

Data-driven reduction of lookup libraries for the simulation of non-equilibrium effects in hypersonic flows

By C. Scherding, T. Sayadi,[†] C. Williams, M. Di Renzo,[‡] G. Rigas[¶] AND P. J. Schmid^{||}

Accurate simulations of hypersonic flows in thermochemical non-equilibrium (accounting for high-temperature effects) rely on detailed thermochemical gas models. While capable of accurately capturing the underlying aerothermochemistry, these models dramatically increase the cost of such calculations. In this work, we revisit two aspects of a recently proposed model-agnostic machine-learning technique, which extracts a reduced thermochemical model of a gas mixture from a library (Scherding et al. 2022), and thereby reduces the cost of the underlying calculation. The lightweight library is constructed using data from a cooled-wall laminar boundary layer direct numerical simulation (DNS) at $Ma = 10$. A physics-based prediction of the temperature is implemented to obtain a stable model. The training data acquisition problem is then analyzed using a self-similar solution, and the low-dimensional manifold is compared with the laminar DNS benchmark. A similar analysis is then performed using the turbulent boundary layer data of Di Renzo & Urzay (2021).

1. Introduction

Non-equilibrium effects play an important role in the accurate simulation of flows at hypersonic conditions and in the computation of design characteristics, such as transition location or thermal loading. Recent studies have identified these effects as causes of order-one changes in growth rates, response behavior, or sensitivities, even though the corresponding variations in first-order flow statistics have been modest (Fedorov 2011; Marxen et al. 2014; Di Renzo & Urzay 2021; Passiatore et al. 2022). These findings have in turn prompted a significant endeavor to augment existing flow solvers with non-equilibrium modules to account for finite-rate aerothermochemical processes.

Simulations in this parameter regime introduce and track a range of species in their neutral or ionized forms (Candler 2019). Complementing the hydrodynamic state vector by chemical components is a well-established technique, for example, in combustion or atmospheric simulations. However, the required modeling of the interspecies interactions for hypersonic applications, such as dissociation, reaction, and recombination, poses great challenges (Anderson 2000). Much of this modeling is accomplished by lookup libraries, which act as repositories of tabulated chemical reactions encountered for a given flow state (Scoggins et al. 2020). When passing state vector components to the library, ampli-

[†] Institut Jean le Rond d'Alembert, France

[‡] Centre European de Recherche et de Formation Avancée en Calcul Scientifique, France

[¶] Department of Aeronautics, Imperial College London, United Kingdom

^{||} Department of Mechanical Engineering, King Abdullah University of Science and Technology, Saudi Arabia

tudes and timescales for various forcing terms are returned that act as exogeneous inputs to the momentum, energy, and species transport equations.

For aerothermochemical non-equilibrium effects in hypersonic flows, the Mutation++ library (MULTicomponent Thermodynamic And Transport properties for IONized gases in C++), developed and maintained at the von Karman Institute, has become the standard for high-fidelity simulations of high-speed and high-enthalpy flows (Scoggins et al. 2020). It can be coupled to existing flow solvers and model a range of partially ionized gas effects, together with non-equilibrium features, energy exchange processes, and gas-surface interactions. The flexibility and scope of the library come at the expense of a computational bottleneck that significantly slows down a typical large-scale simulation. For this reason, simulations of non-equilibrium flows rank among the most inefficient and laborious computations in fundamental hypersonic research.

Recently, Scherding et al. (2022) proposed a data-driven algorithm to extract any thermochemical model using machine-learning techniques. The method was successfully applied to a $Ma = 10$ laminar adiabatic boundary layer as a proof of concept. However, in the adiabatic regime, temperature and concentrations of dissociation products increase monotonically with decreasing distance from the wall, resulting in a rather simplistic manifold of thermochemical states. Here, the algorithm is tested on a $Ma = 10$ cooled-wall boundary layer, at the same condition studied by Di Renzo & Urzay (2021), and a physics-based improvement to the method is assessed to improve stability. One drawback of the original method is the necessity to generate the training set through a costly DNS, which can hinder its use in speeding up the simulations. Here, we propose a training alternative by leveraging the use of locally self-similar solutions to generate the training set.

This report is organized as follows. In Section 2, the governing equations and thermochemical models are briefly recalled, followed by the description of the test case. In Section 3, the data-driven modeling framework and its physics-based improvement are presented. In Section 4, the application of the algorithm to the cooled-wall boundary layer is showcased, and training using the local similarity solutions and turbulent boundary layer data is assessed. Finally, conclusions are drawn in Section 5.

2. Governing equations

In this work, we consider the reactive compressible Navier-Stokes equations for an air mixture of five species $S = \{N, O, NO, N_2, O_2\}$ given as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (2.1)$$

$$\frac{\partial \rho_s}{\partial t} + \nabla \cdot (\rho_s (\mathbf{u} + \mathbf{V}_s)) = \dot{\omega}_s, \quad s \in S, \quad (2.2)$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u}) = -\nabla P + \nabla \tau, \quad (2.3)$$

$$\frac{\partial \rho e_0}{\partial t} + \nabla \cdot (\rho (e_0 + P) \mathbf{u}) = \nabla \cdot (\tau \cdot \mathbf{u}) - \nabla \cdot \mathbf{q}, \quad (2.4)$$

where the corresponding velocity components are $\mathbf{u} = \{u, v, w\}$, t denotes the time coordinate, ρ denotes the mixture density, and $\rho_s = \rho Y_s$ and Y_s are the partial density and mass fraction of species $s \in S$, respectively. These equations are integrated numerically

for all but one species $s \in S$ in a three-dimensional Cartesian coordinate system $\{x, y, z\}$ with the origin placed at the leading edge of the flat plate, and x, y , and z point in the streamwise, wall-normal, and spanwise directions, respectively. In the momentum equation (Eq. (2.3)), P stands for pressure, and $\boldsymbol{\tau}$ is the viscous stress tensor, defined for a Newtonian fluid as

$$\boldsymbol{\tau} = \mu(\nabla \mathbf{u} + \nabla \mathbf{u}^T - 2(\nabla \cdot \mathbf{u})\mathbf{I}/3), \quad (2.5)$$

where μ is the dynamic viscosity and \mathbf{I} is the identity tensor. In the energy equation (Eq. (2.4)), $e_0 = e + |\mathbf{u}|^2/2$ is the stagnation internal energy, with e denoting the specific internal energy, defined from the species-specific enthalpies h_s as

$$e = \sum_{s \in S} h_s Y_s - P/\rho. \quad (2.6)$$

The heat flux vector takes the form

$$\mathbf{q} = -\kappa \nabla T + \sum_{s \in S} \rho_s h_s \mathbf{V}_s, \quad (2.7)$$

where T denotes the temperature. The term \mathbf{V}_s , appearing in Eqs. (2.2) and (2.7), denotes the diffusion velocity vector of species s . Finally, the net mass production rate of species $\dot{\omega}_s$ in Eq. (2.2), considering all reactions, is computed using Park's (1989) five-reactions chemical mechanism for dissociated air. For a detailed description of the numerical method employed to solve these equations, refer to Margaritis et al. (2022).

2.1. Thermochemical model

The governing equations are closed using the equation of state

$$P = \rho R_u T / \overline{W}, \quad (2.8)$$

where R_u is the universal gas constant. The transport (μ, κ, D_s) , thermodynamic (P, T, h_s) , and chemical $\dot{\omega}_s$ properties defined above are generally a function of two independent thermodynamic state variables and the mixture composition, and need to be modeled accordingly. Different modeling approaches exist in the literature (Di Renzo & Urzay 2021; Passiatore et al. 2022), with the main drawback being that the thermochemical model has to be hard-coded. Hence, any update in the model, for instance, to add complexity or to simulate a different mixture, comes at a human cost in terms of implementation, testing, and validation. These limitations led to the development of the library Mutation++ (Scoggins et al. 2020), which offers a flexible high-level application programming interface to model the physico-chemical properties of mixtures in different levels of non-equilibrium. A wide range of algorithms for the calculation of the individual and mixture-averaged properties are supported. The library is easily coupled to any computational fluid dynamics solver as an input/output problem $\mathbf{z} = f(\mathbf{x})$. More precisely, given the local state vector

$$\mathbf{q}_{\text{th}} = [\rho, \rho_s, \rho e] \in \mathbb{R}^D, \quad (2.9)$$

the library returns all physico-chemical properties needed to close the governing equations

$$\mathbf{z} = f(\mathbf{q}_{\text{th}}) = [P, T, \mu, \kappa, D_s, h_s, \dot{\omega}_s] \in \mathbb{R}^{D_z}. \quad (2.10)$$

2.2. Test case

In this work, the algorithm is tested on the simulation of a two-dimensional, Mach-10, cooled-wall hypersonic boundary layer, based on the conditions studied by Di Renzo




FIGURE 1. Contours of normalized (a) streamwise velocity, and (b) temperature in the laminar cooled-wall Mach-10 boundary layer.

& Urzay (2021). Lengths are non-dimensionalized using the displacement thickness δ_0^* at the inflow, with a corresponding Reynolds number $Re_{\delta_0^*} = 6,000$. The freestream temperature is $T_\infty = 1039$ K, and the wall temperature is set to $T_w = 1700$ K. The computational domain extends from $x = 65\delta_0^*$ to $265\delta_0^*$ in the streamwise direction, with 960 grid points uniformly distributed. In the wall-normal direction, the domain extends from $x = 0$ to $5\delta_0^*$, and 211 grid points are used, with clustering near the wall to properly resolve the large gradients in that region. At the inflow, a locally self-similar solution in chemical non-equilibrium is provided (see appendix A of Di Renzo & Urzay (2021)). At the outflows, numerical sponges damp the solution to a reference unperturbed solution.

Figure 1(a,b) shows, respectively, the streamwise velocity and temperature contours after convergence of the numerical residuals to machine precision. Despite the cooled-wall, aerodynamic heating produces temperatures of the order of $T_{max} \approx 4T_\infty$ at a wall-normal distance of $y = 0.8$, as shown in Figure 1(b). This temperature is sufficiently high to promote dissociation of O_2 into its atomic counterpart, as well as production of NO. However, it is too low to produce significant amounts of atomic N, as shown in Figure 2(c). The temperature then decreases towards the wall after the peak aerodynamic heating, due to wall cooling. Mass fractions of O and NO, on the other hand, stay relatively high, owing to the diffusion of radicals away from the wall.

3. Data-driven model

The goal of the training is to obtain an efficient and accurate surrogate model g of the library f . Let $\mathbf{X} \in \mathbb{R}^{N \times D}$ be a set of N scaled inputs to the library and $\mathbf{Z} \in \mathbb{R}^{N \times D_z}$ the corresponding outputs, both sampled from the simulation presented in Section 2.2. Hence, the goal is to find g , where $\hat{\mathbf{Z}} = g(\mathbf{X})$, and $\|\hat{\mathbf{Z}} - \mathbf{Z}\|$ is minimized. The training of the data-driven model is performed in three steps. In the first step, dimensionality reduction, the input vector \mathbf{X} is encoded into a low-dimensional space that accounts for most of the features of the output space through an input/output encoder (IO-E). The IO-E architecture consists of two sequential deep neural networks. The first network, the encoder, projects the inputs of the library in a latent space of dimension $d < D$. The second network, the decoder, predicts the outputs of the library from this latent space. The training is done through back-propagation of the L_2 norm of the error $\|\hat{\mathbf{Z}} - \mathbf{Z}\|_2$ through the full IO-E, where $\hat{\mathbf{Z}}$ denotes the prediction of the network. However, only the encoder part of the network is used. In fact, the radial basis function networks described later have shown better accuracy than the decoder on the case considered. In the second

step, clustering, the low-dimensional representation of the input vector, $\mathbf{Y} \in \mathbb{R}^{N \times d}$, is clustered using Newman’s algorithm (Newman 2006). This algorithm detects N_c coherent thermodynamic communities within the low-dimensional sub-space. In the third step, surrogate model construction, a radial basis function network is trained with N_R kernels to predict the outputs of the library ($\hat{\mathbf{Z}}$) on the low-dimensional subspace spanned by each cluster.

These steps are chosen to optimize the performance of the algorithm while maintaining the high accuracy needed for high-fidelity simulations of hypersonic flows in chemical non-equilibrium. We refer the reader to Scherding et al. (2022) for a detailed description of the algorithm.

3.1. Physics-based prediction of temperature

In boundary layer flows, pressure is (nearly) constant along the wall-normal direction. This low sensitivity of pressure to the input thermochemical state vector \mathbf{q}_{th} (which drastically varies across the boundary layer) results in a highly accurate pressure prediction \hat{P} through the data-driven model. However, in our thermochemical framework, temperature appears in the equation only through its gradient (Eq. (2.7)), and any disturbance in the model’s prediction will be enhanced through differentiation. To alleviate spurious oscillations, the equation of state (Eq. (2.8)) is used instead to predict the temperature \hat{T} from the predicted pressure \hat{P} and the other model inputs ($\rho_s, s \in S$).

4. Results

In this section, the performance of the algorithms (pre- and post-training) when applied to a Mach-10 cooled-wall boundary layer is assessed. Then, a self-similar solution is used for training and the resulting surrogate is compared with the DNS data. Finally, the turbulent boundary layer data of Di Renzo & Urzay (2021) are used in training, and the results are compared with both laminar DNS and the self-similar solution.

4.1. Cooled-wall Mach-10 DNS data

To train the data-driven surrogate model, $N = 10^5$ thermodynamic state vectors are sampled from the converged solution. The corresponding outputs of the library are collected simultaneously. The model obtained has the following specifications: $d = 2, N_c = 2$ and $N_R = 250$. Moreover, the model performs 50% faster than the original library, drastically decreasing the time to solution of such calculations. Figure 2(d) presents the wall-normal temperature gradient profile. The baseline model (solid cyan line) has a high error with high-frequency content. However, using the physics-based prediction of temperature leads to a smoother gradient. Ultimately, after restarting the simulation with the data-driven model (i.e., in a closed loop), the solution remains stable after four flow-through times using the physics-enhanced version of the model (Figure 2(a-c)).

4.2. Self-similar solution for training

Another drawback of the algorithm as described by Scherding et al. (2022) is the generation of training data using costly DNS calculations. We propose here to leverage locally self-similar solutions in chemical non-equilibrium to generate the training dataset, eliminating the previous computational bottleneck associated with DNS. See Williams et al. (2021) and Lees (1956) for a mathematical and numerical description of these solutions.

Locally self-similar solutions at the same thermodynamic and freestream conditions are generated using the same Mutation++ thermochemical model for a range of Re_x

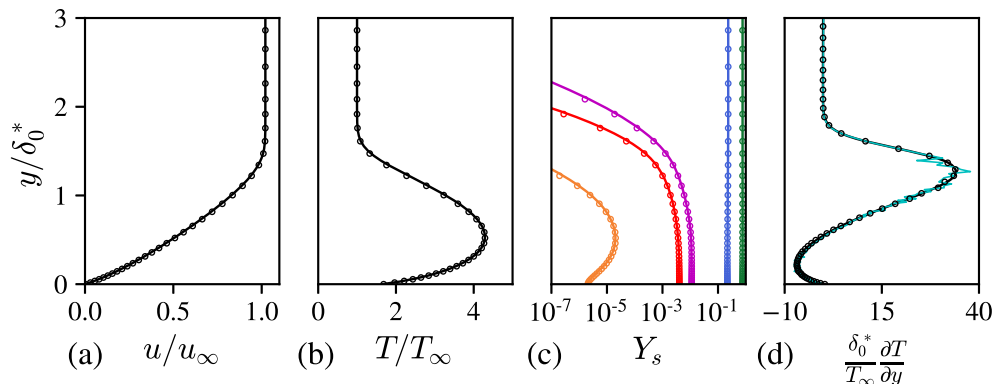


FIGURE 2. Profiles of (a) streamwise velocity, (b) temperature, (c) species mass fractions (from left to right: N, O, NO, N₂, O₂), and (d) temperature gradient at $Re_x = 600,000$. Solid lines and symbols correspond to the solutions generated using Mutation++ and the data-driven model, respectively. In panel (d), the blue curve presents the original model with no physics-based correction.




FIGURE 3. Profiles of (a) streamwise velocity, (b) temperature, and (c) species mass fractions (from left to right: N, O, NO, N₂, O₂), at $Re_x = 600,000$. Solid and dash-dotted lines correspond to the DNS and self-similar solutions, respectively.

representative of the computational domain considered in Section 2.2. Figure 3 displays a comparison of the DNS and self-similar profiles at $Re_x = 600,000$. The self-similar solutions were generated using 0.5 core hours on the Yellowstone cluster, compared with 1,500 core hours needed to obtain the converged laminar DNS. Good agreement is observed for the velocity, temperature, and mass fraction profiles. The most notable difference is in the species mass fractions, where higher dissociation is observed in the self-similar solution than in the DNS. This can be explained by the streamwise species diffusion flux that is neglected in the locally self-similar solution.

Locally self-similar profiles are then interpolated on the numerical grid of the DNS, and input/output pairs are sampled. To assess the viability of using the self-similar solutions as a training basis for the thermochemical model, we compare the distribution of the input thermodynamic states from the DNS with that of the self-similar solutions, using kernel density estimation (KDE) (Rosenblatt 1956). We assume that the training dataset




FIGURE 4. Pairwise kernel density of the selected thermodynamic state vector components ρ , e , and ρ_O : black lines represent laminar DNS; blue lines represents self-similar solutions. Extreme contour lines represent a probability of 1%.

\mathbf{X} consists of N thermodynamic states \mathbf{x}_i , $i = 1, \dots, N$, where $\mathbf{X} \in \mathbb{R}^{N \times D}$ are independent and identically distributed samples drawn from the same multivariate distribution p . The kernel density estimator of p is derived as

$$\hat{p}_h(\mathbf{x}) = \frac{1}{hN} \sum_{i=1}^N K\left(\frac{\|\mathbf{x} - \mathbf{x}_i\|}{h}\right), \quad (4.1)$$

where $K(r) = e^{-r^2}$ denotes the Gaussian kernel and h is the bandwidth. The latter is estimated according to Scott (2015) as $h = N^{-1/(D+4)}$.

Figure 4 shows the pairwise KDE plots of the selected thermodynamic state vector components obtained when sampling points from either the laminar DNS (black) or the self-similar solution (blue). Considering the diagonal plots, both distributions are bimodal along each component of the thermodynamic state vector. This bimodality comes from the freestream and the peak aerodynamic heating regions. Very good agreement is found for density (Figure 4(a)). A slight deviation in the high energy e and high radical's partial density (e.g., ρ_O) is observed. The latter is a direct consequence of the higher level of dissociated species shown in Figure 3. Since temperature distributions are almost identical (Figure 3 (b)), the internal energy increases (Eq. (2.6)), due to higher specific enthalpies of atomic species compared with their molecular counterpart. Looking at the bimodal distribution (Figure 4(b,d,e)), the self-similar distribution (blue) extends to higher e and ρ_O . Overall, the thermodynamic distribution spanned by the laminar flow




FIGURE 5. Contours of normalized (a) streamwise velocity, and (b) temperature at $z = 10\pi\delta_0^*$, in the turbulent cooled-wall Mach 10 boundary layer (Di Renzo & Urzay 2021).

is mostly contained in the region spanned by the locally self-similar solutions. Therefore, the new samples generated in an actual laminar simulation will be within the training set obtained from the self-similar solutions, resulting in reasonable accuracy of the surrogate model. The model could then be fine-tuned later on for the minority of states not seen during training.

4.3. Turbulent DNS

Transition to turbulence is a critical consideration in the design of successful aerospace missions in the hypersonic regime. In fact, skin friction and heat flux at the wall are enhanced in a turbulent boundary layer compared with a laminar one. However, the DNS of turbulent boundary layers requires billions of degrees of freedom (Di Renzo & Urzay 2021; Passiatore *et al.* 2022). The additional computational cost incurred by a complex thermochemical model could be better utilized to study higher Reynolds number hypersonic reacting flows, for instance. Hence, an efficient data-driven thermochemical model would be highly beneficial in these scenarios. Providing training data directly from the turbulent simulations, however, would represent a suboptimal approach. It is thus prudent to investigate the possibility of approximating the manifold (at least partly) using the laminar self-similar solutions in chemical non-equilibrium.

To this end, we consider a fully turbulent section of the boundary layer simulated by Di Renzo & Urzay (2021). The velocity and temperature fields on a streamwise slice (at half the spanwise extent) are plotted in Figure 5(a,b). We see that higher temperatures are present within the boundary layer compared with the laminar solution, partly due to the presence of shocklets within the boundary layer, increasing local temperatures.

N thermodynamic states are sampled out of the domain and an approximation of the probability density function is obtained with KDE (Eq. (4.1)). We compare the distribution from the self-similar (blue) and turbulent (red) solution in Figure 6. The turbulent boundary layer presents a bimodal distribution along each component of the thermodynamic state vector except energy. In fact, the high-energy peak disappears in the turbulent boundary layer due to sweeps and injection of cold air in the peak heating region. Higher energy levels are also observed due to the presence of shocklets. The most notable difference is a shift of the freestream density peak to lower densities. This is the result of a succession of pressure waves emanating from the boundary layer to the freestream as it transitions to turbulence. The self-similar manifold predominantly covers the thermodynamic states seen in the turbulent boundary layer section considered.




FIGURE 6. Pairwise kernel density of selected thermodynamic state vector components ρ , e , and ρ_O : blue lines represent self-similar solutions; red lines represent turbulent DNS. Extreme contour lines represent a probability of 1%.

These results also suggest that locally self-similar solutions could be used to warm-start the training of a data-driven model adapted to hypersonic turbulent boundary layer simulations with finite-rate chemistry.

5. Conclusions

In this work, a data-driven framework to build an efficient and lightweight thermochemical model (Scherding et al. 2022) for simulating hypersonic flows in chemical non-equilibrium has been revisited. Through a physics-based improvement to the temperature prediction, the method has been successfully applied to a Mach-10, cooled-wall boundary layer in chemical non-equilibrium, yielding a 50% CPU time reduction. Self-similar solutions provide a good approximation of the six-component thermodynamic state manifold of the laminar DNS, at a fraction of the cost, thus proving to be a viable substitute for training. Preliminary comparisons between the turbulent and laminar manifolds suggest that the locally self-similar solutions may also serve as an effective method for rapid estimation of thermodynamic manifolds. This study paves the way for further development of the model to adaptively learn new states on the fly, as the simulation progresses. Finally, the algorithm has been tested only in a closed loop on hypersonic boundary layers without shocks. Future work will also consider the application to high-enthalpy shock-wave boundary layer interaction in chemical non-equilibrium.

Acknowledgments

The authors acknowledge the use of computational resources from the Yellowstone cluster awarded by NSF to CTR and the MeSU computing platform at Sorbonne University.

REFERENCES

- ANDERSON, J. D. 2000 *Hypersonic and High Temperature Gas Dynamics*. AIAA.
- CANDLER, G. V. 2019 Rate effects in hypersonic flows. *Annu. Rev. Fluid Mech.* **51**, 379–402.
- DI RENZO, M. & URZAY, J. 2021 Direct numerical simulation of a hypersonic transitional boundary layer at suborbital enthalpies. *J. Fluid Mech.* **912**, A29.
- FEDOROV, A. V. 2011 Transition and stability of high-speed boundary Layers. *Annu. Rev. Fluid Mech.* **43**, 79–95.
- LEES, L. 1956 Laminar heat transfer over blunt-nosed bodies at hypersonic flight speeds. *J. Jet Propul.* **26**, 259–269.
- MARGARITIS, A. T., SCHERDING, C., MARXEN, O., SCHMID, P. J. & SAYADI, T. High-fidelity computational tool for chemically reacting hypersonic flow simulations. *arXiv:2210.05547 [physics.flu-dyn]*.
- MARXEN, O., IACCARINO, G. & MAGIN, T. E. 2014 Direct numerical simulations of hypersonic boundary-layer transition with finite-rate chemistry. *J. Fluid Mech.* **755**, 35–49.
- MCBRIDE, B. J. 2002 *NASA Glenn coefficients for calculating thermodynamic properties of individual species*. Tech. Pub., John H. Glenn Research Center, NASA.
- NEWMAN, M. E. 2006 Modularity and community structure in networks. *P. Natl. Acad. Sci. USA* **103**, 8577–8582.
- PARK, C. 1989 A review of reaction rates in high temperature air. *AIAA Paper 1989*, 1740.
- PASSIATORE, D., SCIACOVELLI, L., CINNELLA, P. & PASCAZIO, G. 2022 Thermochemical non-equilibrium effects in turbulent hypersonic boundary layers. *J. Fluid Mech.* **941**, A21.
- ROSENBLATT, M. 1956 Remarks on some nonparametric estimates of a density function. *Ann. Math. Stat.* **27**, 832–837.
- SCHERDING, C., RIGAS, G., SIPP, D., SCHMID, P. & SAYADI, T. 2022 Data-driven framework for input/output lookup tables reduction — with application to hypersonic flows in chemical non-equilibrium. *arXiv:2210.04269 [physics.flu-dyn]*.
- SCOGGINS, J. B., LEROY, V., BELLAS-CHATZIGEORGIS, G., DIAS, B. & MAGIN, T. E. 2020 Mutation++: multicomponent thermodynamic and transport properties for ionized gases in C++. *SoftwareX* **12**, 100575.
- SCOTT, D. W. 2015 *Multivariate Density Estimation: Theory, Practice, and Visualization*. John Wiley & Sons.
- WILKE, C. 1950 A viscosity equation for gas mixtures. *J. Chem. Phys.* **18**, 517–519.
- WILLIAMS, C., DI RENZO, M., MOIN, P. & URZAY, J. 2021 Locally self-similar formulation for hypersonic laminar boundary layers in thermochemical nonequilibrium. *Annual Research Briefs*, Center for Turbulence Research, Stanford University, pp. 119–128.