

2019 WSSCI Fall Technical Meeting
Organized by the Western States Section of the Combustion Institute
October 14–15, 2019
Albuquerque, New Mexico

Constructing and Accessing Tabulated Chemistry for Fire Scenarios

Elizabeth Armstrong^{1,}, Michael A. Hansen², Robert C. Knaus², John C. Hewson³,
and James C. Sutherland¹*

¹*Department of Chemical Engineering, University of Utah, Salt Lake City, UT*

²*Computational Thermodynamics and Fluid Mechanics, Sandia National Laboratories,
Albuquerque, NM*

³*Fire Science and Technology, Sandia National Laboratories, Albuquerque, NM*

**Corresponding author: elizabeth.armstrong@chemeng.utah.edu*

Abstract: Tabulated chemistry models are widely used in practical simulations of fires with RANS and LES. We address options for constructing and accessing non-adiabatic tabulated chemistry suitable for fire scenarios. In order to capture strong radiative heat losses in fires, we use transient flamelet solutions under strong heat-loss conditions for building structured libraries. There are various techniques for handling the implicit dependence of the enthalpy deficit independent coordinate of the resulting non-adiabatic flamelet table on the mixture fraction coordinate when performing structured table lookups. We compare computational cost and accuracy of these techniques in simulations of laminar fire configurations with substantial radiative heat loss and soot formation. We also investigate the use of machine learning methods for creating and evaluating fully unstructured tables as an alternative for both current and higher dimensional tables of interest.

Keywords: *tabulated chemistry, non-adiabatic flamelets*

1. Introduction

Combustion simulations can be computationally expensive, especially when primitive variable equations are being solved, which can involve hundreds of chemical species participating in thousands of reactions at very small timescales relative to the flow field. Chemistry tabulation methods have been developed to minimize this cost by taking advantage of time-scale separation between fluid mixing rates and the many fast chemical rates. For traditional methods, a range of possible thermochemical states are tabulated *a priori* over a reduced set of reaction coordinates. This reduced set of coordinates is then transported in the simulation in place of the original and used to update the original properties by querying the tables through interpolation.

Flamelet models are frequently used in computing these states for tabulation in large-eddy simulations (LES) and Reynolds-averaged Navier-Stokes (RANS) of combustion processes. In particular, many tabulation methods use the steady laminar flamelet model (SLFM) based on the equations derived by Peters [1]. However, the fire scenarios of interest experience large heat loss, which requires a non-adiabatic flamelet formulation. This can be achieved through adding an enthalpy or enthalpy deficit coordinate to the reduced set of reaction coordinates for tabulation. In this paper, the enthalpy deficit is defined as the difference between the actual enthalpy and the enthalpy of an adiabatic flame. In general, a non-adiabatic flamelet library with enthalpy deficit can be generated

either by modifying the available enthalpy, such as through changing boundary conditions, or by modifying the enthalpy source term in the flamelet equations.

In this first class of enthalpy reduction methods, a common approach is to reduce the boundary temperature to reach the required enthalpy deficit [2–4]. However, this can often lead to unrealistically low temperatures at the boundaries for large heat loss. This can be countered by either shrinking the solution domain in mixture fraction space or changing the composition at the boundaries. It is also possible to simply specify a uniform enthalpy deficit across a laminar flamelet, but this may result in unphysical flame structures.

The second class of enthalpy reduction methods involves reducing the enthalpy in the solution via a heat-loss source term. This class of problems begins by solving an initial adiabatic steady-state with a given scalar dissipation rate. A transient approach then solves the unsteady flamelet equations under strong heat-loss conditions enforced with the added source term. A quasi-steady approach instead solves a series of steady-state problems for a corresponding series of modifications to the source term through, for example, changing a coefficient. In either case, the process is repeated for a range of scalar dissipation rates, leading to an unstructured, non-adiabatic dataset that can be formed into a structured table. Since the quasi-steady approach does not generate an adequate range of heat loss for some fire applications with strong heat losses, this work focuses on the transient approach.

Examples of using radiative source terms in the transient formulation include [5–7]. However, the source terms in these papers are not appropriate for heavily smoking flames where low intensity mixing and strong radiation lead to radiative quenching. Therefore, this paper uses an enthalpy source term linear in temperature, also discussed in [8], which provides a more uniform heat loss as the transient flamelet evolves through heat-loss-driven extinction. This is relevant to fire scenarios where radiative heat losses and interactions with cold surfaces can lead to local temperatures very near ambient. Details of this method are presented in the following sections.

Once the non-adiabatic flamelet states are obtained, they can be organized into structured tables to optimize access speed during runtime. However, in instances where the mixture fraction, scalar dissipation rate, and enthalpy or enthalpy deficit are chosen as coordinates, flamelet properties cannot be directly tabulated in a structured manner over the independent coordinates due to implicit dependencies of these last two coordinates on mixture fraction. Instead, a reference value for the scalar dissipation rate and enthalpy deficit, such as the stoichiometric value, may be used to provide structured lookup over these coordinates in order to generate three independent structured table dimensions. Details of how to then access the table at the correct state given a mixture fraction, scalar dissipation rate, and enthalpy deficit from the simulation, instead of the reference values required as table query inputs, are often left out of papers. Our goal is to shed some light on options for accessing the flamelet states in these structured tables. In this paper, we evaluate the cost and accuracy of various techniques for accessing structured tables in laminar fire simulations with strong radiative heat loss and soot formation. We also investigate the option of using machine learning techniques for creating and evaluating fully unstructured tables, which can offer an advantage over structured tables for higher dimensional problems.

2. Flamelet Library Generation

Non-adiabatic flamelet libraries are generated with Spitfire [9], a Python and C++ code for solving ODEs and diffusion-reaction PDEs with complex chemical kinetics. This library generation involves first saving off steady-state adiabatic flamelet solutions for a range of scalar dissipation rates, followed by saving transient solutions to the unsteady flamelet equations with an additional heat-loss term, as briefly mentioned in the introduction.

The adiabatic flamelet equations Spitfire solves for species mass fractions Y_i and temperature T in Lagrangian time t and mixture fraction Z are

$$\frac{\partial Y_i}{\partial t} = \frac{\chi}{2} \frac{\partial^2 Y_i}{\partial Z^2} + \frac{\omega_i}{\rho} \quad (1)$$

$$\frac{\partial T}{\partial t} = -\underbrace{\frac{1}{\rho c_p} \sum_{i=1}^n \omega_i h_i}_{\text{classical flamelet equations}} + \underbrace{\frac{\chi}{2} \frac{\partial^2 T}{\partial Z^2} + \frac{\chi}{2} \frac{1}{c_p} \frac{\partial T}{\partial Z} \frac{\partial c_p}{\partial Z}}_{\text{variable } c_p} + \underbrace{\frac{\chi}{2} \frac{\partial T}{\partial Z} \sum_{i=1}^n \frac{c_{p,i}}{c_p} \frac{\partial Y_i}{\partial Z}}_{\text{distinct } c_{p,i}} \quad (2)$$

with boundary conditions defined by the oxidizer and fuel states

$$T(t, 0) = T_{\text{oxy}}, \quad Y_i(t, 0) = Y_{i,\text{oxy}}, \quad T(t, 1) = T_{\text{fuel}}, \quad Y_i(t, 1) = Y_{i,\text{fuel}}. \quad (3)$$

In (1) and (2), ω_i is the net mass production rate of species i due to chemical reactions, h_i is the specific internal enthalpy of species i , ρ is the mixture density, c_p is the specific, isobaric heat capacity of the mixture, and χ represents the scalar dissipation rate which depends on mixture fraction as

$$\chi(Z) = \chi_{\max} \exp(-2[\text{erf}^{-1}(2Z-1)]^2) = \chi_{\max} F_{\chi}(Z) \quad (4)$$

for a maximum scalar dissipation rate χ_{\max} . The last two terms in (2) are important for maintaining consistent temperature and enthalpy profiles in the flamelet solution. Equations (1) and (2) are also shown for unity Lewis number, which is appropriate for turbulent hydrocarbon simulations, the main application.

The only difference between the adiabatic and non-adiabatic equations Spitfire solves is in the temperature equation, which for non-adiabatic flamelets has an additional source term:

$$\frac{\partial T}{\partial t} = \frac{\partial T}{\partial t} \Big|_{\text{adiabatic}} - \frac{H \chi_{\max}}{\rho c_p} \frac{1 - Z_{\text{st}}}{Z_{\text{st}}} \frac{T - T_{\infty}}{T_{\max} - T_{\infty}} \quad (5)$$

where T_{∞} is the background temperature defined as $T_{\infty} = T_{\text{oxy}} + Z(T_{\text{fuel}} - T_{\text{oxy}})$ and H is related to the heat released for the fuel-air mixture under consideration. The magnitude of H is selected to create an overall flame heat loss capable of cooling eventually to near the background temperature, and values of $H = O(10^7 \text{ J m}^{-3})$ are typical for hydrocarbon fuels but some variability with specific fuels is required. Because heat release rates are proportional to χ , χ_{\max} also appears in the heat-loss term to provide sufficient cooling for flames with higher mixing rates. As mentioned in the introduction, choosing a linear temperature profile for the source term provides a more uniform heat loss as the transient flamelet evolves through heat-loss-driven extinction. Scaling the heat loss by $(T_{\max} - T_{\infty})^{-1}$ helps to avoid an increase in stoichiometric enthalpy from reactant mixing.

This form for the source term is recommended for fire scenarios where radiative heat losses and interactions with cold surfaces can lead to local temperatures very near ambient.

Spitfire is used to compute and save steady-state solutions to the adiabatic equations (1) and (2) for a range of stoichiometric scalar dissipation rates, $\chi_{st} = \chi(Z_{st})$. Each of these adiabatic solutions then becomes the initial condition to the solution of the non-adiabatic equations (1) and (5). The transient evolution of equations (1) and (5) creates the non-adiabatic states for the full range of flame temperatures from adiabatic to within 5% of the maximum T_∞ . Any properties needed for the flow evolution and any source terms for radiation and soot needed for the fire simulation are also computed and saved. This process results in an unstructured non-adiabatic flamelet library since each transient solution covers a slightly different range of stoichiometric enthalpy deficit, $\gamma_{st} = \gamma(Z_{st})$, values. Example enthalpy deficit, γ , profiles for χ_{st} of 10^{-3} and 10.8 s^{-1} are shown in Figure 1. We notice that the profiles are more curved for the lower χ_{st} and appear more triangular for larger χ_{st} .

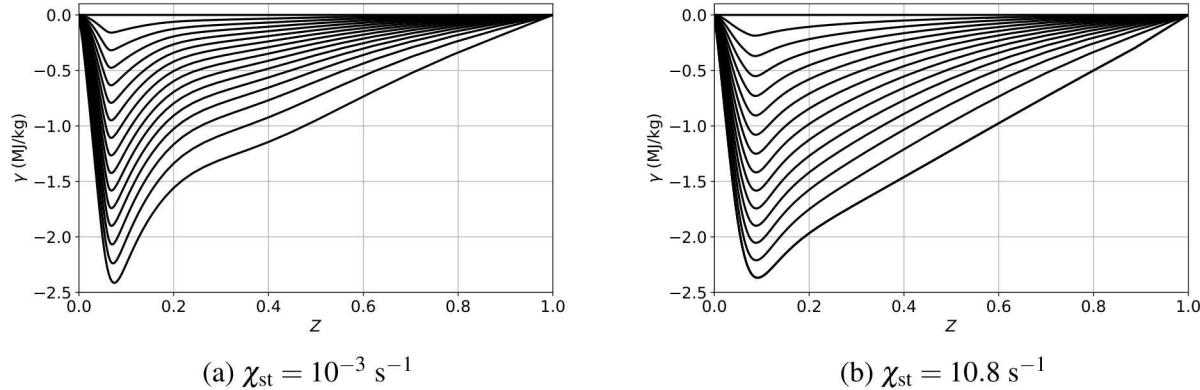


Figure 1: Example transient enthalpy deficit, γ , profiles over Z .

3. Tabulation

Once an unstructured non-adiabatic flamelet library is compiled it can be advantageous to organize the data into a structured table in order to optimize access speed. Since χ and γ are both functions of mixture fraction, the flamelet solutions must be tabulated over reference values, which we choose to be χ_{st} and γ_{st} . Because the flamelet data is unstructured in the γ_{st} direction, it is first interpolated onto a uniform grid based on the largest range in γ_{st} from the unstructured data. However, the range of γ_{st} values varies with χ_{st} ; for example, in the profiles of Figure 1 the lowest value of γ_{st} for $\chi_{st} = 10^{-3} \text{ s}^{-1}$ is -2.41 MJ/kg while for $\chi_{st} = 10.8 \text{ s}^{-1}$ it is -2.16 MJ/kg . In the events where the uniform grid extends beyond the last available γ_{st} of the unstructured data, the last available property value is repeated to fill out the structured dimension of the table. We refer to this as *internal clipping*.

For LES, the Favre-filtered properties must be tabulated. These are calculated through a convolution of the unfiltered property with the joint probability density function (PDF) of the independent coordinates. The filtered mixture fraction variance is also added as an independent coordinate to

the table for turbulent flows where the mixture fraction is not well resolved. Assuming statistical independence of the independent coordinates, their joint PDF can be approximated as the product of single-variate PDFs which are modeled using a presumed PDF approach. A clipped-Gaussian-PDF or beta-PDF is used for the mixture fraction, which is constructed using the mixture fraction mean and variance; and a Dirac delta PDF is used for the scalar dissipation rate and enthalpy deficit.

Once the library data is in a structured table format, cubic B-splines are used to interpolate the data over the table coordinates. These table interpolants can then be queried during a simulation. Using cubic splines ensures continuous first and second derivatives, which are important for robustness of projection methods, for example, used in low-Mach number combustion. If a table is queried outside of its defined independent coordinate boundaries in a certain dimension, the nearest value in that dimension is returned. This aims to keep the thermochemical state of the simulation near the well defined states in the table and we refer to it as *external* clipping. In general, both internal and external clipping force constant values outside of the unstructured flamelet data.

4. Accessing Non-Adiabatic Tables

4.1 Structured Tables

In a laminar flame simulation, transport equations for the mixture fraction and enthalpy (total and adiabatic enthalpies from which an enthalpy deficit can be calculated) are solved. The instantaneous laminar scalar dissipation rate can be calculated from the mixture fraction by

$$\chi = 2\mathcal{D} \frac{\partial Z}{\partial x_i} \frac{\partial Z}{\partial x_i} \quad (6)$$

where \mathcal{D} is a molecular mass diffusion coefficient. For turbulent simulations, the filtered variables are transported, a filtered form of (6) is computed, and an additional equation for the filtered mixture fraction variance is solved. In each of these cases, we need some way of going from the simulation variables to the independent table coordinate variables for looking up property values.

In the laminar case, going from χ to χ_{st} is relatively straightforward since there is a well defined form for $\chi(Z)$ given by (4) that was used in constructing the flamelet library. We can therefore compute χ_{st} from

$$\chi_{st} = \frac{\chi F_\chi(Z_{st})}{F_\chi(Z)} \quad (7)$$

given a χ and Z from the simulation. In the turbulent case, a filtered stoichiometric scalar dissipation rate can be calculated from the filtered quantities of the simulation similarly to (7).

However, finding the correct γ_{st} coordinate for table queries is not as simple since there is not a well defined functional form of $\gamma(Z)$. Figure 1 demonstrates how much $\gamma(Z)$ changes shape as χ_{st} and γ_{st} change. Various techniques have, therefore, been developed for finding a γ_{st} given γ in order to query the tables. We introduce two classes of these methods below.

4.1.1 Consistent Enthalpy Reconstruction

We refer to the first method for accessing structured non-adiabatic tables as a consistent enthalpy reconstruction (CER). With CER the correct location in the table is found directly using a (Newton-based) root-finding method in the γ direction, requiring γ be explicitly retained in the table. The CER method ensures that the queried enthalpy is *consistent* with the simulation enthalpy, provided the correct root is found. This is considered the most accurate way of querying a property table; however, it can have challenges associated with convergence and can be expensive depending on the number of required spline evaluations to find the desired root.

One root-finding challenge present for the Newton algorithm is the degeneracy in the pure oxidizer ($Z = 0$) and pure fuel ($Z = 1$) states of the table, which is overcome by selecting the adiabatic state if the mixture fraction is within 10^{-12} of either boundary. Coarse domain discretization can result in B-spline interpolants with artificial extrema that also challenge Newton root-finding; table construction needs to have sufficient resolution to ensure additional extrema are not created by the interpolation.

The internal and external clipping of the structured tables described in §3 requires an adjustment to the Newton algorithm. The standard for property evaluation outside of the unstructured flamelet data, as introduced above, forces constant values where input information does not exist, nominally introducing zero gradients into the Newton algorithm with associated divide-by-zero problems. Furthermore, the non-uniqueness of boundary γ values introduced by clipping prohibits a nonlinear solver from finding a unique root for large γ values. We address these issues in two ways. First, we linearly extrapolate the γ property instead of clipping to a constant value. Second, we combine Newton's method with a backtracking line-search algorithm and provide a significant penalty (i.e. a 10^{16} multiplier) for residuals outside of the structured table bounds. These approaches are a critical component of the CER implementation.

4.1.2 Approximate Enthalpy Reconstruction

We refer to the second method for accessing structured non-adiabatic tables as an approximate enthalpy reconstruction (AER). The idea behind AER is to define a functional form for the enthalpy deficit in terms of the independent structured table coordinates, which would allow for easy calculation of γ_{st} , similar to χ_{st} in (7). The reconstructed enthalpy deficit is therefore only *approximate* compared to the observed value because there is no fundamental form of $\gamma(Z)$ embedded in the flamelet solutions, and its shape depends on the interaction of chemical heat release, the dissipative mixing, and the heat loss. This chosen functional form for γ must be zero in the pure oxidizer and pure fuel streams and should maintain the general shape of the γ profiles which tend to have larger radiative losses to the environment near Z_{st} , as seen in Figure 1. This method requires a minimal number of polynomial evaluations and is computationally less expensive than the CER method described above, but it is also less accurate since the chosen functional form of γ is just an approximation. In the following description, we discuss potential accuracy issues associated with two functional forms.

One option for the functional form of γ is a piecewise linear, or triangular (T-AER), form given by

$$\gamma = \gamma_{st} F_\gamma(Z, Z_{st}) \quad (8)$$

where

$$F_\gamma(Z, Z_{st}) = \begin{cases} \frac{Z}{Z_{st}}, & Z \leq Z_{st} \\ \frac{1-Z}{1-Z_{st}}, & Z > Z_{st} \end{cases} \quad (9)$$

This comes from the generally triangular shapes of the γ profiles, especially seen for large χ_{st} , as demonstrated in Figure 1b. However, the linear profile is just an approximation for the actual heat loss, which has considerable variation over the possible range of γ and χ space due to heat release associated with the detailed chemistry, especially for rich mixtures. This variation is demonstrated in Figure 2 which displays γ scaled by γ_{st} and colored by $\log(\chi_{st})$. The data shown in Figure 2 is the unstructured flamelet library generated for ethylene [10] and air at 300 K for $10^{-3} \leq \chi_{st} \leq 10^2 \text{ s}^{-1}$. We see that even for a given χ_{st} and Z , the scaled γ/γ_{st} values can still cover a considerable range. However, (8) offers a simple and computationally cheap inversion to calculate γ_{st} ; a threshold is placed on the denominator of this inversion to avoid division by zero near the pure oxidizer and pure fuel states

$$\gamma_{st} = \frac{\gamma}{\max[F_\gamma(Z, Z_{st}), 10^{-12}]} \quad (10)$$

In the turbulent case, a filtered stoichiometric enthalpy deficit can be calculated from the filtered quantities of the simulation in a manner similar to (10).

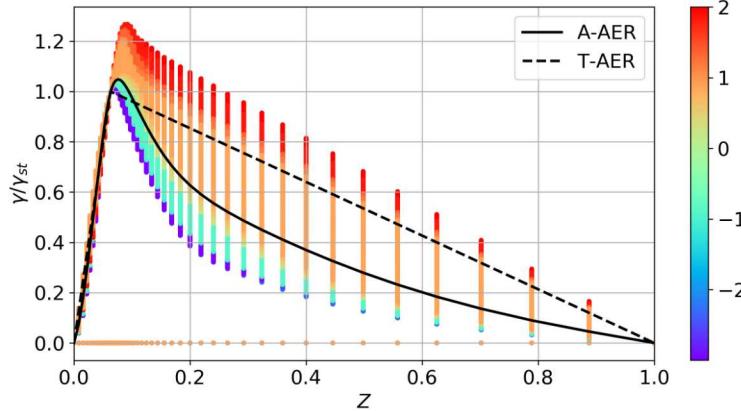


Figure 2: Unstructured scaled enthalpy deficit data for pure ethylene [10] and air at 300 K, colored by $\log(\chi_{st})$ and plotted with two AER forms of $F_\gamma(Z, Z_{st})$.

Other functional forms may be chosen for $F_\gamma(Z, Z_{st})$ and used in (10) as well. One such option is to compute an average for γ/γ_{st} and fit the result to a cubic B-spline to be evaluated as $F_\gamma(Z, Z_{st})$. This approach is referred to as the average approximate enthalpy reconstruction (A-AER); the profile for $F_\gamma(Z, Z_{st})$ is obtained by averaging γ/γ_{st} over the complete set of input flamelets (averaging over both χ_{st} and γ_{st}) with the resulting A-AER profile for ethylene also included in Figure 2. This method tends to be slightly more accurate than the T-AER method for lower χ_{st} values where the γ profiles are more curved than linear as documented in the Results section below. However, as noted earlier, both AER methods have more error than the CER approach because there is still considerable variation in the range of γ/γ_{st} profiles these methods are trying to approximate with a single line. We use the flame scenarios in section §5 to provide an example quantification of the

differences in the computational cost and accuracy of using the CER and AER methods described here.

4.2 Unstructured Tables

One option for avoiding the complexities of accessing structured tables described in section §4.1 is to build an unstructured table directly from the non-adiabatic flamelet data. This unstructured table could be constructed by regressing the flamelet data over Z , χ_{st} , and enthalpy, h , using machine learning techniques. With this format, property values could be queried directly using the simulation enthalpy instead of through an additional relationship between γ and χ_{st} .

Artificial neural networks (ANNs), for example, can provide smooth representations of the variables of interest with modest memory requirements since only the network's synaptic weights and architecture are stored. Some work has been done [11, 12] to address the non-trivial process of finding an optimal network representation in the context of combustion; but even with these optimized processes it can be hard to guarantee correct behavior. One of the reasons ANNs aren't more widely used in place of structured tabulation methods is the lack of error bounds. Another disadvantage of using ANNs is that querying the networks can be computationally more expensive than querying the interpolated tables.

However, the studies in [11, 12] found that optimal ANNs often require less memory in order to obtain equivalent accuracy compared to traditional chemistry tabulation. In addition, a major potential advantage of using ANNs is the ease of moving to higher dimensional problems. Structured tabulation is limited to around five-dimensional problems, mainly due to memory requirements. Many combustion problems, such as those involving multiple mixture fractions with heat loss, can often exceed five dimensions, especially in turbulent cases. With the smaller memory requirements of ANNs also comes the advantage of being able to run larger combustion simulations on next-generation computing architectures such as general-purpose graphics processing units (GPUs). Therefore, we explore replacing structured tabulation methods with unstructured tabulation using machine learning techniques to map dependent to independent table variables.

Our exploration of the potential for ANNs to replace structured tabulation is done through the open-source neural network library Keras [13]. The activation functions and number of layers for the ANNs are decided through trial-and-error rather than through the techniques of [11, 12], but future work should certainly consider these newer methods if our ANNs show promise. The ANNs are optimized using stochastic gradient descent with a mean squared error loss function. Further details of the ANN calibration are presented in §5.2 along with an analysis of errors from using these ANNs compared to the structured table access methods.

5. Results

5.1 Comparing Methods for Accessing Structured Tables

While the methods described here are generally used in turbulent simulations, some issues associated with structured table lookup methods are masked in turbulent flames and most easily identified in laminar flames. For this purpose we consider an ethylene laminar coflow, sooting jet with a diameter of 11 mm; the bulk velocities of ethylene and air are 4 cm/s at 485 K and 8.66 cm/s at 300

K respectively. The size of the domain is 8.6 jet diameters wide and 18 diameters high; the mesh used for the cylindrical geometry contained 48480 points. Simulations are run using Fuego [14], a part of the SIERRA thermal-fluids code suite from Sandia National Laboratories that handles unstructured low-Mach number turbulent combustion. A two-equation soot model is used for predicting soot formation. A non-adiabatic structured table for ethylene [15] and air at 485 and 300 K respectively was generated for the simulation with 98 points in Z , 48 points in χ_{st} spanning $10^{-3} \leq \chi_{st} \leq 10^2 \text{ s}^{-1}$, and 48 points in γ_{st} . Both simple radiation (SR) and participating media radiation (PMR) models are considered. SR models a uniform incident radiative heat flux as $4\sigma T_{\text{ref}}^4$ for the Stefan-Boltzmann constant σ . In other words, each element in the domain exchanges radiation with the surroundings at T_{ref} , which we specify as 300 K, and does not interact with other elements. PMR transport is predicted using the discrete ordinates method.

Figure 3 displays the axisymmetric laminar SR ethylene jet steady-state flame profiles in the two-dimensional plane along the cylinder's axis for all the structured table access methods. We see the largest amount of heat loss occurring near the center of the flame and the temperature, T , and soot mass fraction, Y_{soot} , decreasing as the flow moves downstream and exits the domain. The T-AER method appears to lead to slightly larger areas of heat loss for a given contour compared to the A-AER and CER methods. The difference in soot production for each method is more apparent, highlighted by the black contour representing 10% of the maximum soot mass fraction. We can see this contour extends out past 14 diameters up the domain for the T-AER method, while it stops near seven diameters for A-AER and around six diameters up the flame for CER. This suggests the T-AER flame is emitting smoke more than the others.

Figure 4 shows the temperature over the radius at different locations moving up the flame for the AER methods compared to CER. These show similar trends between the AER methods, but the A-AER method seems to match the CER temperature more closely for the various heights. Both AER methods are first over-predicting the temperature at the centerline until around 3 diameters up when they start under-predicting the temperature at this location. Interestingly, the AER methods then again slightly over-predict the temperature at much larger heights.

The error in the AER methods compared to the CER method is shown in Figure 5 for each of the quantities displayed in Figure 3, where red signifies an over-prediction and blue represents an under-prediction. We first notice that the A-AER method leads to overall less error than the T-AER method, seen by the softer contour colors, but both AER methods show similar trends in the locations for the under- or over-predictions. Temperature is directly extracted from the table, but the differences here are compounded by small differences in flame shape associated with each method that are primarily affected by the table extraction differences for the density. Looking at the top plot for temperature, the AER methods show a pattern of over-predicting temperature near the fuel inlet followed by under-predicting temperature around 4 diameters up the flame, and then again over-predicting temperature until the end of the vertical domain. This is the same observation made in Figure 4. As a contrast to this non-monotonic trend in the error, both AER methods give rise to more heat loss and more soot production than CER along the full domain moving up through the center of the flame. Both enthalpy (heat loss) and soot production are evolved and profile differences arise through differences in the temperature dependent source terms. The areas for the largest over-predictions in temperature seem to correspond to the areas of the largest over-predictions in the heat loss and soot mass fraction as well, which happen in the fuel-rich regions

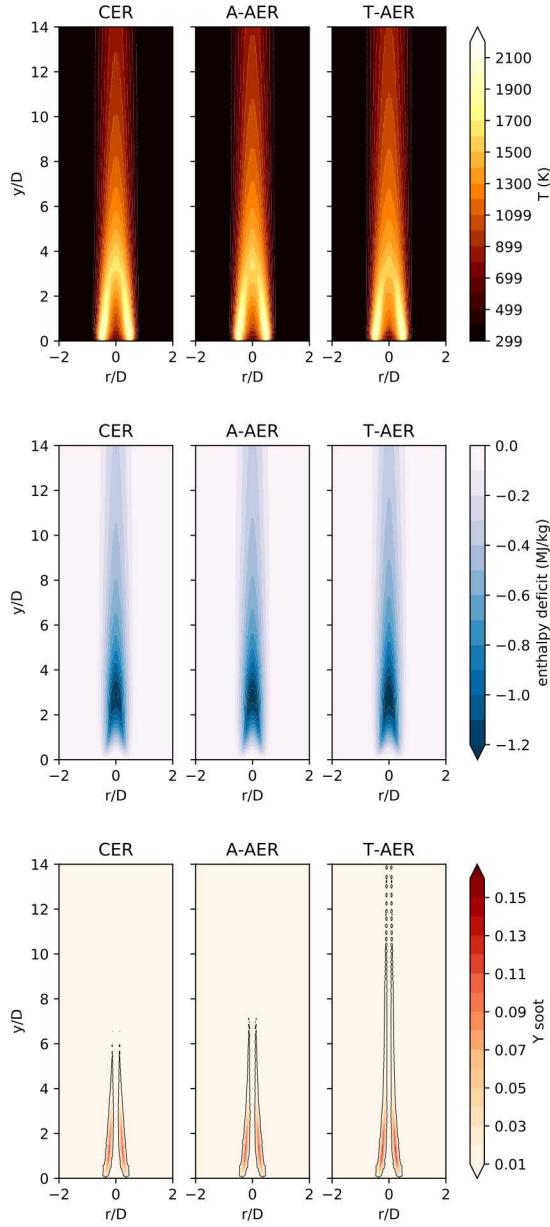


Figure 3: Steady-state profiles for the laminar SR ethylene jet. D represents the jet diameter of 11 mm, r the radius of the domain, and y the height of the domain.

of the flame near the inlet. This could be due to the increase in temperature there increasing the source terms for radiative heat losses and soot production. However, this relationship does not continue in the more lean regions of the domain because the enthalpy there is entrained higher in the flame. Differences between the methods in general become smaller near the pure streams, as expected since the property profiles become degenerate at these pure streams.

There is also less of a difference in the heat-loss predictions among the structured table access methods the closer to adiabatic the heat loss becomes. This makes sense since $\gamma = 0$ for all Z and

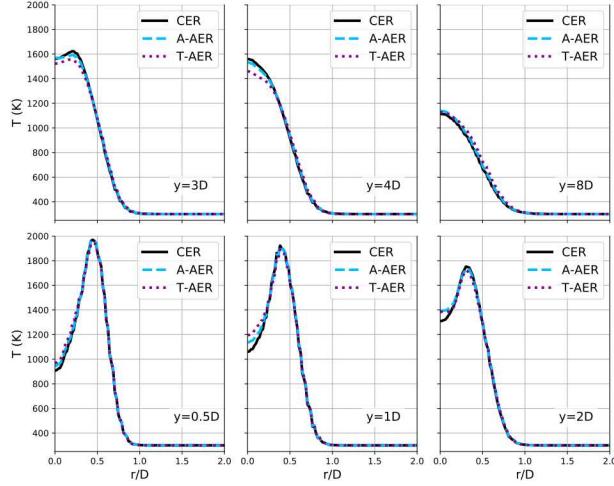


Figure 4: AER and CER steady-state temperature profiles for the laminar SR ethylene jet at different heights in the domain.

χ_{st} when the solution is adiabatic. All the structured table access methods return γ_{st} of zero for this adiabatic profile. This means the closer to zero γ becomes, the closer the calculated γ_{st} of each table access method should become. Therefore, the AER methods should have less error for scenarios with small heat loss and would be computationally less expensive than CER. Conversely, as seen in Figure 5, the more heat loss that is present in a scenario, the more differences there will be between CER and AER results.

We also note that the errors between the AER and CER methods are visible in the steady-state profiles to different extents. The difference in the maximum temperature for T-AER and CER is only around 1%, but the difference in the maximum soot mass fraction is 10%. The soot contours in Figure 3 also have more easily observable qualitative differences than the temperature contours. A-AER seems to predict closer to CER than to T-AER, only over-predicting the maximum soot mass fraction by 2% compared to CER, which is less than half of the over-prediction T-AER gives.

The differences in T-AER and CER temperature predictions can be explained with *a priori* errors shown in Figure 6. These errors were calculated using (10), given the structured table γ values, to obtain γ_{st} values that are then used to query the structured temperature table. The difference between this predicted temperature and the observed table temperature for a given γ , Z , and χ_{st} is shown in the contours of Figure 6. Figure 6a shows the *a priori* temperature errors for $\chi_{st} = 0.58 \text{ s}^{-1}$, which is near the observed steady-state χ_{st} values corresponding to the large T-AER over-predictions near the inlet. Looking at Figure 6a, we see that T-AER should over-predict temperature for all mixture fractions rich of about $Z_{st} = 0.063$ in the range of observed heat loss for this location. This is exactly the trend seen in the steady-state temperatures near the inlet.

This over-prediction in temperature for T-AER is being countered by areas of under-prediction. The areas of under-predicted temperatures in Figure 5 occur near Z_{st} where the values of observed χ_{st} decrease to around 0.05 s^{-1} . This is the χ_{st} value considered for the *a priori* temperature errors in Figure 6b. Over the range of observed γ , we can see in Figure 6b that the T-AER method leads

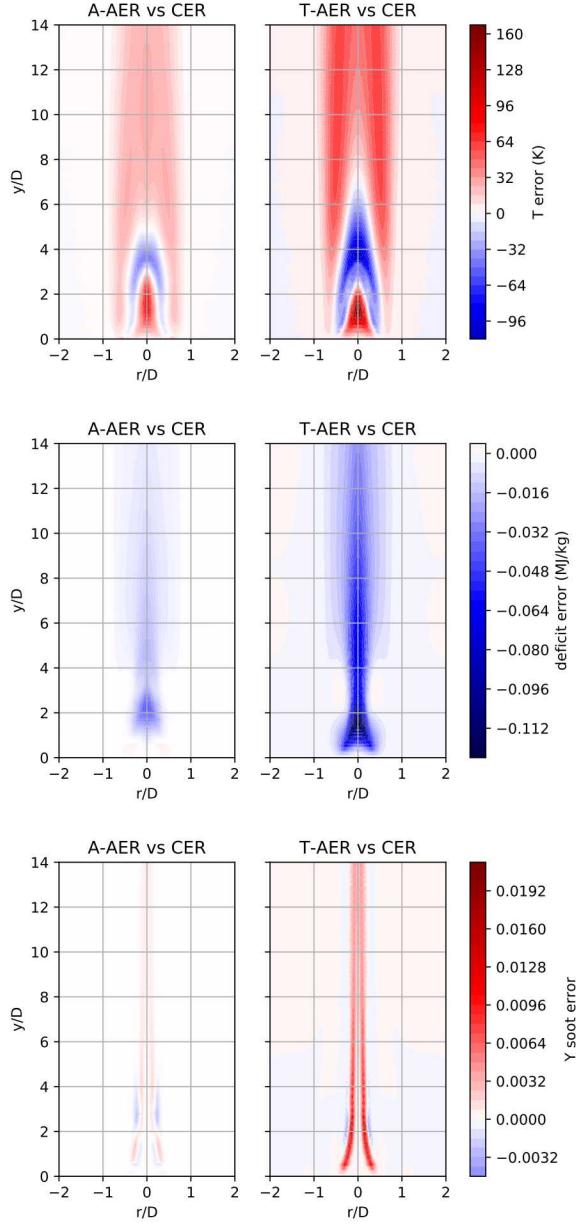


Figure 5: Difference in the AER steady-state profiles for the laminar SR ethylene jet compared to CER.

to slight under-predictions of temperature around Z_{st} . The magnitude of error in this region is also smaller than that observed at the higher χ_{st} , which matches observations in Figure 5 as well as a comparison of Figures 6a and 6b in the regions of interest.

An additional *a priori* error analysis is shown in Figure 7. This plot shows the reconstructed temperature profiles over Z for the transient flamelet solution at two values of χ_{st} . The CER method leads to near perfect reconstruction, with at most 2% error, and so is not included in these plots. These plots demonstrate the AER methods having less error in the lean regions and larger

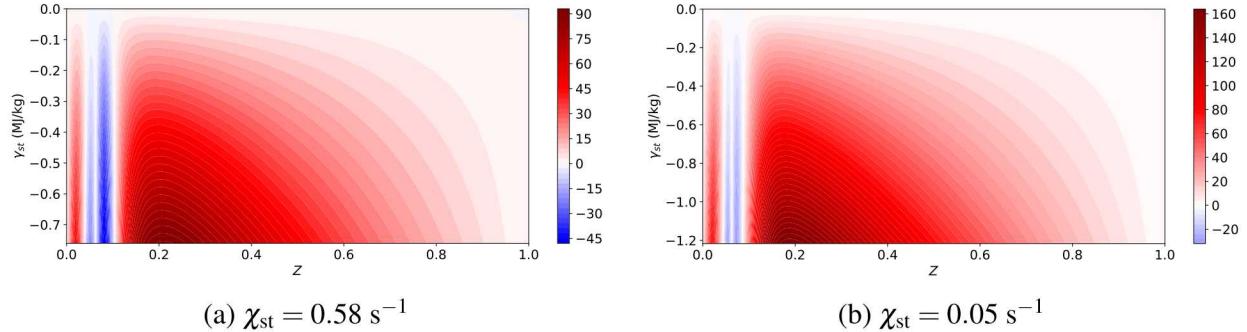


Figure 6: *A priori* error (predicted - observed) in predicting temperature for the T-AER structured table access method.

errors in rich regions, with errors disappearing at the pure streams. We also see that the A-AER method has larger errors than the T-AER method for larger χ_{st} (Figure 7a) and has less error than the T-AER method for smaller χ_{st} (Figure 7b). This can be explained by the heat-loss profiles having a more triangular shape at large χ_{st} , seen in Figure 1b, and so being better represented by T-AER. Conversely, γ profiles for smaller χ_{st} , shown in Figure 1a, are more curved, and so are better represented with the A-AER method using the average displayed in Figure 2. We note that the A-AER method could be improved for larger χ_{st} by simply calculating an average profile skewed toward those γ profiles. Since the laminar jet simulation contained more regions of lower χ_{st} , the average profile for A-AER in our case was calculated from flamelet data containing a large amount of small χ_{st} solutions, which lead to less error than the T-AER method overall in Figure 5.

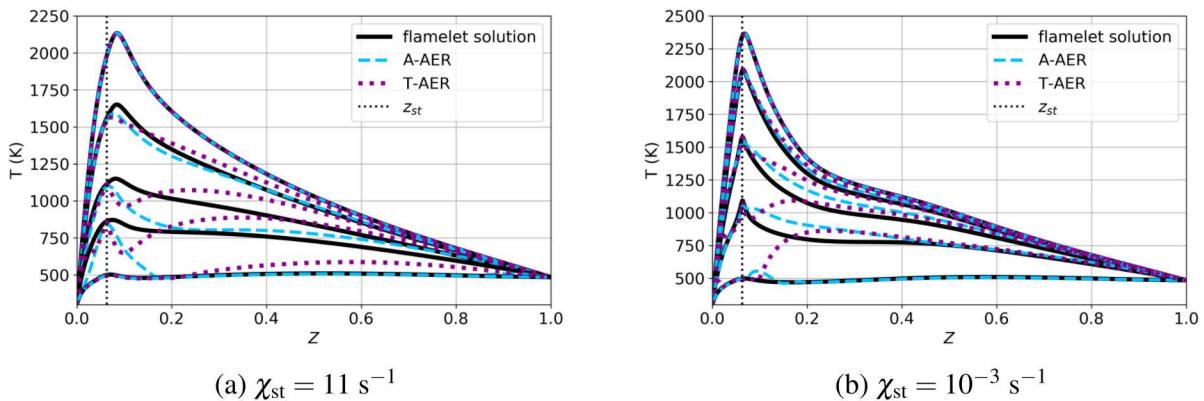


Figure 7: Reconstructed temperature profiles over Z using the AER methods compared to the original unstructured flamelet solution.

Figure 7a also shows the same non-monotonicity of error predictions moving from lean to rich regions seen in Figure 6a. This trend also exists in the smaller χ_{st} cases, but to less of an extent; instead the AER methods are mainly over-predicting temperature, which can be observed in both Figures 6b and 7b. We can also see the difference in the temperature profiles increasing moving down the temperature curves, which corresponds to increasing heat loss, as expected. Figures 6 and 7 help relate the errors we expect to see from the AER methods in various regions to the errors

observed in the simulation in Figure 5. They also show a range of possible errors reached over Z , γ_{st} , and χ_{st} which could be used to determine whether or not the error will be acceptable in the regions of interest for a simulation.

As an example of the effect the radiation model can have on the jet simulation steady state, Figure 8 shows the difference in steady-state temperature and soot mass fraction for the CER method with PMR and SR. In the regions near the inlet, the temperature is higher for PMR than SR, and this corresponds to a larger soot mass fraction. Overall, there is only at most a 5% difference in temperature, but there is an 11% higher maximum observed soot mass fraction compared to SR. Interestingly, this difference is around the same order of difference observed in comparing CER and T-AER for the SR simulation of Figure 3. Since soot production can be very sensitive to temperature, it makes sense that any small difference in temperature between SR and PMR can be amplified in the soot mass fraction data. Soot production is also more coupled to PMR than SR, so higher levels of soot would be expected when using PMR. While PMR changed the overall quantities of the steady-state properties compared to SR in Figure 3, it did not dramatically change the relative differences between the CER, T-AER, and A-AER methods. Therefore the analysis done in comparing the AER and CER methods for SR holds true for PMR as well.

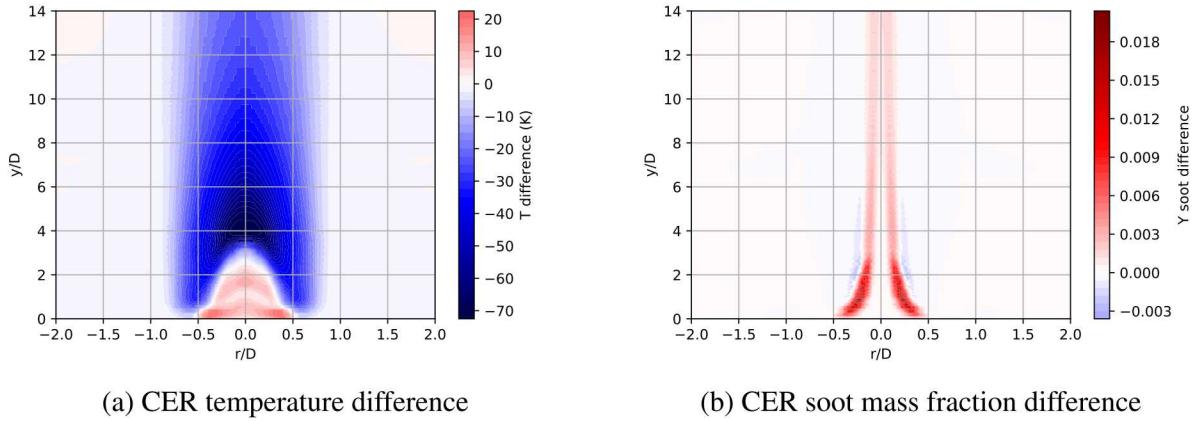


Figure 8: PMR quantities compared to SR for the CER method steady-state data.

As a final consideration for laminar fire simulations, we ran the same configuration described for the ethylene jet but replaced the fuel with heptane [16]. The resulting steady-state errors for the AER methods compared to CER are shown in Figure 9 for SR. These plots show similar trends to those explained in Figure 3 for ethylene. However, with heptane there is only about a 5% over-prediction of the maximum soot mass fraction by T-AER compared to CER. The amount of heat loss for heptane is also smaller than the amount that occurs for ethylene. This supports the earlier discussion about AER methods having less error for scenarios with small heat loss; and perhaps being worth the small errors for the less expensive computational cost compared to CER. We also continue to see the A-AER predictions sitting in between the CER and T-AER predictions, demonstrating it as a better approximation than the triangular form.

The question then becomes: how much more expensive is CER compared to the AER methods? The CER approach results in around four times as many table queries with its backtracking line

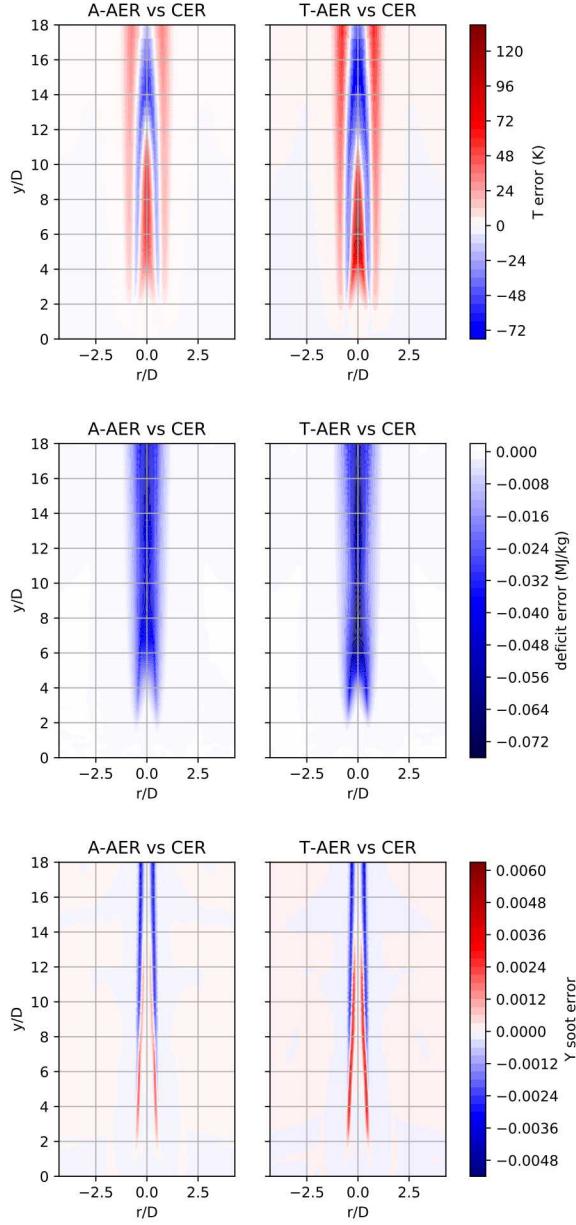


Figure 9: Difference in the AER steady-state profiles for the laminar SR heptane jet compared to CER.

search and Newton's method compared to the T-AER method in calculating the *a priori* errors of Figure 6. This factor, however, includes all the possible (Z , χ_{st} , γ_{st}) space used to build the structured table, which may not necessarily be accessed during the simulation. Therefore, this shows the potential of CER to be very expensive but does not guarantee it. The total Fuego simulation time for the PMR ethylene jet using the CER method was around twice the time for the T-AER simulation. However, this time includes performing the Newton solve for every property table query instead of optimizing to save off the query points for one property table Newton search and use those inputs to look up other property values. Therefore, an optimal redesign should result in

about a 5-10% increase in simulation time when using CER versus T-AER. This increase seems fairly reasonable, suggesting the CER method would be worth the extra cost for the added accuracy in many cases.

The results from the laminar cases seem to suggest that the difference in predictions for each of the structured table access methods are larger in scenarios with larger heat loss. Therefore, using the AER methods, which are less expensive than the CER approach, seem appropriate to use in cases with small heat loss. However, using the more accurate CER method is potentially only around a 5-10% increase in the simulation time. Both AER methods showed qualitative and quantitative differences compared to the CER approach, but the A-AER method showed less of a difference than the T-AER method.

5.2 Using Machine Learning Techniques to Map Unstructured Inputs

As mentioned in §4.2, using machine learning techniques to construct expressions for dependent variables as a function of unstructured independent variables can offer benefits including a reduction in memory requirements. The ANNs built for our analysis include seven nonlinear hidden layers, with the number of neurons in each layer decreasing from 200 down to 10, followed by a linear output layer. These are trained on an unstructured non-adiabatic flamelet temperature and density dataset for ethylene and air at the same conditions used in the laminar ethylene jet simulation. We consider two neural networks: one trained on the entire flamelet dataset (referred to as ANN1), and another trained on every other point in the dataset (referred to as ANN2). ANN2 was built in order to test the ANN's ability to perform in between the data points used to train it. The R^2 values for evaluating the ANNs on the entire unstructured flamelet dataset were above 0.999 in both cases, suggesting a reasonable network design. Figure 10 shows the parity plots for the ANN temperature predictions corresponding to these R^2 values.

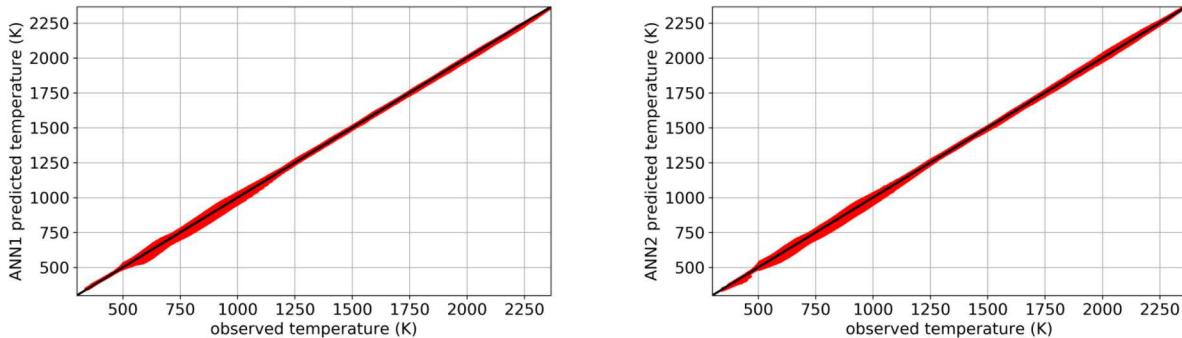


Figure 10: Temperature parity plots for artificial neural networks ANN1 and ANN2.

Figure 11 shows the reconstructed temperature profiles for both ANN1 and ANN2 over Z compared to the same flamelet solutions shown in Figure 7 for four different χ_{st} values. We can see both ANNs matching the flamelet data fairly well in all cases, especially in lean regions. There is more error in the ANN predictions for the small χ_{st} values, but not nearly as much as observed with the AER structured table access methods in Figure 7b. There is a slight difference between

ANN1 and ANN2 at $\chi_{st} = 10^{-3} \text{ s}^{-1}$, but the two are mostly indistinguishable. This means ANN2 is doing well at predicting the temperature at locations in between the data it was trained on and hasn't introduced any unphysical features. These plots show the ANNs having more error than the CER method, which reconstructs the flamelet solution data to within 2% error as mentioned earlier; but this error is still relatively small, especially compared to the AER method errors in Figure 7. Using machine learning techniques therefore has potential to be used in simulations and could be further optimized to increase its accuracy using methods such as those described in [11, 12].

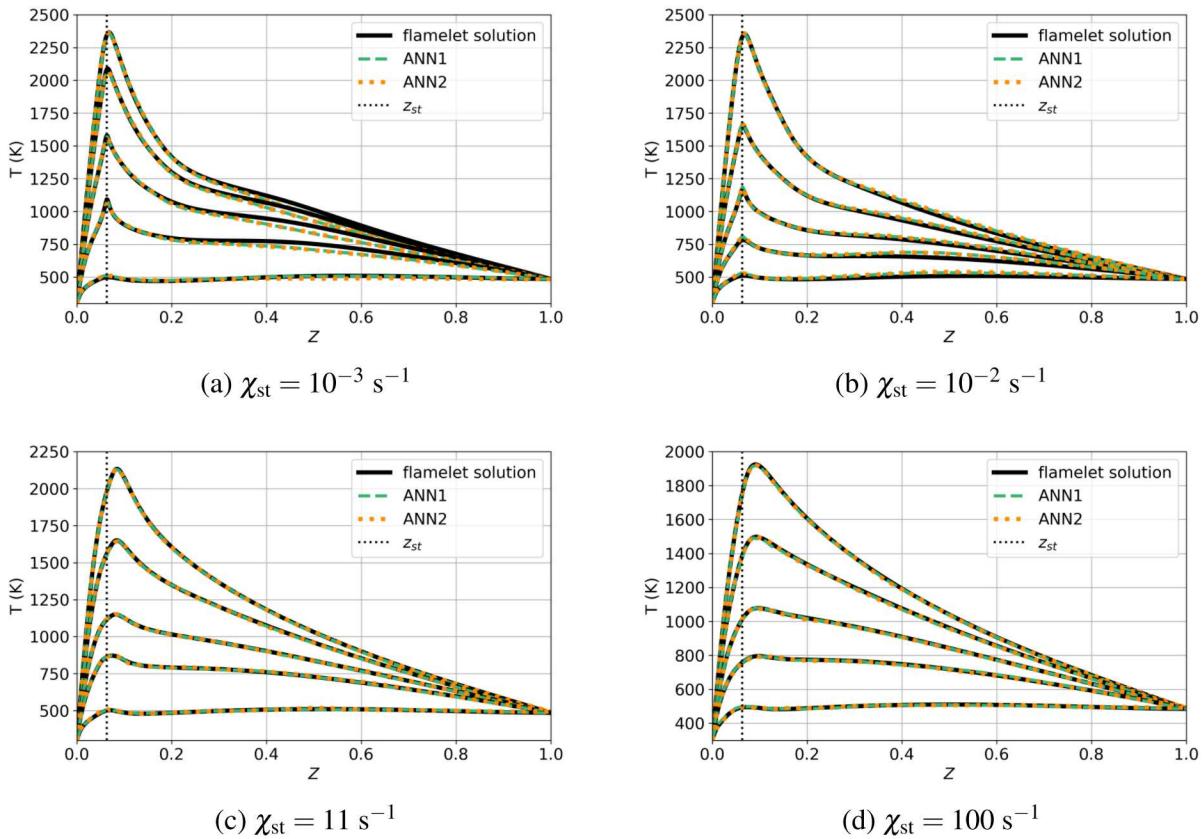


Figure 11: Reconstructed temperature profiles over Z using ANN1 and ANN2 compared to the same unstructured flamelet data as shown in Figure 7.

ANNs can also offer the advantage of less memory requirements. The ANNs trained on temperature and density data for the analysis above were only 1 MB in size. The size of a structured table built from the same unstructured flamelet data as the ANN, including the enthalpy deficit values needed for the CER approach, is about 28 MB by comparison. The size of these structured tables can also dramatically increase as they include more variables required by the flow simulation and are refined for accuracy. For fire application-scale turbulent simulations, these tables can reach a few GB in size when retaining all the required filtered quantities. The storage size for the ANNs will not increase as dramatically and often be smaller than a structured table built from the same data. This makes them more attractive than structured table methods for higher dimensional problems.

These results demonstrate the potential of machine learning techniques to be used in constructing unstructured tables for combustion simulations, especially in the context of high dimensional problems. Future work will explore the use of optimal ANNs for easier and more accurate construction as well as methods for ensuring bounded errors.

6. Conclusions

Chemistry tabulation methods make higher fidelity combustion simulations more feasible for challenging applications like fires with significant heat losses. One approach to generating non-adiabatic flamelet information involves solving unsteady flamelet equations with an appropriate heat-loss source term as used here. There are various ways to relate the variables evolved in the simulation to tabulated quantities including structured tables, involving the CER and AER methods introduced in this paper, and unstructured approaches. These structured table approaches are shown to have trade-offs in terms of accuracy and computational cost. Results show qualitative and quantitative differences using the CER, T-AER, and A-AER methods, including around 10% differences in predicted steady-state temperature, soot mass fraction, and heat loss. These differences become smaller as the heat loss nears adiabatic, as expected; therefore, the computationally less expensive AER methods can be used with negligible error in scenarios with low heat loss. In high heat-loss scenarios, CER may be needed for accurate predictions and, if optimized, is estimated to only take around 5-10% more time to run than simulations with the T-AER method. Using ANNs to build unstructured tables also shows promise as an alternative to structured tables, especially with its lower memory requirements and ability to handle higher dimensional problems. Future work should include exploring the use of optimal ANNs as well as performing a more detailed cost comparison for the AER, CER, and ANN approaches.

7. Acknowledgements

The authors gratefully acknowledge funding from the Department of Energy National Nuclear Security Administration Office of Advanced Simulation and Computing. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

- [1] N. Peters, *Turbulent Combustion*, Cambridge Monographs on Mechanics, Cambridge University Press, 2000.
- [2] B. Marracino and D. Lentini, *Radiation Modelling in Non-Luminous Nonpremixed Turbulent Flames*, *Combustion Science and Technology* 128 (1997) 23–48.
- [3] J. van Oijen and L. de Goey, *Modelling of Premixed Laminar Flames using Flamelet-Generated Manifolds*, *Combustion Science and Technology* 161 (2000) 113–137.

- [4] J. C. J. M. Hossain and W. Malalasekera, Modelling of a Bluff-Body Nonpremixed Flame using a Coupled Radiation/Flamelet Combustion Model, *Flow, Turbulence and Combustion* 67 (2001) 217–234.
- [5] M. Ihme and H. Pitsch, Modeling of radiation and nitric oxide formation in turbulent non-premixed flames using a flamelet/progress variable formulation, *Physics of Fluids* 20 (2008) 055110.
- [6] D. Carbonell, C. Perez-Segarra, P. Coelho, and A. Oliva, Flamelet mathematical models for non-premixed laminar combustion, *Combustion and Flame* 156 (2009) 334–347.
- [7] M. E. Mueller and H. Pitsch, LES model for sooting turbulent nonpremixed flames, *Combustion and Flame* 159 (2012) 2166–2180.
- [8] H. Koo, J. C. Hewson, and R. C. Knaus, LES Soot-Radiation Predictions of Buoyant Fire Plumes, Western States Section of the Combustion Institute Spring Technical Meeting, Bend, OR (2018).
- [9] M. A. Hansen, Spitfire, 2019, URL: <https://gitlab.multiscale.utah.edu/common/spitfire>.
- [10] Z. Luo, C. Yoo, E. Richardson, J. Chen, C. Law, and T. Lu, Chemical explosive mode analysis for a turbulent lifted ethylene jet flame in highly-heated coflow, *Combustion and Flame* 159 (2012) 265–274.
- [11] M. Ihme, C. Schmitt, and H. Pitsch, Optimal artificial neural networks and tabulation methods for chemistry representation in LES of a bluff-body swirl-stabilized flame, *Proceedings of the Combustion Institute* 32 (2009) 1527–1535.
- [12] M. Ihme, A. Marsden, and H. Pitsch, Generation of optimal artificial neural networks using a pattern search algorithm: Application to approximation of chemical systems, *Neural computation* 20 (2008) 573–601.
- [13] F. Chollet et al., *Keras*, <https://keras.io>, 2015.
- [14] Sandia National Laboratories SIERRA Thermal/Fluid Development Team, SIERRA Low mach Module: Fuego User Manual - Version 4.52, SAND 2019-3787, (2019).
- [15] H. Wang and A. Laskin, A Comprehensive Kinetic Model of Ethylene and Acetylene Oxidation at High Temperatures, *tech. rep. Report No.*, 1999.
- [16] S. Liu, J. C. Hewson, J. H. Chen, and H. Pitsch, Effects of strain rate on high-pressure non-premixed n-heptane autoignition in counterflow, *Combustion and Flame* 137 (2004) 320–339.