of this section will therefore be finite-dimensional. This optimization problem is challenging, for it requires computing the gradients of dependent variables and derived quantities with respect to the neural network parameters  $\theta$ . A naïve approach would attempt to evaluate

$$\nabla_{\theta} J = \frac{\partial \rho^{\top}}{\partial \theta} \frac{\partial J}{\partial \rho} + \frac{\partial u^{\top}}{\partial \theta} \frac{\partial J}{\partial u} + \frac{\partial T^{\top}}{\partial \theta} \frac{\partial J}{\partial T} = \frac{\partial \mathbf{u}_p^{\top}}{\partial \theta} \frac{\partial J}{\partial \mathbf{u}_p} \quad (14)$$

on the discrete computational mesh. The partial derivatives  $\partial \mathbf{u}_p / \partial \theta \in \mathbb{R}^{3n_x \times n_{\theta}}$  could be evaluated by deriving and solving PDEs for each  $\partial \mathbf{u} / \partial \theta_i$ , where  $i = 1, \dots, n_{\theta}$ , or by introducing finite-difference perturbations  $\theta_i + \theta'_i$ . However, the number of neural network parameters  $n_{\theta}$  is often large (e.g.,  $\sim 10^5$  parameters), such that both of these approaches would require an intractable number of calculations.

Instead, for  $n$  forward equations,  $n$  adjoint equations may be solved to provide the necessary gradients. Thus the overall cost will not be significantly higher for optimization than for prediction. To derive the adjoint equations for a

steady problem, note that

$$\nabla_{\theta} \mathbf{F}_x = \frac{\partial \mathbf{F}_x}{\partial \mathbf{u}_p} \frac{\partial \mathbf{u}_p}{\partial \theta} + \frac{\partial \mathbf{F}_x}{\partial \theta} = 0, \quad (15)$$

where  $\mathbf{F}_x$  is the discrete representation of the right-hand side (RHS)  $\frac{\partial}{\partial x} (\mathbf{F}_d - \mathbf{F}_c)$ .

Let us now introduce the adjoint variables  $\hat{\mathbf{u}} = [\hat{\rho}, \hat{u}, \hat{T}]^{\top}$ , which satisfy the adjoint equation

$$\frac{\partial J}{\partial \mathbf{u}_p} + \frac{\partial \mathbf{F}_x}{\partial \mathbf{u}_p}^{\top} \hat{\mathbf{u}} = 0, \quad (16)$$

with Dirichlet boundary conditions  $(\hat{\rho}, \hat{u}, \hat{T}) = 0$  at  $x = \pm\infty$ . Multiplying (15) by  $\hat{\mathbf{u}}^{\top}$  produces

$$0 = \hat{\mathbf{u}}^{\top} \frac{\partial \mathbf{F}_x}{\partial \mathbf{u}_p} \frac{\partial \mathbf{u}_p}{\partial \theta} + \hat{\mathbf{u}}^{\top} \frac{\partial \mathbf{F}_x}{\partial \theta}. \quad (17)$$

From (16), we have that  $\hat{\mathbf{u}}^{\top} \frac{\partial \mathbf{F}_x}{\partial \mathbf{u}_p} = -\frac{\partial J}{\partial \mathbf{u}_p}^{\top}$  which, when substituted into (17), yields

$$\frac{\partial J}{\partial \mathbf{u}_p}^{\top} \frac{\partial \mathbf{u}_p}{\partial \theta} = \hat{\mathbf{u}}^{\top} \frac{\partial \mathbf{F}_x}{\partial \theta}. \quad (18)$$

The relationship (18) combined with (14) provides a formula for the gradient in terms of the adjoint solution:

$$\nabla_{\theta} J = \frac{\partial \mathbf{F}_x}{\partial \theta}^{\top} \hat{\mathbf{u}}. \quad (19)$$

Therefore, we can solve the adjoint equation (16) for  $\hat{\mathbf{u}}$  and use its solution to evaluate the gradient of the loss function with respect to the model parameters ( $\nabla_{\theta} J$ ) using (19). The dimension of (16) is equal to the dimension of the forward equations (8), so the overall optimization is not significantly greater than the cost of solving the forward equations. It is important to note that the validity of (19) is contingent on the right-hand side of (15) being zero. In turn, this requires that the discrete residual of the forward equations is close to zero before taking an optimization step.

The discrete adjoint equations are solved as a linear system. This requires the computation of the Jacobian  $\partial \mathbf{F}_x / \partial \mathbf{u}_p$ , which is calculated using algorithmic differentiation as provided by the *PyTorch* library [62]. After converging the forward and adjoint equations, gradient-descent steps are taken for parameters  $\theta^{(k)}$  according to

$$\theta^{(k+1)} = \theta^{(k)} - \alpha^{(k)} \nabla_{\theta} J(\theta^{(k)}), \quad (20)$$

where  $\alpha^{(k)}$  is the learning rate at the iteration  $k \in \{0, 1, \dots, N_{\text{iter}}\}$ , where  $N_{\text{iter}}$  is the maximum number of optimization iterations. The learning-rate schedule and other model-training hyperparameters are described in Section IV.D.

### C. Entropy Constraints

Using DL closures unconstrained for the second law of thermodynamics can result in unphysical flow quantities. We constrain the outputs of the neural networks to satisfy the second law of thermodynamics using the Clausius–Duhem inequality [63],

$$\dot{\mathcal{I}} = -\frac{1}{T^2} \mathbf{q}(\mathbf{U}; \theta) \nabla T + \frac{1}{T} \sigma(\mathbf{U}; \theta) : \nabla \mathbf{u}^T \geq 0, \quad (21)$$

where  $\dot{\mathcal{I}}$  is the rate of irreversible entropy production. A “strong” form of the entropy constraint would impose conditions on the model outputs such that (21) is nonnegative. For Approach A (transport-coefficient corrections), ensuring the positivity of  $\mu(\mathbf{U}; \theta)$  and  $\lambda(\mathbf{U}; \theta)$  is sufficient; this is done using an ELU output layer as described in Section IV. A “weak” constraint for the entropy inequality would incorporate  $\dot{\mathcal{I}}$  as a penalty term in the training loss function (13). The application of strong and weak constraints for Approach B (transport-term augmentation) is now discussed separately.

#### 1. Strong Entropy Constraint

The Clausius–Duhem inequality (21) can be rearranged to give

$$\frac{1}{T} \sigma(\mathbf{U}; \theta) : \nabla \mathbf{u}^T \geq -\frac{1}{T^2} \mathbf{q}(\mathbf{U}; \theta) \nabla T, \quad (22)$$



from which a constraint on the viscous stress may be obtained:

$$\sigma(\mathbf{U}; \theta) \geq \frac{\frac{1}{T} \mathbf{q}(\mathbf{U}; \theta) \nabla T}{\nabla \mathbf{u}^T}. \quad (23)$$

A constrained viscous-stress model is therefore

$$\tilde{\sigma}(\mathbf{U}; \theta) = \max(\sigma(\mathbf{U}; \theta), \mathbf{N}), \quad (24)$$

where

$$\mathbf{N} = \frac{\frac{1}{T} \mathbf{q}(\mathbf{U}; \theta) \nabla T}{\nabla \mathbf{u}^T}, \quad (25)$$

and the  $\max()$  function in (24) is applied in a pointwise manner. After applying (24), the constrained closure terms  $\mathbf{q}(\theta)$  and  $\tilde{\sigma}(\theta)$  are substituted into the governing equations. This enforcement of the strong entropy condition makes no assumptions about the signs of  $\nabla T$  and  $\nabla u$ .

For 1D flows, the Approach B closures (5) can be simplified to

$$\begin{aligned} \sigma(\mathbf{U}; \theta) &= \frac{4}{3} \mu \frac{\partial u}{\partial x} - f_1(\mathbf{U}_x; \theta) \\ \mathbf{q}(\mathbf{U}; \theta) &= -\lambda \frac{\partial T}{\partial x} - f_2(\mathbf{U}_x; \theta), \end{aligned} \quad (26)$$

and the Clausius-Duhem inequality (21) reduces to

$$\dot{I} = -\frac{1}{T^2} \mathbf{q}(\mathbf{U}; \theta) \frac{\partial T}{\partial x} + \frac{1}{T} \sigma(\mathbf{U}; \theta) \frac{\partial u}{\partial x} \geq 0. \quad (27)$$

The term  $\mathbf{N}$  (25) reduces to

$$\mathbf{N} = \frac{\frac{1}{T} \mathbf{q}(\mathbf{U}; \theta) \frac{\partial T}{\partial x}}{\frac{\partial u}{\partial x}}, \quad (28)$$

and the strong entropy constraint can be enforced according to (24).

## 2. Weak Entropy Constraint

A weak entropy constraint would penalize the violation of the entropy inequality (21) during the constrained optimization process by adding an extra term to the training loss function (13),

$$J_s(\theta) = W_s \int_{-\infty}^{\infty} [\min(0, f_1(\mathbf{U}_x; \theta)) + \min(0, f_2(\mathbf{U}_x; \theta))] dx, \quad (29)$$

where  $W_s$  is a scalar weight. The net training loss for the weak entropy constraint is thus

$$J_w(\mathbf{u}_p; \theta) = J(\mathbf{u}_p(\theta)) + J_s(\mathbf{u}_p(\theta)). \quad (30)$$

The requirement to simultaneously minimize both components of  $J_w$  prevented weak-constraint cases from matching the testing and training accuracy of equivalent strong-constraint cases. This could be improved by dynamic selection of  $W_s$ , for example, to enforce the  $J_s$  term more strongly in regions of high entropy production.

## D. Model Training

A distributed training approach was implemented to perform the online adjoint optimization. For each of Approaches A and B, models were trained for four different training regimens: three models targeted only a single freestream Mach number ( $M_{\text{train}} = 2$ ,  $M_{\text{train}} = 5$ , and  $M_{\text{train}} = 8$ ), and one model targeted three Mach numbers simultaneously:  $M_{\text{train}} = (2, 5, 8)$ . Training this combined model utilized message-passing interface (MPI) parallelization to average the loss-function gradients  $\nabla_{\theta} J$  between the component simulations before optimizing.

Table 1 provides the training/testing regimen matrix. Freestream Mach numbers  $M_{\infty} = 2, 5, 8$  were out-of-sample when not included in a model's training regimen. Freestream Mach numbers  $M_{\infty} = 3, 6, 9, 10$  were out-of-sample for

Training	$M_\infty = 2$	$M_\infty = 3$	$M_\infty = 5$	$M_\infty = 6$	$M_\infty = 8$	$M_\infty = 9$	$M_\infty = 10$
$M_{\text{train}} = 2$	•						
$M_{\text{train}} = 5$			•				
$M_{\text{train}} = 8$					•		
$M_{\text{train}} = (2, 5, 8)$	•		•		•		
Testing	○	•	○	•	○	•	•

**Table 1 Model training and testing matrix.**  $M_\infty = 2, 5, 8$  (○) were out-of-sample when not used for training.  $M_\infty = 3, 6, 9, 10$  (•) were out-of-sample for all models.

all trained models. This testing arrangement is designed to evaluate the stability and accuracy of the trained models for interpolation ( $M_{\text{train}} = (2, 5, 8)$  only) and extrapolation to lower and higher Mach numbers.

Initial Navier–Stokes viscous-shock solutions for each freestream Mach number listed in Table 1 were obtained by converging from the corresponding inviscid-shock profile. Training and testing cases were initialized from the converged Navier–Stokes solutions. For model training, the parameters of randomly-initialized neural networks were selected using (20) with initial learning rate  $\alpha^{(0)} = 2.5 \times 10^{-6}$ . To ensure convergence, the learning rate was reduced using an adaptive schedule:

$$\alpha = \begin{cases} \alpha, & \epsilon_{\text{rel}} > \epsilon_T \\ 0.75\alpha, & \epsilon_{\text{rel}} \leq \epsilon_T \end{cases}, \quad \epsilon_T = \begin{cases} \epsilon_T, & \epsilon_{\text{rel}} > \epsilon_T \\ 0.75\epsilon_T, & \epsilon_{\text{rel}} \leq \epsilon_T \end{cases},$$

where  $\epsilon_{\text{rel}} = J(\theta)/J_0$  is the objective-function relative error,  $J_0$  is the initial value of the objective function (i.e., for the initial, unmodified Navier–Stokes solution), and the starting target error is  $\epsilon_T = 0.9$ . For models trained simultaneously for  $M_{\text{train}} = (2, 5, 8)$ , the minimum  $\alpha$  and  $\epsilon_T$  were synchronized between training cases. The training process was stopped for change in relative error  $\epsilon_{\text{rel}}$  less than  $1 \times 10^{-5}$  between successive iterations. Algorithm 1 summarizes the training process used for all training regimens.

---

**Algorithm 1** Training process

---

- 1: Restart  $\mathbf{U}$  from unmodified Navier–Stokes solution
  - 2: Randomly initialize DL network parameters  $\theta^0$
  - 3: Calculate loss of unmodified solution  $J_0$
  - 4: Set  $\epsilon_{\text{rel}}^0 = 1.0$
  - 5: **while**  $\Delta\epsilon_{\text{rel}}^n \geq 10^{-5}$  **do**
  - 6:   Converge solution to steady-state using damped Newton solver III.C
  - 7:   Solve adjoint equation (16) for  $\hat{\mathbf{u}}$
  - 8:   Compute  $\nabla_{\theta} J(\mathbf{u}_p(\theta^n))$  using (19)
  - 9:   Update parameters  $\theta$  using gradient descent (20)
  - 10:   Calculate  $\epsilon_{\text{rel}}^{n+1}$
  - 11:   Compute  $\Delta\epsilon_{\text{rel}}^{n+1} = \epsilon_{\text{rel}}^{n+1} - \epsilon_{\text{rel}}^n$
  - 12: **end while**
- 

Table 2 shows the total number of optimization iterations and the average number of Newton iterations per optimization iteration for all models. For all training cases, the relative tolerance of the forward Newton solver was  $1 \times 10^{-15}$  before proceeding to optimization. For both Approach A and Approach B, the higher Mach number cases required more Newton iterations per optimization iteration. Approach A resulted in better training accuracy (lower  $\epsilon_{\text{rel}}$ ) for most training cases.

Training	$N_{\text{iter}}^{\text{opt}}$	$N_{\text{iter}}^{\text{Newt}}$	$\epsilon_{\text{rel}}$
<i>Approach A</i>			
$M_{\text{train}} = 2$	21	8	$5.333 \times 10^{-2}$
$M_{\text{train}} = 5$	19	20	$1.774 \times 10^{-2}$
$M_{\text{train}} = 8$	21	35	$1.069 \times 10^{-2}$
$M_{\text{train}} = (2, 5, 8)$	(100, 100, 100)	(9, 36, 17)	$(6.715 \times 10^{-1}, 9.487 \times 10^{-2}, 6.099 \times 10^{-2})$
<i>Approach B</i>			
$M_{\text{train}} = 2$	46	9	$2.876 \times 10^{-2}$
$M_{\text{train}} = 5$	37	19	$4.855 \times 10^{-2}$
$M_{\text{train}} = 8$	22	41	$7.334 \times 10^{-2}$
$M_{\text{train}} = (2, 5, 8)$	(100, 100, 100)	(6, 22, 33)	$(4.264 \times 10^{-1}, 7.516 \times 10^{-2}, 5.397 \times 10^{-2})$

**Table 2** Number of optimization iterations,  $N_{\text{iter}}^{\text{opt}}$ , average number of Newton iterations per optimization iteration  $N_{\text{iter}}^{\text{Newt}}$ , and training loss (objective-function relative error  $\epsilon_{\text{rel}}$ ) for Approaches A and B. For  $M_{\text{train}} = (2, 5, 8)$ , the training loss is listed separately for the three target Mach numbers.

## V. Results and Analysis

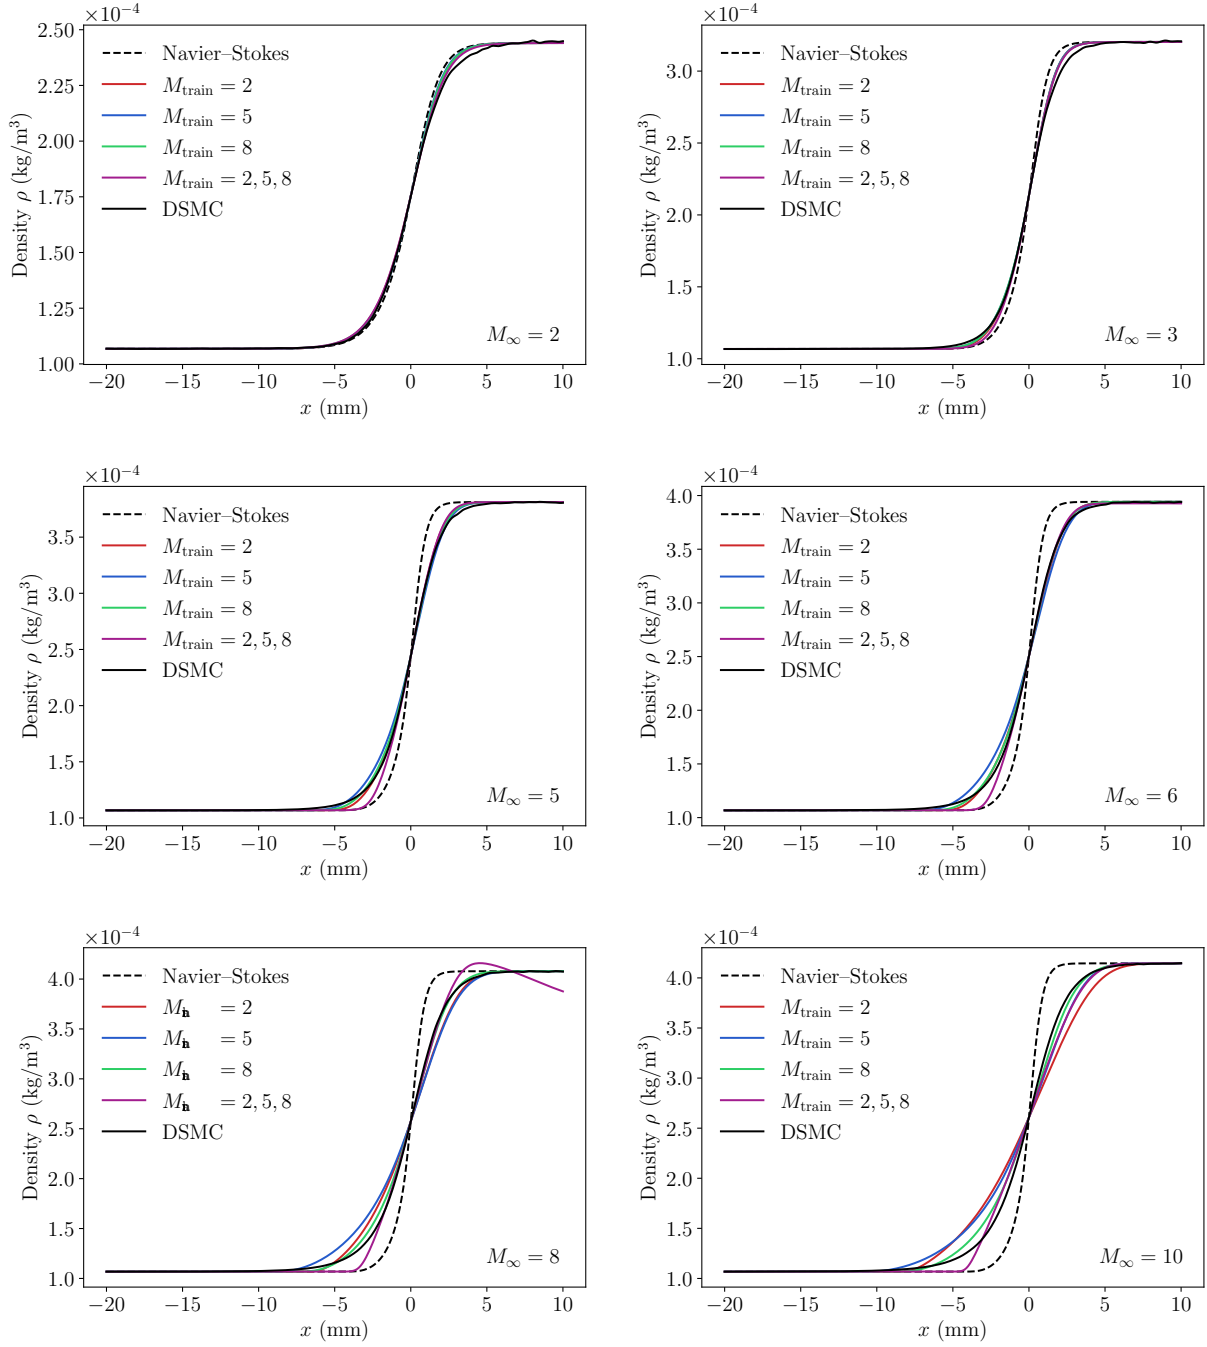
### A. Shock Predictions

Out-of-sample results were obtained for the testing cases listed in Table 1. Without a closure model, the Navier–Stokes equations give qualitatively incorrect predictions for all training and testing cases. Figures 2 and 3 (Approaches A and B, respectively) show the computed density from the unmodified Navier–Stokes equations, the Navier–Stokes solutions with the trained neural networks, and the target DSMC solutions for  $M_{\infty} \in [2, 10]$ . The neural network-augmented Navier–Stokes solutions more closely match the target data for all cases and models. As could be anticipated from the training loss (Table 2), the DL-augmented Navier–Stokes agreement with the DSMC data is qualitatively better for lower Mach numbers and for Mach numbers closer to each model’s training conditions. It is also noteworthy that, for both approaches, the models trained on higher Mach numbers extrapolate better to lower Mach numbers. Similar trends were obtained for the pressure, temperature, and velocity fields.


### B. Prediction Accuracy

Figure 4 shows the relative loss-function error as a function of the testing Mach number. This represents the integrated error of the primitive variables over the entire domain and can be interpreted as a worst-case assessment of a model’s performance. All models trained for a single Mach number using Approach A perform better (lower  $\epsilon_{\text{rel}}$ ) than the corresponding models using the Approach B. For Approach B, the model trained for  $M_{\text{train}} = 2$  is generally more accurate at lower than higher testing Mach numbers, and the model trained for  $M_{\text{train}} = 8$  is generally more accurate at higher than lower testing Mach numbers, as could be anticipated. The Approach B model trained simultaneously for  $M_{\text{train}} = (2, 5, 8)$  is reasonably accurate across the range of testing Mach numbers, with comparable accuracy to the in-sample Mach 2, 5, and 8 models, though its extrapolation accuracy diminishes at  $M_{\infty} = 9, 10$ . Interestingly, the single-case  $M_{\text{train}} = 8$  model reached the best in-sample local minimum and gave the lowest relative errors for lower testing Mach numbers. This model is also reasonably accurate for higher testing Mach numbers.


The extrapolation trends of Approach A are similar to those of Approach B. The model trained for  $M_{\text{train}} = 8$  extrapolates well to lower Mach number cases, with comparable accuracy to the  $M_{\text{train}} = 2$  case. However, the model trained simultaneously for  $M_{\text{train}} = (2, 5, 8)$  performs poorly for this case with errors higher than all other models for almost all Mach numbers. The model trained simultaneously for for  $M_{\text{train}} = (2, 5, 8)$  is reasonably accurate for  $M_{\infty} = 2$  through 8 but has increased errors when extrapolating to the higher  $M_{\infty} = 9$  and 10.



**Fig. 2 Approach A: In- and out-of-sample freestream density predictions. Results are shown for the unmodified Navier–Stokes equations, the Navier–Stokes equations with DL models for Approach A trained for  $M = 2$ ,  $M = 5$ ,  $M = 8$ , and  $M = (2, 5, 8)$ , and the DSMC target data. The  $M = 3$ ,  $M = 6$ , and  $M = 10$  cases are out-of-sample for all trained DL models.**



**Fig. 3 Approach B: In- and out-of-sample freestream density predictions. Results are shown for the unmodified Navier–Stokes equations, the Navier–Stokes equations with DL models for Approach B trained for  $M = 2$ ,  $M = 5$ ,  $M = 8$ , and  $M = (2, 5, 8)$ , and the DSMC target data. The  $M = 3$ ,  $M = 6$ , and  $M = 10$  cases are out-of-sample for all trained DL models.**



**Fig. 4** In- and out-of-sample objective-function relative error for the four DL models trained for **Approach A** (left) and **Approach B** (right). In-sample cases for each model are indicated by red circles. The unmodified Navier–Stokes solution has  $\epsilon_{rel} = 1.0$ .

### C. Shock Thickness

The shock thickness is evaluated using the maximum density gradient,

$$\delta = \frac{\max(\rho) - \min(\rho)}{\max(\partial\rho/\partial x)}. \quad (31)$$

Figure 5 plots the inverse shock thickness normalized by the mean free path [44],

$$\lambda_{\infty} = \frac{16}{5} \left( \frac{\gamma}{2\pi} \right)^{\frac{1}{2}} \frac{\mu_0}{\rho_{\infty} a_{\infty}}, \quad (32)$$

where  $a_{\infty}$  is the freestream speed of sound and  $\lambda_{\infty} = 1.098$  mm for the temperature and pressure conditions considered here [44]. As shown in Fig. 5, all of the DL models improve the Navier–Stokes predictions’ accuracy over the unmodified continuum solutions, which overpredict the shock thickness for all freestream Mach numbers at this very low freestream density. The models trained for Approach A are generally more accurate than those trained for Approach B, which is counterintuitive, as one would expect the greater modeling flexibility of Approach B to result in higher overall accuracy, though this is generally consistent with the  $\epsilon_{rel}$  trends (Fig. 4). The accuracy reduction for Approach B is likely due to inconsistencies in the Clausius–Duhem inequality when applied as a constraint for subcontinuum flows. Conversely, for Approach A, the simple positivity constraint on  $\mu$  and  $k$  is a stricter requirement but does not necessarily require the validity of the Chapman–Enskog expansion of the Boltzmann equation. For both Approaches A and B, the models trained simultaneously for  $M_{train} = (2, 5, 8)$  is reasonably accurate for shock-thickness predictions, especially at their in-sample conditions. The Approach A model trained for  $M_{train} = 8$  gives the most accurate inverse shock thickness predictions over the entire range of testing Mach numbers and is used for subsequent comparisons.

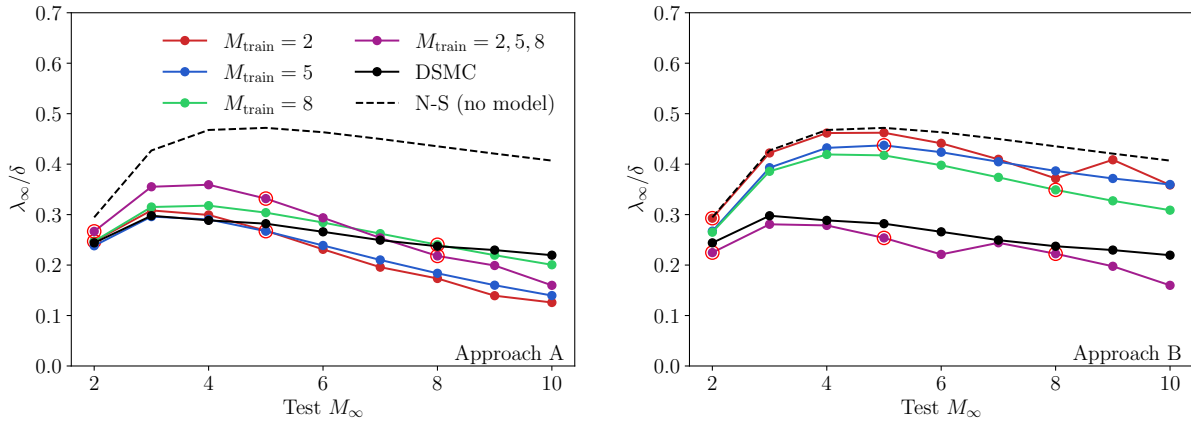
Figure 6 compares the shock-thickness predictions of the Approach A,  $M_{train} = 8$  model to the DSMC targets and experimental data [44, 52–55]. The DL model follows the DSMC trends to within the experimental uncertainty. Additionally, Fig. 6 extends the testing range below  $M_{\infty} = 2$ , at which conditions the DL-augmented Navier–Stokes, unmodified Navier–Stokes, DSMC, and experimental data converge.

### D. Shock Profile Asymmetry

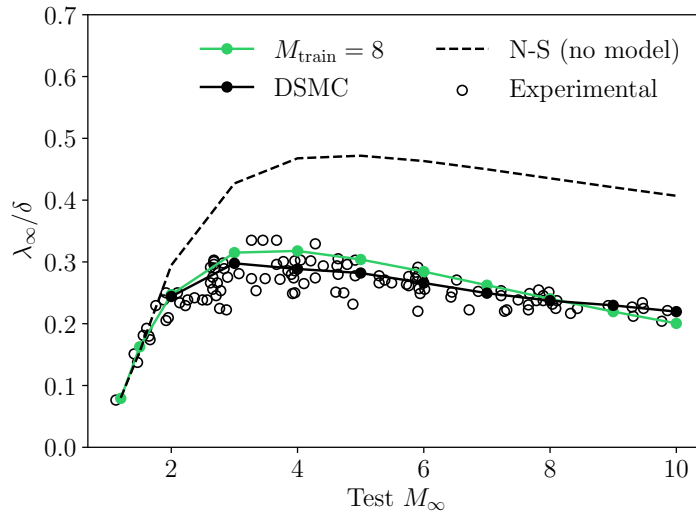
A second derived quantity, which provides a better assessment of the overall shape, is the density asymmetry quotient

$$Q_p = \frac{\int_{-\infty}^0 \rho^* dx}{\int_0^{\infty} [1 - \rho^*] dx}, \quad (33)$$


where  $\rho^*$  is the normalized density, the maximum gradient of which is centered at  $x = 0$ . The asymmetry quotient measures the skewness of a shock density profile relative to its midpoint; a perfectly symmetric profile has  $Q_p = 1$ .



**Fig. 5** Predicted inverse shock thickness from the unmodified (no-model) Navier–Stokes equations, the augmented NS equations for the four DL models, and the DSMC target data for Approach A (left) and Approach B (right).



**Fig. 6** Experimental data from [44, 52–55] compared to predicted inverse shock thickness from the unmodified (no-model) Navier–Stokes equations, the augmented NS equations for  $M_{\text{train}} = (2, 5, 8)$  using Approach A, and the DSMC target data.



**Fig. 7 Predicted density asymmetry quotient for the unmodified (no-model) Navier–Stokes equations, the augmented Navier–Stokes equations with the Approach A closure trained for  $M_{\text{train}} = 8$ , the DSMC target data, and experimental data [44].**

	Testing $M_\infty$								
	2	3	4	5	6	7	8	9	10
$N_{\text{iter}}^{\text{opt}} = 100$	0.026	0.071	0.117	0.148	0.156	0.199	0.199	0.205	0.250
$N_{\text{iter}}^{\text{opt}} = 300$	0.037	0.066	0.115	0.137	0.135	0.180	0.231	0.266	0.297


**Table 3 Objective-function relative error  $\epsilon_{\text{rel}}$  for  $N_{\text{iter}}^{\text{opt}} = 40$  and  $N_{\text{iter}}^{\text{opt}} = 300$ . Shown for the Approach A model trained for  $M_{\text{train}} = (2, 5, 8)$ .**

Figure 7 compares the computed  $Q_p$  using the unmodified Navier–Stokes equations, the Approach A model trained for  $M_{\text{train}} = 8$ , the DSMC target data, and published experimental data [44]. The experimental data shows downstream-skewed shocks ( $Q_p < 1$ ) for freestream Mach numbers less than 2.5 and upstream-skewed shocks ( $Q_p > 1$ ) for higher Mach numbers, while the unmodified Navier–Stokes equations predict upstream-skewed shocks for all Mach numbers. The DSMC target profiles capture the overall experimental trends but lie between the experimental and unmodified Navier–Stokes profiles; this is consistent with other DSMC results using the VHS model [64–67]. The DL model leads to skewnesses reasonably close to the target DSMC profiles near the training  $M_{\text{train}} = 8$  condition, though these are still qualitatively different from the experimental measurements at higher Mach numbers.

### E. Training Convergence

The convergence of  $\epsilon_{\text{rel}}$  with the number of optimization iterations is shown in Fig. 8 for the Approach A model trained simultaneously for  $M_{\text{train}} = (2, 5, 8)$ . A model trained for  $N_{\text{iter}}^{\text{opt}} = 100$  iterations was used for the preceding tests, as it had the lowest total  $\epsilon_{\text{rel}}$  between the three training Mach numbers. It can be seen that, as the model was trained for more iterations, its predictions became more accurate for the intermediate  $M_\infty = 5$  and less accurate for  $M_\infty = 2$  and 8. The testing errors of this model for all testing Mach numbers are listed in Table 3 for  $N_{\text{iter}}^{\text{opt}} = 100$  and  $N_{\text{iter}}^{\text{opt}} = 300$ . It is evident that the prediction error for the longer-trained model is significantly higher for Mach numbers away from  $M_\infty = 5$ , while interpolation to Mach numbers close to it are marginally lower. This is possibly evidence of the model overfitting to  $M_\infty = 5$ .





**Fig. 8** Convergence of  $\epsilon_{\text{rel}}$  (left) and the sum of  $\epsilon_{\text{rel}}$  for the three training Mach numbers (right) as a function of the number of optimization iterations  $N_{\text{iter}}^{\text{opt}}$  using Approach A, for the training regimen  $M_{\text{train}} = (2, 5, 8)$ . The trained model at  $N_{\text{iter}}^{\text{opt}} = 100$  (---) is used for in- and out-of-sample prediction.


### F. Physical Interpretation of the Closure

Figures 9 and 10 for Approaches A and B, respectively, compare the DL closure models to the unmodified continuum closures and the corresponding closures integrated from the DSMC data [68]. Approaches A and B modify the viscous stress and heat flux to similar extents relative to the standard continuum models in the vicinity of the shock and are in excellent agreement with the DSMC profiles, particularly upstream and in the vicinity of the shock. Overall, the standard models underpredict molecular transport at these nonequilibrium conditions, especially the post-shock magnitude of the transport terms at higher freestream Mach numbers. Near the shock, the DL closure models and DSMC profiles have close magnitudes and are spatially more diffuse than the standard continuum models. This is consistent with the resulting steeper shock profiles (Fig. 6).


Fig. 9, for Approach A, shows that the modified heat-flux and viscous-stress terms have the same slope away from the shock as the DSMC terms, unlike for Approach B. This is not surprising, as Approach B augments the continuum closures with nonlinear terms which, are functions of the (nearly zero) flow gradients away from the shock. Significantly, the DL-predicted profiles for  $\sigma$  and  $\mathbf{q}$  differ, which reflects the fact that viscosity and thermal conductivity are continuum manifestations of different sub-continuum processes. Conversely, the temperature-dependent models have the same functional form (7) with only slightly different coefficients. (Indeed, a common approximation for  $\lambda(T)$  is based upon multiplying  $\mu(T)$  by a constant Prandtl number.) This modeling choice renders the standard models unable to account for the differing degrees of momentum and thermal diffusion at transition-continuum conditions. The Approach A and B models have comparable maxima to the DSMC profiles, but the Approach B models generally underpredict  $\sigma$  and  $\mathbf{q}$  downstream of the shock, which is consistent with their relatively inaccurate shock-thickness predictions. Furthermore, for both Approaches A and B, the modeled  $\tau$  and  $\mathbf{q}$  are more asymmetric (upstream-skewed) than the DSMC, particularly for high Mach numbers, which leads to asymmetry quotients  $Q_p > 1$  (see Fig. 7).

Approach A (4) modifies the viscosity and thermal conductivity; it is therefore instructive to consider the form of the model. Fig. 11 compares the Approach A closures for  $\mu(\mathbf{U}; \theta)$  and  $\lambda(\mathbf{U}; \theta)$  (3), evaluated using the neural network trained for  $M_{\text{train}} = (2, 5, 8)$ , to the standard, temperature-dependent models (7) for the range of in- and out-of-sample testing Mach numbers. The differences between the two models are pronounced at higher  $M_{\infty}$  and both models predict significantly higher maximum values than the temperature-dependent models. It is also evident that downstream of the shock, the DL model predicted transport coefficients almost exactly match temperature dependent models, thus it accurately recovers continuum behavior away from the shock.

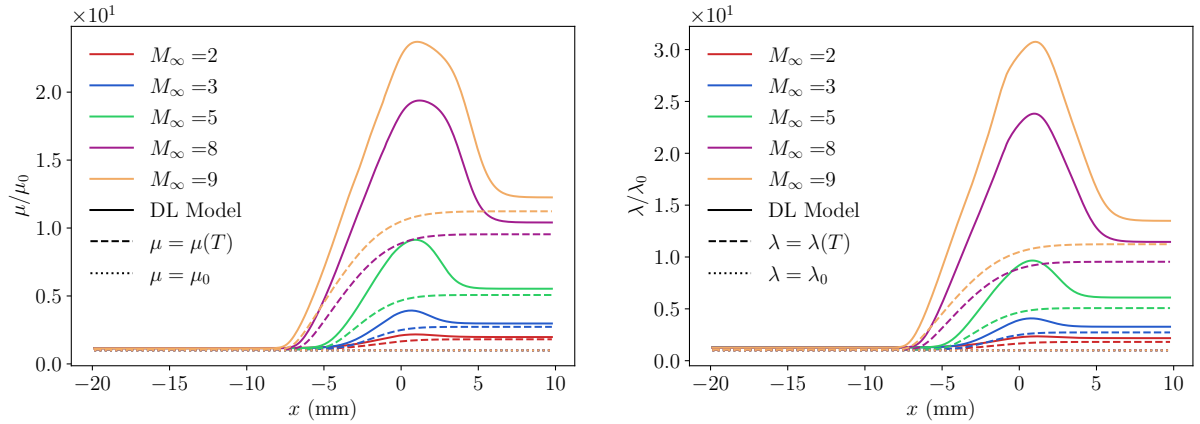
The general trends captured by the DL models could be incorporated into algebraic models similar to (7), which would enable computationally efficient predictions on a wide range of out-of-sample Mach numbers and geometric configurations, even with higher dimensionality than the present shock configurations. This is more probable for simpler models, such as our Approach A, that interact only with a scalar coefficient, and so could be more portable to different dimensionalities, flow geometries, and/or alignments of key features (e.g., shocks) with the computational mesh. The more complicated models that introduce vector/tensor modifications, such as Approach B, are more likely to train in a



**Fig. 9** DL-modeled viscous stress and heat flux (—) for Approach A versus the standard continuum models (---). Shown for the Approach A DL model trained for  $M_{\text{train}} = 8$ .



**Fig. 10** DL-modeled viscous stress and heat flux (—) for Approach B versus the standard continuum models (---). Shown for the Approach B DL model trained for  $M_{\text{train}} = 8$ .



**Fig. 11** DL-modeled viscosity and thermal conductivity (—; Eq. 3) versus standard temperature-dependent models (---; Eq. 7). Shown for the Approach A (4) DL model trained for  $M_{\text{train}} = 8$ .

geometry-specific manner and likely require re-training for new geometries. The development and validity of more generally applicable models, as well as the extrapolability of trained models to different freestream thermodynamic conditions and gas compositions, will be the subject of future work.

## VI. Conclusion

This paper presents an adjoint-based deep learning augmentation method to extend the validity of the continuum flow equations to the transition-continuum regime. Two modeling approaches are compared: Approach A modifies only the viscosity and thermal conductivity, while Approach B directly augments the continuum diffusive fluxes with corrective neural networks. These modeling approaches have less flexibility than outright replacement of the fluxes but were more stable during optimization. Model training is performed by optimizing over the Navier–Stokes equations while targeting trusted flow data from DSMC solutions of the Boltzmann equation.

An online adjoint optimization method enables this PDE-constrained optimization to be performed efficiently. The trained models improve the accuracy of Navier–Stokes predictions in this regime and exhibit stability and extrapolability for Mach numbers higher and lower than those used for training, with concomitant reduction of accuracy when used far out-of-sample (although still not of lower accuracy than that of the unmodified Navier–Stokes equations). Models trained using Approach A have marginally better accuracy compared to models trained using Approach B, and part of Approach A’s success may be attributed to the distinctions it identifies between the different transport-coefficient profiles and their different growth trends with respect to the upstream Mach number. This results in different profiles for the DL-predicted viscous-stress and heat-flux closures.

Entropy constraints ensure that the learned closure models satisfy the second law of thermodynamics. The strong entropy constraint, which poses an algebraic constraint on the output of the DL model, resulted in stable, accurate predictions over the range of testing Mach numbers. The weak entropy constraint, which penalizes the violation of the entropy inequality in the optimization loss function, is attractive for its flexibility but was less stable extrapolating to higher Mach numbers. Further study of the weighting factor for the entropy violation term is necessary to realize the full potential of the weak entropy condition.

Notably, the loss function for the online adjoint optimization can be arbitrarily defined, can include flow variables (the PDE solution or derived quantities), and does not need to explicitly include the unclosed term(s). This is in contrast to *a priori* training, which conforms the model directly to the unclosed terms. The success of the online adjoint method with “third-party” DSMC target data (i.e., not directly obtained from the flow PDEs) is striking, though not unexpected, given the success of adjoint-based methods for aerodynamic shape design and turbulence modeling, among others. One hope is that this method will enable model development against experimental target data with comparable stability and accuracy, though this remains to be tested. Alternatively, or in addition, the loss function for the flow model can target derived or integral quantities that are functions of the flow variables, such as the shock thickness or, for wall-bounded flows, the boundary layer thickness and drag coefficient.

Future work will focus on extending the current model to more-complex geometric scenarios (2D oblique and curved shocks, 3D shocks) and flows with additional physical complexity (nonmonatomic gases, reacting flows). Particularly useful for cases with additional physical complexity, the adjoint-based DL approach can account for multiple, coupled constitutive models simultaneously, which would enable, for example, coupling the transport coefficients with the chemical kinetics.

## Acknowledgments

The authors would like to thank Prof. Iain Boyd for sharing his DSMC solver. This work was supported by the Office of Naval Research under award N00014-22-1-2441. The research of M.P. is supported by the Vannevar Bush Faculty Fellowship. The authors acknowledge computational time on resources supported by the University of Notre Dame Center for Research Computing (CRC).

## References

- [1] Paolucci, S., and Paolucci, C., “A second-order continuum theory of fluids,” *Journal of Fluid Mechanics*, Vol. 846, No. 22, 2018, p. 686–710. <https://doi.org/10.1017/jfm.2018.291>.
- [2] Bird, G. A., “Approach to translational equilibrium in a rigid sphere gas,” *Physics of Fluids*, Vol. 6, 1963, p. 1518. <https://doi.org/10.1063/1.1710976>.

- [3] Bird, G. A., “Direct simulation of the Boltzmann equation,” *Physics of Fluids*, Vol. 13, 1970, p. 2676–2681. <https://doi.org/10.1063/1.1692849>.
- [4] Boyd, I. D., “Modeling of associative ionization reactions in hypersonic rarefied flows,” *Physics of Fluids*, Vol. 19, No. 9, 2007, p. 096102. <https://doi.org/10.1063/1.2771662>.
- [5] Boyd, I. D., “Computation of Hypersonic Flows Using the Direct Simulation Monte Carlo Method,” *Journal of Spacecraft and Rockets*, Vol. 52, No. 1, 2015, pp. 38–53. <https://doi.org/10.2514/1.A32767>.
- [6] Chapman, S., and Cowling, T. G., *The Mathematical Theory of Non-Uniform Gases*, Cambridge University Press, 1970.
- [7] Ferziger, J. H., and Kaper, H. G., *Mathematical Theory of Transport Processes in Gases*, Elsevier Science Publishing, 1976.
- [8] Levermore, C. D., “Moment closure hierarchies for kinetic theories,” *Journal of Statistical Physics*, Vol. 83, No. 5-6, 1996, pp. 1021–1065. <https://doi.org/10.1007/BF02179552>.
- [9] Agarwal, R. K., Yun, K.-Y., and Balakrishnan, R., “Beyond Navier–Stokes: Burnett equations for flows in the continuum transition regime,” *Physics of Fluids*, Vol. 13, 2001, p. 3061–3085. <https://doi.org/10.1063/1.1397256>.
- [10] Woods, L. C., *An introduction to the kinetic theory of gases and magnetoplasmas*, Oxford University Press, 1993.
- [11] Grad, H., “On the kinetic theory of rarefied gases,” *Communications on Pure and Applied Mathematics*, Vol. 2, No. 4, 1949, p. 331–407. <https://doi.org/10.1002/cpa.3160020403>.
- [12] Grad, H., “The profile of a steady plane shock wave,” *Communications on Pure and Applied Mathematics*, Vol. 5, No. 3, 1952, p. 257–300. <https://doi.org/10.1002/cpa.3160050304>.
- [13] Jin, S., and Slemrod, M., “Regularization of the Burnett equations via relaxation,” *Journal of Statistical Physics*, Vol. 103, No. 5-6, 2001, p. 1009–1033. <https://doi.org/10.1023/A:1010365123288>.
- [14] Struchtrup, H., and Torrilhin, M., “Regularization of Grad’s 13 moment equations: derivation and linear analysis,” *Physics of Fluids*, Vol. 15, No. 9, 2003, p. 2668–2680. <https://doi.org/10.1063/1.1597472>.
- [15] Hash, D. B., and Hassan, H. A., “Assessment of schemes for coupling Monte Carlo and Navier-Stokes solution methods,” *Journal of Thermophysics and Heat Transfer*, Vol. 10, No. 2, 1996, pp. 242–249. <https://doi.org/10.2514/3.781>.
- [16] Aktas, O., and Aluru, N., “A Combined Continuum/DSMC Technique for Multiscale Analysis of Microfluidic Filters,” *Journal of Computational Physics*, Vol. 178, No. 2, 2002, pp. 342–372. <https://doi.org/10.1006/jcph.2002.7030>.
- [17] Panesi, M., Magin, T. E., Bourdon, A., Bultel, A., and Chazot, O., “Electronic Excitation of Atoms and Molecules for the FIRE II Flight Experiment,” *Journal of Thermophysics and Heat Transfer*, Vol. 25, No. 3, 2011, pp. 361–374. <https://doi.org/10.2514/1.50033>.
- [18] Stephani, K., Goldstein, D., and Varghese, P., “Consistent treatment of transport properties for five-species air direct simulation Monte Carlo/Navier-Stokes applications,” *Physics of Fluids*, Vol. 24, No. 7, 2012, p. 077101. <https://doi.org/10.1063/1.4729610>.
- [19] Stephani, K., Goldstein, D., and Varghese, P., “A non-equilibrium surface reservoir approach for hybrid DSMC/Navier–Stokes particle generation,” *Journal of Computational Physics*, Vol. 232, No. 1, 2013, pp. 468–481. <https://doi.org/10.1016/j.jcp.2012.08.017>.
- [20] Schwartzenuber, T. E., and Boyd, I. D., “A hybrid particle-continuum method applied to shock waves,” *Journal of Computational Physics*, Vol. 215, No. 2, 2006, pp. 402–416. <https://doi.org/10.1016/j.jcp.2005.10.023>.
- [21] Schwartzenuber, T. E., Scalabrin, L. C., and Boyd, I. D., “Hybrid particle-continuum simulations of nonequilibrium hypersonic blunt-body flowfields,” *Journal of Thermophysics and Heat Transfer*, Vol. 22, No. 1, 2008, pp. 29–37. <https://doi.org/10.2514/1.30216>.
- [22] Schwartzenuber, T. E., Scalabrin, L. C., and Boyd, I. D., “Hybrid particle-continuum simulations of hypersonic flow over a hollow-cylinder-flare geometry,” *AIAA Journal*, Vol. 46, No. 8, 2008, pp. 2086–2095. <https://doi.org/10.2514/1.36681>.
- [23] Schwartzenuber, T. E., Scalabrin, L. C., and Boyd, I. D., “A modular particle–continuum numerical method for hypersonic non-equilibrium gas flows,” *Journal of Computational Physics*, Vol. 225, No. 1, 2007, pp. 1159–1174. <https://doi.org/10.1016/j.jcp.2007.01.022>.

- [24] Burt, J. M., and Boyd, I. D., “A hybrid particle approach for continuum and rarefied flow simulation,” *Journal of Computational Physics*, Vol. 228, No. 2, 2009, pp. 460–475. <https://doi.org/10.1016/j.jcp.2008.09.022>.
- [25] Munafó, A., Torres, E., Haack, J., Gamba, I., and Magin, T., “Investigation of nonequilibrium effects across normal shock waves by means of a spectral-Lagrangian Boltzmann solver,” *51st AIAA Aerospace Sciences Meeting*, 2013, p. 0305. <https://doi.org/10.2514/6.2013-305>.
- [26] Brenner, M. P., Eldredge, J. D., and Freund, J. B., “Perspective on machine learning for advancing fluid mechanics,” *Physical Review Fluids*, Vol. 4, No. 10, 2019, p. 100501. <https://doi.org/10.1103/PhysRevFluids.4.100501>.
- [27] Duraisamy, K., Iaccarino, G., and Xiao, H., “Turbulence modeling in the age of data,” *Annual Review of Fluid Mechanics*, Vol. 51, 2019, pp. 357–377. <https://doi.org/10.1146/annurev-fluid-010518-040547>.
- [28] Duraisamy, K., “Perspectives on machine learning-augmented Reynolds-averaged and large eddy simulation models of turbulence,” *Physical Review Fluids*, Vol. 6, No. 5, 2021, p. 050504. <https://doi.org/10.1103/physrevfluids.6.050504>.
- [29] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., and Yang, L., “Physics-informed machine learning,” *Nature Reviews Physics*, 2021, pp. 1–19. <https://doi.org/10.1038/s42254-021-00314-5>.
- [30] Kochkov, D., Smith, J., Alieva, A., Wang, Q., Brenner, M., and Hoyer, S., “Machine learning-accelerated computational fluid dynamics,” *Proceedings of the National Academy of Sciences*, Vol. 118, No. 21, 2021, p. 2101784118. <https://doi.org/10.1073/pnas.2101784118>.
- [31] Ling, J., Jones, R., and Templeton, J., “Machine learning strategies for systems with invariance properties,” *Journal of Computational Physics*, Vol. 318, 2016, pp. 22–35. <https://doi.org/10.1016/j.jcp.2016.05.003>.
- [32] Wu, J.-L., Xiao, H., and Paterson, E., “Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework,” *Physical Review Fluids*, Vol. 3, No. 7, 2018, p. 074602. <https://doi.org/10.1103/PhysRevFluids.3.074602>.
- [33] Beetham, S., Fox, R. O., and Capecelatro, J., “Sparse identification of multiphase turbulence closures for coupled fluid–particle flows,” *Journal of Fluid Mechanics*, Vol. 914, 2021, p. A11. <https://doi.org/10.1017/jfm.2021.53>.
- [34] Sirignano, J., MacArt, J. F., and Freund, J. B., “DPM: A deep learning PDE augmentation method with application to large-eddy simulation,” *Journal of Computational Physics*, Vol. 423, 2020, p. 109811. <https://doi.org/10.1016/j.jcp.2020.109811>.
- [35] Holland, J. R., Baeder, J. D., and Duraisamy, K., “Towards integrated field inversion and machine learning with embedded neural networks for RANS modeling,” *AIAA Scitech 2019 Forum*, 2019, p. 1884. <https://doi.org/10.2514/6.2019-1884>.
- [36] MacArt, J. F., Sirignano, J., and Freund, J. B., “Embedded training of neural-network subgrid-scale turbulence models,” *Physical Review Fluids*, Vol. 6, 2021, p. 050502. <https://doi.org/10.1103/physrevfluids.6.050502>.
- [37] Sirignano, J., and MacArt, J. F., “Deep Learning Closure Models for Large-Eddy Simulation of Flows Around Bluff Bodies,” *Journal of Fluid Mechanics*, 2023 (accepted). <https://doi.org/10.48550/arXiv.2208.03498>.
- [38] Sirignano, J., MacArt, J. F., and Spiliopoulos, K., “PDE-constrained models with neural network terms: Optimization and global convergence,” *Journal of Computational Physics*, 2023, p. 112016. <https://doi.org/10.1016/j.jcp.2023.112016>.
- [39] Blonigan, P. J., and Wang, Q., “Multiple shooting shadowing for sensitivity analysis of chaotic dynamical systems,” *Journal of Computational Physics*, Vol. 354, 2018, pp. 447–475. <https://doi.org/10.1016/j.jcp.2017.10.032>.
- [40] Ni, A., Wang, Q., Fernandez, P., and Talnikar, C., “Sensitivity analysis on chaotic dynamical systems by finite difference non-intrusive least squares shadowing (FD-NILSS),” *Journal of Computational Physics*, Vol. 394, 2019, pp. 615–631. <https://doi.org/10.1016/j.jcp.2019.06.004>.
- [41] Taghizadeh, S., Witherden, F. D., and Girimaji, S. S., “Turbulence closure modeling with data-driven techniques: physical compatibility and consistency considerations,” *New Journal of Physics*, Vol. 22, No. 9, 2020, p. 093023. <https://doi.org/10.1088/1367-2630/abadb3>.
- [42] Zhao, Y., Akolekar, H. D., Weatheritt, J., Michelassi, V., and Sandberg, R. D., “RANS turbulence model development using CFD-driven machine learning,” *Journal of Computational Physics*, Vol. 411, 2020, p. 109413. <https://doi.org/10.1016/j.jcp.2020.109413>.

- [43] Liu, W., Fang, J., Rolfo, S., Moulinec, C., and Emerson, D. R., “An iterative machine-learning framework for RANS turbulence modeling,” *International Journal of Heat and Fluid Flow*, Vol. 90, 2021, p. 108822. <https://doi.org/10.1016/j.ijheatfluidflow.2021.108822>.
- [44] Alsmeyer, H., “Density profiles in argon and nitrogen shock waves measured by the absorption of an electron beam,” *Journal of Fluid Mechanics*, Vol. 74, No. 3, 1976, pp. 497–513. <https://doi.org/10.1017/S0022112076001912>.
- [45] Reese, J. M., Woods, L. C., Thivet, F. J. P., and Candel, S. M., “A second-order description of shock structure,” *Journal of Computational Physics*, Vol. 117, No. 2, 1995, pp. 240–250. <https://doi.org/10.1006/jcph.1995.1062>.
- [46] Lumpkin, F. E., and Chapman, D. R., “Accuracy of the Burnett equations for hypersonic real gas flows,” *AIAA Journal*, Vol. 6, No. 3, 1992, pp. 419–425. <https://doi.org/10.2514/3.377>.
- [47] Muckenfuss, C., “Some aspects of shock structure according to the bimodal mode,” *Physics of Fluids*, Vol. 5, No. 9, 1962, pp. 1325–1336. <https://doi.org/10.1063/1.1706528>.
- [48] Levermore, C. D., and Morokoff, W. J., “The Gaussian moment closure for gas dynamics,” *Journal of Applied Mathematics*, Vol. 59, No. 1, 1998, pp. 72–96. <https://doi.org/10.1137/S0036139996299236>.
- [49] Cercignani, C., Frezzotti, A., and Grosfils, P., “The structure of an infinitely strong shock wave,” *Physics of Fluids*, Vol. 11, No. 9, 1994, pp. 2757–2764. <https://doi.org/10.1063/1.870134>.
- [50] Myong, R., “A Computational Method for Eu’s Generalized Hydrodynamic Equations of Rarefied and Microscale Gasdynamics,” *Journal of Computational Physics*, Vol. 168, No. 1, 2001, pp. 47–72. <https://doi.org/10.1006/jcph.2000.6678>.
- [51] Xu, K., and Tang, L., “Nonequilibrium Bhatnagar–Gross–Krook model for nitrogen shock structure,” *Physics of Fluids*, Vol. 16, No. 10, 2004, pp. 3824–3827. <https://doi.org/10.1063/1.1783372>.
- [52] Linzer, M., and Hornig, D. F., “Structure of Shock Fronts in Argon and Nitrogen,” *Physics of Fluids*, Vol. 6, No. 12, 1963, pp. 1661–1668. <https://doi.org/10.1063/1.1711007>.
- [53] Camac, M., “Argon and nitrogen shock thicknesses,” *Proceedings of the Fourth International Symposium on Rarefied Gas Dynamics*, Vol. 1, 1965, p. 240. <https://doi.org/10.2514/6.1964-35>.
- [54] Schmidt, B., “Electron beam density measurements in shock waves in argon,” *Journal of Fluid Mechanics*, Vol. 39, No. 2, 1969, pp. 361–373. <https://doi.org/10.1017/S0022112069002229>.
- [55] Garen, W., Synofzik, R., and Frohn, A., “Shock Tube for Generating Weak Shock Waves,” *AIAA Journal*, Vol. 12, No. 8, 1974, pp. 1132–1134. <https://doi.org/10.2514/3.49425>.
- [56] Poinso, T. J., and Lele, S. K., “Boundary conditions for direct simulations of compressible viscous flows,” *Journal of Computational Physics*, Vol. 101, No. 1, 1992, pp. 104–129. [https://doi.org/10.1016/0021-9991\(92\)90046-2](https://doi.org/10.1016/0021-9991(92)90046-2).
- [57] Roe, P. L., “Characteristic-based schemes for the Euler equations,” *Annual Review of Fluid Mechanics*, Vol. 18, No. 1, 1986, pp. 337–365. <https://doi.org/10.1146/annurev.fluid.18.1.337>.
- [58] Boyd, I. D., and Schwartzentruber, T. E., *Nonequilibrium Gas Dynamics and Molecular Simulation*, Cambridge Aerospace Series, Cambridge University Press, 2017. <https://doi.org/10.1017/9781139683494>.
- [59] Valentini, P., and Schwartzentruber, T. E., “Large-scale molecular dynamics simulations of normal shock waves in dilute argon,” *Physics of Fluids*, Vol. 21, No. 6, 2009, p. 066101. <https://doi.org/10.1063/1.3149866>.
- [60] Wang, S., “Extensions to the Navier–Stokes equations,” *Physics of Fluids*, Vol. 34, No. 5, 2022, p. 053106. <https://doi.org/10.1063/5.0087550>.
- [61] Sirignano, J., and Spiliopoulos, K., “Online Adjoint Methods for Optimization of PDEs,” *Applied Mathematics and Optimization*, 2022, p. 18. <https://doi.org/10.1007/s00245-022-09852-5>.
- [62] Steiner, B., DeVito, Z., Chintala, S., Gross, S., Paszke, A., Massa, F., Lerer, A., Chanan, G., Lin, Z., Yang, E., Desmaison, A., Tejani, A., Kopf, A., Bradbury, J., Antiga, L., Raison, M., Gimselshin, N., Chilamkurthy, S., Killeen, T., Fang, L., and Bai, J., “PyTorch: An imperative style, high-performance deep learning library,” *NeurIPS*, 2019.
- [63] Powers, J. M., *Combustion Thermodynamics and Dynamics*, Cambridge University Press, 2016.

- [64] Torrilhon, M., and Struchtrup, H., “Regularized 13-moment equations: shock structure calculations and comparison to Burnett models,” *Journal of Fluid Mechanics*, Vol. 513, 2004, pp. 171–198. <https://doi.org/10.1017/S0022112004009917>.
- [65] Bentley, B. I., “An Investigation of Shock Wave Physics via Hybrid CFD-BGK Solution Methods for Nonequilibrium Flows,” Tech. rep., Air Force Inst. Tech., 2009.
- [66] Bentley, B. I., and Greedyke, R., “Using the Unified Flow Solver to Investigate the Normal Shock Wave Structure,” *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, 2009, p. 1541. <https://doi.org/10.2514/6.2009-1541>.
- [67] Bentley, B. I., and Greedyke, R. B., “Shock Structure Solution by Hybrid Flowfield Methods,” *Journal of Spacecraft and Rockets*, Vol. 49, No. 6, 2012, pp. 1058–1068. <https://doi.org/10.2514/1.A32179>.
- [68] Vincenti, W. G., and Kruger, C. H., *Introduction to Physical Gas Dynamics*, New York, Wiley, 1965.