

Development of Artificial Neural Network Chemistry Framework for Turbulent Combustion of Non-premixed Flames

by

Jacob Chateauneuf

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Mechanical and Mechatronics Engineering

Waterloo, Ontario, Canada, 2025

© Jacob Chateauneuf 2025

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Detailed chemical kinetics are required to accurately model pollutant formation in turbulent combustion simulations. However, existing methods for representing detailed turbulent combustion chemistry can have prohibitive computational cost. Machine learning (ML) methods are gaining attention for representing detailed chemistry in a computationally efficient manner, however the difficulties in preparing these models limit their application. In this work, artificial neural networks (ANNs) are used to represent non-premixed turbulent combustion chemistry in the Conditional Source-term Estimation (CSE) model, which accounts for detailed kinetics and turbulence-chemistry interactions (TCI). This work represents the first application of ML to CSE.

The objectives of the present study are to (i) develop a methodology for representing any chemistry dataset with ANNs in the context of the CSE combustion model, and (ii) apply ANN-CSE to the simulation of two turbulent non-premixed methane jet flames in a Reynolds-Averaged Navier-Stokes (RANS) framework.

The ANNs are developed using the MATLAB Deep Learning Toolbox. Two tabulated chemistry datasets are considered for ANN development: Trajectory-Generated Low-Dimensional Manifolds (TGLDM) and samples from direct integration (DI). In addition, pure and diluted methane fuels are considered. Detailed chemistry via GRI-Mech 3.0 and reduced chemistry via Smooke’s mechanism is used. A data preparation procedure for ANN development is outlined. The sensitivity of various ANN parameters are also investigated to optimize the ANNs to each dataset. It is found that the ANNs can predict species mass fractions, reaction rates, temperature and heat release rate with good accuracy for each case. The storage requirement is also reduced by over 50% for each case.

ANN-CSE is applied to turbulent non-premixed jet flames with both sets of data and either pure or diluted methane. The predictions of conditional and Favre averages of temperature, species mass fractions and source terms from ANN-CSE are compared to those from CSE with tabulated chemistry (TGLDM-CSE) or direct chemistry (DI-CSE) to verify ANN-CSE. Identical computational settings for each flame are used to verify the influence of the chemistry implementation on the flame structure. For each case, ANN-CSE is generally able to capture the trends of each conditional and Favre-averaged quantity at various locations. In addition, ANN-CSE requires less memory and is faster than regular CSE for each case.

This study shows that CSE can be effectively coupled with ANNs for chemistry representation of different mechanisms and sources of data. This will enable CSE for simulating more detailed cases. Future work may involve more complex fuels, more sophisticated ANNs and data preparation, or the Large Eddy Simulation (LES) framework.

Acknowledgements

There have been a number of people who have assisted me throughout my graduate career and made this work possible. I would like to take this opportunity to acknowledge them and thank them for their support.

First, I would like to thank my supervisor, Professor Cécile Devaud, for her support throughout my graduate studies. She introduced me to the conditional source-term estimation model and machine learning techniques, and her support has enhanced my knowledge of turbulent combustion modeling and computational fluid dynamics. Her guidance and encouragement was unwavering during my degree. My accomplishments would not have been possible without her support.

I would like to thank Professors Vinny Gupta and Robert Nishida for consenting to be members of the thesis committee and for reviewing my thesis. I would also like to acknowledge the financial support from the National Sciences and Engineering Research Council (NSERC).

I would also like to thank my friends and colleagues at the turbulent combustion modeling lab, especially Ahmed Abdalhamid, Amir Mahdipour and Akira Tiesma, for their support and guidance. Finally, I would like to thank my family for their support during my graduate studies and throughout my life.

Table of Contents

Author’s Declaration	ii
Abstract	iii
Acknowledgements	iv
List of Figures	ix
List of Tables	xiv
1 Introduction	1
1.1 Overview	1
1.2 Objectives	3
1.3 Thesis Outline	4
1.4 Author’s Current Contributions	5
2 Background	6
2.1 Governing Equations	6
2.1.1 Conservation of Mass	6
2.1.2 Conservation of Linear Momentum	7
2.1.3 Conservation of Energy	7
2.1.4 Species Transport	9

2.2	Turbulent Flows	9
2.2.1	Description of Turbulence	9
2.2.2	Reynolds-Averaged Navier-Stokes	10
2.2.3	Averaged Transport Equations	11
2.3	Eddy Viscosity Turbulence Models	13
2.3.1	Mixing Length Model	13
2.3.2	k - ϵ Model	14
2.3.3	k - ω Model	15
2.3.4	Shear Stress Transport Model	16
2.4	Turbulent Non-premixed Combustion Modelling	16
2.4.1	Turbulent Combustion Models	16
2.4.2	Mixture Fraction	18
2.4.3	Beta Probability Density Function (PDF)	19
2.5	Conditional Source-term Estimation for Turbulent Non-premixed Combustion	20
2.5.1	CSE Principles	20
2.5.2	Integral Inversion and Regularization	21
2.5.3	Chemistry Representation	23
2.5.4	CSE Implementation	28
2.6	Summary	29
3	Machine Learning Methods for Turbulent Combustion	31
3.1	Machine Learning Models	31
3.1.1	Decision Trees	31
3.1.2	Artificial Neural Networks	33
3.2	Literature Review	34
3.2.1	Early ML Studies	35
3.2.2	ML Chemistry with Flamelet Training Data	37
3.2.3	ML Chemistry with Sampled Training Data	41

3.2.4	ML Chemistry with Alternative Source of Tabulated Training Data	44
3.2.5	Studies Involving Comparisons of ML Methods	47
3.3	Summary	50
4	Machine Learning Methods for Chemistry Representation	52
4.1	ANN Architecture	52
4.2	ANN Data Pre-processing	59
4.2.1	Tabulated Chemistry Data Preparation	60
4.2.2	Direct Chemistry Data Preparation	62
4.3	ANN Training and Testing Process	62
4.4	ANN Training Results	71
4.4.1	TGLDM with GRI-Mech 3.0 for Pure and Diluted Methane	72
4.4.2	TGLDM and DI Samples with Smooke’s Mechanism for Diluted Methane	74
4.5	CSE Implementation with ANN Chemistry	76
4.6	Summary	77
5	ANN-CSE for Turbulent Non-premixed Jet Flames with Pure and Diluted Methane	78
5.1	Case 1: Brookes-Moss Flame	78
5.1.1	Experimental Setup	79
5.1.2	ANN Preparation	79
5.1.3	Computational Details	79
5.2	Case 2: Sandia Flame D	80
5.2.1	Experimental Setup	81
5.2.2	ANN Preparation	82
5.2.3	Computational Details	83
5.3	Results	85
5.3.1	Mixing Field	85

5.3.2	ANN Input Variables: Conditional CO ₂ and H ₂ O Mass Fractions . . .	87
5.3.3	Other Conditional Averages	88
5.3.4	Favre Averages	94
5.3.5	Computational Efficiency	101
5.4	Summary	102
6	ANN-CSE for Sandia Flame D Using Multiple Sources of Data	103
6.1	ANN Preparation	103
6.2	Computational Details	104
6.3	Results	105
6.3.1	Mixing Field	107
6.3.2	ANN Input Variables: Conditional CO ₂ and H ₂ O Mass Fractions . .	107
6.3.3	Other Conditional Averages	108
6.3.4	Favre Averages	112
6.3.5	Computational Efficiency	117
6.4	Summary	118
7	Conclusions	120
7.1	Summary of Main Findings	120
7.2	Summary of Accomplishments	123
7.3	Recommendations and Future Work	123
	References	125

List of Figures

2.1	Illustration of a TGLDM manifold for methane-air combustion at the stoichiometric mixture fraction ($\zeta = 0.055$)	25
2.2	Illustration of a TGLDM manifold with Delaunay triangulation for pure methane-air combustion at the stoichiometric mixture fraction ($\zeta = 0.055$)	26
2.3	Illustration of the carbon dioxide reaction rate mapped onto the TGLDM manifold	27
2.4	Flowchart of CSE implementation with reduced chemistry in the CFD Module	29
4.1	a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for single-layer ANNs with increasing number of neurons. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates 50 neurons	55
4.2	a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for 50-neuron ANNs with increasing number of hidden layers. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates 5 hidden layers	56
4.3	a) MSE, b) R^2 , and c) training time for ANNs with multiple outputs (50x9, blue) or single outputs (450x1, red). The bottom row is a reproduction of the top row in a reduced ζ range. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 across all ζ . The vertical dashed line indicates the stoichiometric ζ of 0.353	58
4.4	Illustration of TGLDM data augmentation for ANN training, considering pure methane-air combustion at $\zeta = 0.2$: a) base TGLDM data, b) overlapping trajectory is removed, c) additional points are added near the boundaries as constraints, d) Delaunay triangulation is used to generate additional points	61

4.5	Influence of velocity amplitude (columns) and frequency (rows) on $\widetilde{Y_{CO}} \zeta$ at $\zeta = 0.353$ for Sandia flame D with DI samples. Data is saved every 250 time steps (top two rows) or 25 time steps (bottom two rows)	63
4.6	a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing batch sizes. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates a mini-batch size of 250	67
4.7	a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing number of epochs. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates 100 epochs	68
4.8	a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing learning rate. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates a learning rate of 0.01	69
4.9	a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing momentum. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates a momentum factor of 0.9	70
4.10	a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing L2 regularization. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates an L2 regularization factor of 0.0001	70
4.11	a) MSE and b) R^2 for ANNs with various learning rate drop factors and periods. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. Horizontal dashed lines indicate results with a constant learning rate	71
4.12	MSE and R^2 for all 50 ANNs using GRI-Mech 3.0 with a) pure methane and b) diluted methane chemistry. The vertical line indicates the stoichiometric mixture fraction	73
4.13	MSE and R^2 for all 50 ANNs using Smooke's mechanism with a) TGLDM solutions and b) DI samples. The vertical line indicates the stoichiometric mixture fraction	75

5.1	Illustration of computational domain for Case 1 (Brookes-Moss flame): a) temperature contour with CSE ensemble boundaries shown in black and axial locations of results shown in red, b) mesh details near the main fuel inlet	81
5.2	Illustration of computational domain for Case 2 (Sandia flame D): a) temperature contour with CSE ensemble boundaries shown in black and axial locations of results shown in red, b) mesh details near the main fuel inlet	85
5.3	Radial profiles of the mean (top) and RMS (bottom) of the mixture fraction at various locations in the Brookes-Moss flame from ANN-CSE (solid line), TGLDM-CSE (dashed line), and experimental data (symbols)	86
5.4	Radial profiles of the mean (top) and RMS (bottom) of the mixture fraction at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line), and experimental data (symbols)	87
5.5	Conditional CO ₂ (top) and H ₂ O (bottom) mass fraction profiles at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line). The stoichiometric mixture fraction is indicated by a vertical dashed line	88
5.6	Conditional CO ₂ (top) and H ₂ O (bottom) mass fraction profiles at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols). The stoichiometric mixture fraction is indicated by a vertical dashed line	89
5.7	Conditional CO (top), H ₂ (middle) and OH (bottom) mass fraction profiles at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line). The stoichiometric mixture fraction is indicated by a vertical dashed line	90
5.8	Conditional CO (top), H ₂ (middle) and OH (bottom) mass fraction profiles at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols). The stoichiometric mixture fraction is indicated by a vertical dashed line	91
5.9	Conditional CO ₂ , H ₂ O, CH ₄ , O ₂ reaction rate and heat release rate profiles at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line). The stoichiometric mixture fraction is indicated by a vertical dashed line	92

5.10	Conditional CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate profiles at various locations in Sandia flame D from ANN-CSE (solid line) and TGLDM-CSE (dashed line). The stoichiometric mixture fraction is indicated by a vertical dashed line	93
5.11	Radial Favre averages of the CO_2 (top) and H_2O (middle) mass fractions and temperature (bottom) at various locations in the Brookes-Moss flame from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols)	95
5.12	Radial Favre averages of the CO_2 (top) and H_2O (middle) mass fractions and temperature (bottom) at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols)	96
5.13	Radial Favre averages of the CO (top), H_2 (middle) and OH (bottom) mass fractions at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line)	97
5.14	Radial Favre averages of the CO (top), H_2 (middle) and OH (bottom) mass fractions at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols)	98
5.15	Radial Favre averages of the CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line)	99
5.16	Radial Favre averages of the CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate at various locations in Sandia flame D from ANN-CSE (solid line) and TGLDM-CSE (dashed line)	100
6.1	Illustration of computational domain for Sandia flame D in Chapter 6: a) temperature contour with CSE ensemble boundaries shown in black, axial locations of results shown in red, and parallel processing domain boundaries in grey, b) mesh details near the main fuel inlet	105
6.2	Radial profiles of the mean (top) and RMS (bottom) of the mixture fraction at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue), with experimental data (symbols)	108

6.3	Conditional CO_2 (top) and H_2O (bottom) mass fraction profiles at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue) with experimental data (symbols). The stoichiometric mixture fraction is indicated by a vertical dashed line	109
6.4	Conditional CO (top), H_2 (middle) and OH (bottom) mass fraction profiles at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue), with experimental data (symbols). The stoichiometric mixture fraction is indicated by a vertical dashed line	110
6.5	Conditional CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate profiles at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue). The stoichiometric mixture fraction is indicated by a vertical dashed line	111
6.6	Radial Favre averages of the CO_2 (top) and H_2O (middle) mass fractions and temperature (bottom) at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue), with experimental data (symbols)	113
6.7	Radial Favre averages of the CO (top), H_2 (middle) and OH (bottom) mass fractions at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue), with experimental data (symbols)	114
6.8	Radial Favre averages of the CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue) . . .	115

List of Tables

3.1	Summary of early ML studies for turbulent combustion chemistry	37
3.2	Summary of ML turbulent combustion chemistry studies involving flamelet training data	42
3.3	Summary of previous ANN thermochemistry studies involving samples for training data	44
3.4	Summary of previous ANN thermochemistry studies involving alternative sources of training data	48
3.5	Summary of previous ML method comparison studies for thermochemistry	49
4.1	MSE, R^2 , training time and storage size for ANNs with various even neuron distributions. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The 5x50 distribution is selected	57
4.2	MSE, R^2 , training time and storage size for ANNs with even and decreasing neuron distributions. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The even distribution is selected	57
4.3	MSE, R^2 , and training time for ANNs with various activation functions. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The ‘ReLU’ activation function is selected	59
4.4	MSE, R^2 , training time and storage size for ANNs with min-max and Z-score normalization (selected). Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$	65
4.5	MSE (top) and R^2 (bottom) of rescaled variables for ANNs with min-max and Z-score normalization (selected). Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$	65

4.6	MSE, R^2 , training time and storage size for ANNs with the SGDM and Adam solvers. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The SGDM solver is selected	66
4.7	Amount of data and storage size for the original TGLDM dataset, augmented TGLDM dataset, and the weights and biases of the TGLDM-based ANNs for both fuels	74
4.8	Amount of data and storage size for the original and augmented TGLDM and DI datasets for diluted methane-air combustion with Smooke's mechanism, and the weights and biases of the ANNs for each case	76
5.1	Boundary conditions for Case 1 (Brookes-Moss flame), split by variable and boundary name	82
5.2	Boundary conditions for Case 2 (Sandia flame D), split by variable and boundary name	84
5.3	Memory usage, CSE routine execution time, and simulation wall-clock time for ANN-CSE and TGLDM-CSE for both flames	101
6.1	Boundary conditions for Sandia flame D (Chapter 6), split by variable and boundary name	106
6.2	Memory usage, CSE routine execution time, and simulation wall-clock time for ANN-CSE and DI-CSE for both flames	117

Chapter 1

Introduction

1.1 Overview

Combustion is a phenomenon involving the reaction of a fuel and an oxidizer with sufficient energy. Combustion of fossil fuels remains relevant for energy usage, accounting for about 80% of the global energy supply [1]. In addition, the global use of energy has increased from about 250 EJ ($=10^{18}$ J) to about 600 EJ in the past five decades [1]. Turbulent combustion in particular is used in various engineering applications, such as engines, gas turbines, furnaces, boilers, and stoves.

Devices that use turbulent combustion must be designed to provide the intended function while being safe and reducing emissions as much as possible. Prototypes can be used to aid in the design process, however prototyping can be expensive and is not always feasible. A general approach is to use turbulent combustion modelling, often with computational fluid dynamics (CFD) software, in tandem with experimental designs. CFD has the ability to iterate through many designs while being much cheaper than experimentation, and it serves to verify results from experimental findings and provide insights on underlying physics or new designs. The accuracy of CFD is therefore important, as it will impact the design and performance of engineered devices.

Turbulent combustion modelling is important for the design and understanding of combustion systems. However, turbulent combustion modelling remains challenging due to the presence of turbulence-chemistry interactions (TCI), where the turbulent fluctuations impact the mixing of reactants while the chemical reaction produces a change in fluid properties. With the presence of TCI, the mean species reaction rates cannot be evaluated

directly from the mean quantities, and must be approximated with a turbulent combustion model. Different approaches exist [2–8], characterized by the assumptions made about TCI. Among these methods is the Conditional Source-term Estimation (CSE) model [9, 10], which is used in this work.

Turbulent combustion models must contain detailed chemistry in order to predict emissions and model the combustion properly. This is difficult however, considering that detailed mechanisms of hydrocarbon fuels may contain hundreds of species and thousands of reactions. It is extremely challenging to solve transport equations for each species due to the numerical stiffness of the reaction rates and the need to model unclosed terms arising from turbulent fluctuations. In practice, the number of species is reduced while retaining information about the entire mechanism. The most common approach is to pre-tabulate the detailed chemical mechanism based on the solution of a canonical combustion problem, and access the chemical quantities during the simulation. Various tabulated chemistry methods exist [11–16] for this purpose. While these methods reduce the computational cost of turbulent combustion CFD, the resulting tables require large amounts of computer (CPU) storage, often on the order of gigabytes (GB). The storage requirement becomes prohibitive with a large number of species or refinement in the data. Turbulent combustion CFD with tabulated chemistry is often modelled with reduced species.

It is desired to have a methodology that can represent the detailed chemistry of any fuel or mechanism while having a minimal impact on the computational performance and storage. Machine learning (ML) has been gaining recent attention [17, 18] in the field of turbulent combustion for this purpose. ML is a subset of artificial intelligence (AI) where models are trained to learn underlying patterns in a given dataset and be applied to independent problems. ML models can be used to perform classification or regression, among other tasks. What makes ML attractive is that the models are computationally simple and require minimal CPU storage. This allows for the opportunity to represent fuels that were previously not possible with either direct or tabulated chemistry. However, ML models are not routinely implemented in practical turbulent combustion CFD applications for various reasons. First, sufficient amounts and quality of training data are needed to properly tune the ML model for a given turbulent combustion problem. ML models are often specific to a given configuration for this reason. Second, the design and optimization of the ML model structure is non-trivial and requires thorough testing. ML chemistry representation is thus a time-consuming process. Despite these setbacks, various studies of ML chemistry representation in different combustion models have been completed [17, 18]. The benefits of pairing ML with CSE, however, are unclear and require further investigation.

CSE is a turbulent combustion model that is computationally efficient and accounts for both TCI and detailed chemistry. CSE has previously been applied successfully to

various flames, including premixed [19–24] and non-premixed combustion [25–33], lifted flames [34, 35], spray flames [36–40], MILD (Moderate or Intense Low-Oxygen Dilution) combustion [41–43], and pool fires [44–46]. These studies also involve various treatments for flame extinction, heat losses, radiation, and soot formation. However, the majority of these studies rely on tabulated chemistry. CSE with direct chemistry was recently applied successfully to Sandia flame D [33], however this required a reduced mechanism to be computationally feasible. Although the flame structure is captured in each case, the use of either tabulated or direct chemistry has restricted previous CSE studies to simple hydrocarbon fuels. The current work expands on the capabilities of the CSE combustion model by incorporating chemistry representation with ML methods for the first time. ML is promising for chemistry representation in CSE, as it is computationally efficient and can incorporate information from detailed chemical kinetics. This method has the potential to represent fuels with complex chemistry, such as alternative fuels, that were previously not possible with CSE, in addition to reducing the storage requirement and improving the speed of CSE. The use of ML in CSE will enable the further application and adoption of the CSE combustion model.

1.2 Objectives

The motivation of the current work is to expand the applicability of CSE to different fuels by addressing the chemistry limitations. This is achieved by using ML methods, specifically the artificial neural network (ANN), to represent combustion chemistry in the CSE model. This work focuses on non-premixed combustion, which is an industrially-relevant combustion condition that has been previously simulated successfully with CSE. Methane is selected as the fuel of study to validate the current approach with previous CSE approaches using either direct or tabulated chemistry. ANNs are selected in this context to represent conditional species mass fractions, source terms, and thermal quantities at a given mixture fraction. The ANN-CSE approach should be adaptable not only to different fuels, but to different sources of data, to enable development and application to different cases. In this work, two methane mechanisms and two sources of data are considered to meet this requirement, with the ANNs being developed specifically for each case. It is important to stress that ANN design is a non-trivial process, in which many parameters must be optimized to fit the problem that the ANNs are applied to. The ANN-CSE approach is validated by simulating various turbulent non-premixed methane jet flames [47–49] that have been previously simulated using CSE with either tabulated chemistry or direct integration (DI). These simulations are carried out using a Reynolds-Averaged Navier-Stokes (RANS)

framework. Results for the predicted quantities are analyzed with both CSE formulations at various locations to assess the accuracy and influence of chemistry representation in CSE. Experimental data is included when available. The computational speed and memory are also compared between both CSE formulations. The successful implementation of ANN-CSE is indicated by the relative accuracy of the predicted quantities from the ANNs, as well as the computational data from each flame simulation, to those from either TGLDM-CSE or DI-CSE. The resulting ANN framework is independent of CSE itself and could be easily extended to premixed flames or Large Eddy Simulation (LES).

The objectives of this study are summarized as follows:

- Develop a methodology for representing any chemistry dataset with ANNs in the context of the CSE combustion model.
- Apply ANN-CSE to the simulation of two turbulent non-premixed methane jet flames in a Reynolds-Averaged Navier-Stokes (RANS) framework.

1.3 Thesis Outline

Chapter 2 presents the foundations of turbulent non-premixed combustion modelling with CSE. This includes the governing equations, turbulence and averaging methods, turbulence models and a detailed description of CSE. A review of previous ML applications for turbulent combustion chemistry is presented in Chapter 3. The role of the current work in the literature is highlighted. Chapter 4 then examines the ANN development and optimization. This includes the ANN architecture, data preparation, and the training and testing process. Measures of accuracy for each case under study are also presented. In Chapter 5, ANN-CSE is evaluated against TGLDM-CSE in two separate non-premixed jet flames: one with pure methane, and one with diluted methane (Sandia flame D). This chapter includes results for the mixing field, conditional and Favre averaged profiles of reactive scalars (composition, temperature, source terms), and computational statistics of both flames. In Chapter 6, ANN-CSE is developed using training data from either TGLDM or DI, and is applied to Sandia flame D. Note that a separate mechanism is used for this chapter. The same set of results in Chapter 5 are used to compare the various approaches. Finally, conclusions and recommendations for future work are made in Chapter 7.

1.4 Author's Current Contributions

It is acknowledged that parts of the content presented in Chapters 2-5 are published in the following article under review:

- Jacob D. Chateaufneuf and Cecile Devaud. Development of Artificial Neural Networks (ANNs) for Chemistry Representation in the Conditional Source-term Estimation (CSE) Combustion Model. Combustion Theory and Modelling, under review, 2025.

Chapter 2

Background

This chapter presents the background related to turbulent combustion modelling with Conditional Source-term Estimation (CSE). The chapter starts with governing equations for turbulent combustion. Turbulent flow and turbulence models are then discussed, with a focus on Reynolds-Averaged Navier-Stokes equations. An overview of turbulent combustion parameters and models is presented, followed by a detailed presentation of CSE for non-premixed combustion.

2.1 Governing Equations

All fluid flows are governed by the conservation of mass, momentum, and energy. The transport of species is also crucial in the case of reacting flows. The fluid flow properties can be obtained by solving these equations.

2.1.1 Conservation of Mass

The conservation of mass (or continuity) describes the spatial and temporal evolution of the fluid in a given control volume. The following equation represents continuity in tensor notation [50]:

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i}(\rho u_i) = 0, \quad (2.1)$$

where ρ is the fluid density and u_i is the fluid velocity in the i -direction. In this context, there is no creation or destruction of mass.

2.1.2 Conservation of Linear Momentum

The conservation of linear momentum relates surface and body forces to the motion and deformation of the fluid. The following equation represents conservation of linear momentum in tensor notation [50]:

$$\frac{\partial}{\partial t}(\rho u_j) + \frac{\partial}{\partial x_i}(\rho u_i u_j) = -\frac{\partial p}{\partial x_j} + \frac{\partial \tau_{i,j}}{\partial x_i} + F_{b_j}, \quad (2.2)$$

where p is the fluid pressure and $\tau_{i,j}$ is the viscous stress tensor. The terms on the left-hand side of Equation 2.2 represent local rate of change and convective transfer of momentum, respectively. The terms on the right-hand side of Equation 2.2 represent forces acting on the fluid, where the first term represents pressure force while the middle term represents forces due to viscous deformation. The F_{b_j} term represents external body forces, which in this study is solely gravitational force. This term can be recast as follows, where g_j is the standard gravity:

$$F_{b_j} = \rho g_j. \quad (2.3)$$

The viscous stress tensor, $\tau_{i,j}$, describes fluid deformation via shearing. It can be recast in terms of the strain rate tensor, S_{ij} , assuming a Newtonian and isotropic fluid:

$$\tau_{i,j} = \mu \left(2S_{ij} - \frac{2}{3} \delta_{ij} \frac{\partial u_k}{\partial x_k} \right), \quad (2.4)$$

where μ is the dynamic viscosity of the fluid and δ_{ij} is the Kronecker delta, which is a tensor whose components equal 1 when $i = j$ or zero when $i \neq j$. The strain rate tensor can be rewritten in terms of the rate of shear, as shown below:

$$S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right). \quad (2.5)$$

Using this information, the momentum equation can be rewritten as follows:

$$\frac{\partial}{\partial t}(\rho u_j) + \frac{\partial}{\partial x_i}(\rho u_i u_j) = -\frac{\partial p}{\partial x_j} + \frac{\partial}{\partial x_i} \left[\mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3} \mu \delta_{ij} \frac{\partial u_k}{\partial x_k} \right] + \rho g_j, \quad (2.6)$$

2.1.3 Conservation of Energy

The conservation of energy relates work and heat to the internal energy of the fluid. The following equation represents conservation of internal energy in tensor notation [50]:

$$\frac{\partial}{\partial t}(\rho e) + \frac{\partial}{\partial x_i}(\rho u_i e) = -\frac{\partial q_i}{\partial x_i} + \dot{\omega}_T + \dot{Q}_{rad} - \frac{\partial}{\partial x_i}(p x_i) + \frac{\partial}{\partial x_j}(u_i \tau_{i,j}). \quad (2.7)$$

where e is the internal energy. In combustion applications, however, it is more convenient to express the conservation of energy in terms of enthalpy, h , as follows:

$$h = e + \frac{p}{\rho}. \quad (2.8)$$

Using Equation 2.8, an equation for the conservation of enthalpy can be derived [50]:

$$\frac{\partial}{\partial t}(\rho h) + \frac{\partial}{\partial x_i}(\rho u_i h) = -\frac{\partial q_i}{\partial x_i} + \dot{\omega}_T + \dot{Q}_{rad} + \frac{\partial p}{\partial t} + \frac{\partial}{\partial x_j}(u_i \tau_{i,j}). \quad (2.9)$$

In Equation 2.9, the left-hand side describes enthalpy rate of change and convective enthalpy transport, respectively. The right-hand side describes work and heat fluxes to the system. $\dot{\omega}_T$ and $\dot{Q}_{rad}$ represent heat additions to the fluid due to the combustion reaction and radiation, respectively. The last two terms on the right-hand side represent work due to pressure and viscosity, respectively. The first term on the right-hand side represents energy dissipation, where q_i is the energy flux and is defined as follows [50]:

$$q_i = -\lambda \frac{\partial T}{\partial x_i} + \sum_{k=1}^n (h_k j_k), \quad (2.10)$$

where λ is the thermal conductivity and j_k is the diffusive species flux. The right hand side represents energy fluxes due to heat conduction (using Fourier's Law) and molecular diffusion, respectively. j_k can be rewritten by applying Fick's first law, as shown below:

$$j_k = -\rho D_k \frac{\partial Y_k}{\partial x_i}, \quad (2.11)$$

where D_k is the molecular diffusivity and Y_k is the mass fraction of species k . Using this information and the definitions of the Prandtl and Schmidt numbers, as well as assuming a Lewis number of unity, leads to the following expression for the energy flux, where ν is the kinematic viscosity and Pr is the Prandtl number:

$$q_i = -\rho \frac{\nu}{Pr} \frac{\partial h}{\partial x_i}. \quad (2.12)$$

Equation 2.9 can then be rewritten as follows:

$$\frac{\partial}{\partial t}(\rho h) + \frac{\partial}{\partial x_i}(\rho u_i h) = -\frac{\partial}{\partial x_i} \left(-\rho \frac{\nu}{Pr} \frac{\partial h}{\partial x_i} \right) + \dot{\omega}_T + \dot{Q}_{rad} + \frac{\partial p}{\partial t} + \frac{\partial}{\partial x_j}(u_i \tau_{i,j}). \quad (2.13)$$

2.1.4 Species Transport

The transport of species is important in the context of combustion, as it describes how the reaction proceeds and how products form. Considering mass fractions, each species k is governed by the following transport equation in tensor notation [50]:

$$\frac{\partial}{\partial t}(\rho Y_k) + \frac{\partial}{\partial x_i}(\rho u_i Y_k) = -\frac{\partial j_{k,i}}{\partial x_i} + \dot{\omega}_k, \quad (2.14)$$

where $j_{k,i}$ is the diffusive species flux for each direction as defined by Equation 2.11, and $\dot{\omega}_k$ is the reaction rate of species k . In addition, the mass fractions of all species must sum to unity as shown below, where N is the number of species:

$$\sum_{k=1}^N Y_k = 1 \quad (2.15)$$

2.2 Turbulent Flows

Equations 2.1, 2.6, 2.13 and 2.14 above describe the reacting fluid flow at any flow regime in theory. However, these equations are difficult to solve numerically due to the chaotic nature of turbulent flow. Turbulent flows are often solved by modelling part of the turbulent behaviour. This section describes turbulence and how it can be modelled in CFD using different frameworks.

2.2.1 Description of Turbulence

The flow regime of a fluid is often dictated by the Reynolds number, Re , which is the ratio of inertial forces to viscous forces in a fluid. All fluid flows can be classified as either laminar, transitional or turbulent. Laminar flow is dictated by smooth motion of fluid with little to no mixing; it occurs at relatively low Re and depends on the geometry. Transitional flow is an intermediate stage that occurs when small instabilities develop in laminar flows and impact the flow structure, creating mixing. Turbulent flow is characterized by the chaotic motion of fluid particles and the mixing of fluid layers; it occurs at relatively large Re and depends on the geometry. Pressure and velocity vary at each point in time due to the chaotic motion of the flow, which makes turbulent flows difficult to solve. However, turbulent flow is prevalent in nature and in many engineering applications. This includes

combustion, where turbulent flow is desired for mixing the fuel and oxidizer to facilitate reactions. Thus, it is important to understand and capture turbulence in combustion CFD.

Turbulence can be thought of as fluctuations superimposed on a mean flow. The fluctuations in the flow lead to the formation of eddies, which are flow structures characterized by swirling and reverse flow in the mean flow current. Turbulent kinetic energy (TKE) is generated by the velocity fluctuations, and TKE generation is associated with the largest eddies. However, these eddies are unstable and subsequently break up into smaller eddies until stability is achieved [51]. This leads to the formation of eddies of various sizes, which can be categorized into different length scales based on the wavenumber. Eddies are generally described by the integral length scale, the Taylor microscale, and the Kolmogorov length scale. The integral length scale is the largest length scale and the eddies in this scale contain the most TKE. It is categorized by the velocity fluctuations and is limited by geometry. TKE is transferred via an energy cascade to the Taylor microscale, which is smaller in size. Viscous effects are more prominent in the Taylor microscale, however dissipation of TKE does not occur. TKE is further transferred to the Kolmogorov scale, which is the smallest scale of energy. Viscous dissipation of TKE occurs at this scale due to the significance of molecular diffusion. Turbulent flow can therefore be thought of as a spectrum of eddies with varying length scales, where TKE is transferred from the largest to the smallest eddies before dissipation [51]. Note that turbulent flows also have a range of time scales, which are associated with the length scale ranges.

Different frameworks exist for representing turbulent flows in CFD. These frameworks can be classified based on the range of length scales that are modelled. Direct numerical simulation (DNS) is a framework in which the entire range of spatial and temporal turbulent scales are resolved. DNS is the most true relative to the applied formulas since it does not involve turbulence models. However, extremely fine spatial and temporal grids are needed to resolve the Kolmogorov scales and chemical time scales, making DNS computationally prohibitive for high Re flows. LES is a framework in which the small scales are modelled while the large scales are resolved directly. Various methods exist for filtering the scales and modelling the unclosed terms that arise from the filtering operation. LES has lower computational cost than DNS, but is nonetheless computationally expensive for most industrial applications. The Reynolds-Averaged Navier-Stokes (RANS) framework is used in this study, and is described in the following subsection.

2.2.2 Reynolds-Averaged Navier-Stokes

Reynolds-Averaged Navier-Stokes (RANS) is a framework in which all turbulent scales are modelled. This is achieved by applying decomposition to the instantaneous quantities

in the governing equations, in which these quantities are represented by a mean and a fluctuating component. The transport equations are then averaged to obtain expressions for the mean quantities. Turbulence models are required to close the Reynolds stresses, which arise from the decomposition and represent the turbulence.

The computational grid spacing in RANS is not explicitly related to the turbulence scales, although it is useful to refine the grid to capture wall effects and other combustion phenomena and obtain grid-independent results. This means that RANS is less computationally expensive than LES or DNS. RANS is also more straightforward to use than LES or DNS, as there is less user-input for aspects such as calculating the filter width or grid spacing. For these reasons, RANS is widely used for combustion CFD in industry and in some areas of research. RANS is selected for the current study, as it can provide reasonable accuracy of non-premixed turbulent jet flames [26–33]. The following subsection compares averaging methods for RANS.

2.2.3 Averaged Transport Equations

In the RANS framework, decomposition is applied to the governing equations to obtain transport equations for the mean quantities. This subsection describes and compares two commonly-used averaging techniques: Reynolds and Favre averaging.

For Reynolds averaging (or Reynolds decomposition), each transported quantity ϕ is split into a mean and a fluctuating component, as shown below:

$$\phi(x_i, t) = \bar{\phi}(x_i) + \phi'(x_i, t), \quad (2.16)$$

where $\phi(x_i, t)$ is an instantaneous quantity, $\bar{\phi}(x_i)$ is the time-averaged quantity, and $\phi'(x_i, t)$ is the fluctuation of the quantity about the mean. This decomposition is applied to the flow variables (i.e., u_i , h , Y_k and $\dot{\omega}_k$) in the governing equations. The RANS equations are obtained by taking the mean of each term (not shown for brevity).

The decomposition process leads to the Reynolds stress term $\frac{\partial}{\partial x_j}(\overline{\rho u'_i u'_j})$, which is unclosed and must be modelled. Reynolds averaging is more useful for cases with small density variations, such as incompressible flows. However, the density varies significantly in combustion due to the nature of the flow and the different types of species involved. The inclusion of density fluctuations would lead to many additional unclosed terms in all four governing equations, all of which require models or correlations [50]. For this reason, Reynolds averaging is not used in the current study.

An alternative approach to decomposition is Favre averaging [52], which splits an instantaneous quantity into mean and fluctuating components based on a density-weighted average, as shown below:

$$\phi(x_i, t) = \tilde{\phi}(x_i) + \phi''(x_i, t). \quad (2.17)$$

In this equation, $\phi''(x_i, t)$ is the fluctuation of a given quantity about the Favre average, $\tilde{\phi}(x_i)$, which is defined below:

$$\tilde{\phi}(x_i) = \frac{\overline{\rho\phi}(x_i)}{\bar{\rho}}. \quad (2.18)$$

The governing equations are decomposed using Equation 2.17. The following identities are used for simplification:

$$\tilde{\tilde{\phi}}(x_i) = \tilde{\phi}(x_i), \tilde{\phi}''(x_i, t) = 0, \overline{\rho\phi''}(x_i, t) = 0. \quad (2.19)$$

This results in the following equations for conservation of mass, linear momentum, enthalpy, and species transport, respectively [50]:

$$\frac{\partial \bar{\rho}}{\partial t} + \frac{\partial}{\partial x_i}(\bar{\rho}\tilde{u}_i) = 0, \quad (2.20)$$

$$\frac{\partial}{\partial t}(\bar{\rho}\tilde{u}_i) + \frac{\partial}{\partial x_i}(\bar{\rho}\tilde{u}_i\tilde{u}_j) = -\frac{\partial \bar{p}}{\partial x_j} + \frac{\partial \overline{\tau_{i,j}}}{\partial x_j} + \bar{\rho}g_j - \frac{\partial}{\partial x_i}(\bar{\rho}\widetilde{u_i''u_j''}), \quad (2.21)$$

$$\frac{\partial}{\partial t}(\bar{\rho}\tilde{h}) + \frac{\partial}{\partial x_i}(\bar{\rho}\tilde{u}_i\tilde{h}) = -\frac{\partial \bar{q}_i}{\partial x_i} + \tilde{\omega}_T + \tilde{Q}_{rad} + \frac{\partial \bar{p}}{\partial t} + \overline{\tau_{ij}\frac{\partial u_i}{\partial x_j}} - \frac{\partial}{\partial x_i}(\bar{\rho}\widetilde{u_i''h''}), \quad (2.22)$$

$$\frac{\partial}{\partial t}(\bar{\rho}\tilde{Y}_k) + \frac{\partial}{\partial x_i}(\bar{\rho}\tilde{u}_i\tilde{Y}_k) = \frac{\partial}{\partial x_i}\left(\overline{\rho D_k\frac{\partial Y_k}{\partial x_i}}\right) + \tilde{\omega}_k - \frac{\partial}{\partial x_i}(\bar{\rho}\widetilde{u_i''Y_k''}). \quad (2.23)$$

The use of Favre averaging removes unclosed terms associated with density fluctuations. There are also fewer unclosed terms for the Reynolds stresses when expanded. For these reasons, Favre averaging is preferred in turbulent combustion CFD and is used in the current study. Models are required to close the three fluctuating terms (last term on Equations 2.21 to 2.23) and the reaction rate terms ($\tilde{\omega}_k$ in Equation 2.23, and $\tilde{\omega}_T$ and $\tilde{Q}_{rad}$ in Equation 2.22). There are multiple ways to close each term. The following section discusses and compares turbulence models for closing the Reynolds stress term $\frac{\partial}{\partial x_i}(\bar{\rho}\widetilde{u_i''u_j''})$.

2.3 Eddy Viscosity Turbulence Models

There are various ways to close the Reynolds stresses arising from RANS using a turbulence model. This section describes and compares four common turbulence models, with an emphasis on the k - ϵ model. These models rely on the eddy viscosity hypothesis proposed by Boussinesq, where the eddy viscosity models the energy cascade through the range of eddies. This hypothesis relates the Reynolds stresses to the mean fluid deformations, resulting in the following expression for the Reynolds stress:

$$-\widetilde{u_i''u_j''} = \nu_t \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3}k\delta_{ij}, \quad (2.24)$$

where k in this context is the TKE and ν_t represents the kinematic eddy viscosity. k is commonly represented in terms of velocity fluctuations, as shown below:

$$k = \frac{1}{2} \sum_{n=1}^3 \widetilde{(u_n''^2)}. \quad (2.25)$$

Note that ρ is divided for convenience in Equation 2.24, resulting in $\nu_t = \frac{\mu_t}{\rho}$, where μ_t is the dynamic eddy viscosity. The predominant focus of turbulence models is to provide relations for ν_t to close the Reynolds stresses.

2.3.1 Mixing Length Model

One of the first turbulence models is the mixing length model, which was proposed by Prandtl in 1925 [53]. This model was derived from wall-bounded flows (and ties in with boundary layer theory), where it assumes that the eddy viscosity varies with distance from the wall. The eddy viscosity is modelled as the product of a velocity scale and a length scale. It is assumed that the only significant velocity gradient is the streamwise component proportional to the wall (i.e., $\widetilde{\frac{\partial u_i}{\partial x_j}}$). The kinematic eddy viscosity is then represented with the following relation, where l_m is the mixing length. Note that the absolute value of the velocity gradient is taken to produce a positive viscosity.

$$\nu_t = l_m \left| \widetilde{\frac{\partial u_i}{\partial x_j}} \right|. \quad (2.26)$$

The mixing length model is computationally inexpensive since it does not require additional transport equations. It is also straightforward to use since l_m is the only free parameter.

However, the model does not capture complex turbulence phenomena due to the assumption that the rates of production and dissipation of TKE are equal. The accuracy also depends on the selection of l_m , which depends on prior knowledge [51]. For these reasons, the mixing length model is not used in the current study.

2.3.2 k - ϵ Model

The k - ϵ model [54] is one of the most common turbulence models in CFD. It was originally developed as an improvement to the mixing length model by incorporating physical mechanisms that affect the TKE and avoiding the need to specify a mixing length [51]. It is a two-equation model, which indicates that transport equations are solved for TKE (k) and the dissipation rate of TKE (ϵ); this makes it more computationally expensive than the mixing length model. The eddy viscosity is modelled as a product of velocity and length scales as in the mixing length model, however these scales are represented in terms of k and ϵ in the k - ϵ model. This results in the following expression for ν_t , where C_μ is a dimensionless constant:

$$\nu_t = C_\mu \frac{k^2}{\epsilon}. \quad (2.27)$$

Using this identity, the transport equations for k and ϵ can be represented with the following semi-empirical equations, respectively:

$$\frac{\partial k}{\partial t} + \frac{\partial}{\partial x_i}(\tilde{u}_i k) = \frac{\partial}{\partial x_i} \left[\left(\nu + \frac{\nu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_i} \right] + P_k - \epsilon, \quad (2.28)$$

$$\frac{\partial \epsilon}{\partial t} + \frac{\partial}{\partial x_i}(\tilde{u}_i \epsilon) = \frac{\partial}{\partial x_i} \left[\left(\nu + \frac{\nu_t}{\sigma_\epsilon} \right) \frac{\partial \epsilon}{\partial x_i} \right] + C_{\epsilon_1} \frac{\epsilon}{k} P_k - C_{\epsilon_2} \frac{\epsilon^2}{k}. \quad (2.29)$$

In these equations, P_k represents the production of TKE, as shown below:

$$P_k = -\widetilde{u''_i u''_j} \frac{\partial \tilde{u}_i}{\partial x_j}. \quad (2.30)$$

Equations 2.27, 2.28 and 2.29 contain dimensionless constants C_μ , σ_k , σ_ϵ , C_{ϵ_1} and C_{ϵ_2} that are adjustable. However, the following standard values have been determined from extensive testing:

$$C_\mu = 0.09, \sigma_k = 1.0, \sigma_\epsilon = 1.30, C_{\epsilon_1} = 1.92, C_{\epsilon_2} = 1.44. \quad (2.31)$$

The standard k - ϵ model (as formulated above) is one of the simplest two-equation turbulence models. The standard k - ϵ model specifically performs well for free-shear flows,

including mixing layers, plane jets and boundary layers. It does not perform well in cases with large pressure gradients, and wall functions are required to capture the production of TKE. One well-known anomaly is the spreading rate of round jets [51], which can be addressed by adjusting the constants in Equation 2.31. Despite these limitations, the standard k - ϵ model is commonly used since it is relatively robust for various industrial and environmental flows.

The standard k - ϵ model with modified coefficients is selected in this study, following previous work [26–33]. However, alternative formulations of the k - ϵ model exist in an effort to address the limitations of the standard k - ϵ model. Two common alternatives include the RNG k - ϵ model [55] and the realizable k - ϵ model [56]. The RNG (re-normalization group) k - ϵ model is derived from statistics and includes an additional term in the ϵ transport equation to account for strain. The definition of each constant is modified, with analytical expressions being provided for σ_k and σ_ϵ . The RNG k - ϵ model performs better than the standard k - ϵ model for high-strain flows, and can account for swirl with some modifications to ν_t . The realizable k - ϵ model imposes mathematical constraints on the Reynolds stresses to match physical expectations. In this model, ν_t is limited such that the normal stress ($\overline{u''^2}$) remains positive. In addition, $C\mu$ becomes a variable quantity, which is determined from the strain rate, and the ϵ transport equation is derived from an exact equation rather than being posed from empirical reasoning. The realizable k - ϵ model addresses the round jet anomaly and has been shown to give improved results with large pressure gradients.

2.3.3 k - ω Model

The k - ω model is a two-equation turbulence model proposed by Wilcox in 1988 [57]. The k - ω model considers the turbulent frequency, which is defined as $\omega = \frac{k}{\epsilon}$, rather than ϵ itself. The eddy viscosity is modified from the k - ϵ formulation to contain ω , as shown below. Note that ν_t does not contain an adjustable constant.

$$\nu_t = \frac{k}{\omega}. \quad (2.32)$$

The transport equations for k and ω are formulated as follows:

$$\frac{\partial k}{\partial t} + \frac{\partial}{\partial x_i}(\tilde{u}_i k) = \frac{\partial}{\partial x_i} \left[(\nu + \nu_t \sigma_k) \frac{\partial k}{\partial x_i} \right] + P_k - \beta_k k \omega, \quad (2.33)$$

$$\frac{\partial \omega}{\partial t} + \frac{\partial}{\partial x_i}(\tilde{u}_i \omega) = \frac{\partial}{\partial x_i} \left[(\nu + \nu_t \sigma_\omega) \frac{\partial \omega}{\partial x_i} \right] + \frac{\alpha}{\nu_t} \tau_{ij} \frac{\partial \tilde{u}_j}{\partial x_i} - \beta_\omega \omega^2. \quad (2.34)$$

Equations 2.33 and 2.34 contain dimensionless constants σ_k , σ_ω , β_k , β_ω , and α that are adjustable. The following values are taken from the original k - ω paper, however the constants can be specified using the strain rate tensor [58]:

$$\sigma_k = 0.5, \sigma_\omega = 0.5, \beta_k = 0.09, \beta_\omega = \frac{3}{40}, \alpha = \frac{5}{9}. \quad (2.35)$$

The k - ω model does not require wall functions and is more accurate near walls than the standard k - ϵ model, however a finer computational grid is needed with the k - ω model since ω is sensitive near walls. The k - ω model is reasonable for a variety of turbulent flows, however it is sensitive to ν_t in the free stream.

2.3.4 Shear Stress Transport Model

The Shear Stress Transport (SST) turbulence model was proposed by Menter in 1994 [59]. The SST model is a combination of the standard k - ϵ model and the k - ω model, where k - ω is used within the boundary layer and k - ϵ is used in the free stream. The eddy viscosity is modelled with a blending function in the SST model. Transport equations for k and ω can be derived by combining the individual transport equations from the k - ϵ and k - ω models. The SST model combines the benefits of the k - ϵ and k - ω models and is better at predicting separation. However, the SST model contains more constants than both of the k - ϵ or k - ω models, which makes the SST model more difficult to use.

2.4 Turbulent Non-premixed Combustion Modelling

The chemical source terms in Equations 2.22 and 2.23 are highly non-linear and require a turbulent combustion model for closure. This section first briefly describes and compares various combustion models. Important parameters for non-premixed turbulent combustion are also introduced here.

2.4.1 Turbulent Combustion Models

Various turbulent combustion models exist, which may be classified into groups based on the underlying assumptions. Models that assume infinitely-fast chemistry include the Eddy Break-up Model (EBU) [2, 3], Eddy Dissipation Model (EDM) [4] and Eddy Dissipation Concept (EDC) [5]. These models assume that the chemical scales are fast enough such that

they are decoupled from the turbulence scales, thus the reaction rate is mixing controlled. These models are commonly used due to their simplicity and small computational cost. However, these models cannot account for detailed chemical kinetics, and the ignition is difficult to control.

Flamelet models assume that chemical time scales are short enough such that reactions take place in thin reactive-diffusive layers known as ‘flamelets’, which are smaller than the smallest turbulent scales. With this assumption, the flamelets are not affected by turbulence and can thus be considered laminar [15]. For non-premixed combustion, this allows for a coordinate transform of the scalar transport equations, where mixture fraction Z becomes the independent variable. The flamelet solution can be pre-tabulated by parameterizing Z and scalar dissipation rate χ and assuming fixed boundary conditions. This constitutes the basis of the non-premixed steady laminar flamelet model (SLFM). The main advantage of flamelet models is the inclusion of detailed chemistry with relatively small computational cost due to the pre-tabulation. However, they are limited by the small chemical time scale assumption, which is often not verified. This implicates the assumption of laminar flow on the model and hence the formulation it relies on.

Another set of models are constructed on probability density function (PDF) correlations between flow variables. The transported PDF approach [6, 60, 61] involves the solution of a joint-PDF that describes the state of the reacting medium. The velocity components, species mass fractions and enthalpy are often considered in the joint-PDF. A transport equation for the joint-PDF is then solved. This model directly resolves the chemical source terms at the expense of creating additional unclosed terms that are difficult to model. Transported PDF is computationally expensive due to the high-dimensionality of the joint PDF, often requiring stochastic simulations. To reduce computational cost, the shape of the joint-PDF is often presumed, leading to the following presumed-PDF methods. Conditional Source-term Estimation (CSE) [9, 10] and Conditional Moment Closure (CMC) [7] use conditional averages of reactive scalars to close the chemical source terms. The Favre averages are related to the conditional averages through the PDF. Both CSE and CMC rely on the first-order closure assumption (described in Section 2.5.1 below), but differ in the determination of the conditional averages. Multiple Mapping Conditioning [8] is effectively a combination of the CMC and transported PDF approaches, in which the CMC first-order closure assumption is applied to limit the dimensionality of the joint-PDF by correlating certain variables together. CSE for non-premixed combustion is described thoroughly in Section 2.5.

2.4.2 Mixture Fraction

The current work involves non-premixed combustion, which is indicated by separate fuel and oxidizer streams. Non-premixed combustion is a mixing-controlled process since the mixing time scale is larger than the chemical time scale. For this reason, it is useful and reasonable to use a scalar that describes the state of mixing to indicate where combustion occurs. This is referred to as the mixture fraction, Z , which is defined to have a value of zero for pure oxidizer and a value of unity for pure fuel. The formal definition of Z from Bilger [62] is adopted, as shown below:

$$Z = \frac{Y_i - Y_{i,O}}{Y_{i,F} - Y_{i,O}}. \quad (2.36)$$

In this equation, Y_i is the mass fraction of element i , while $Y_{i,F}$ and $Y_{i,O}$ refer to the mass fractions of element i in the fuel and oxidizer streams, respectively. The elements considered in Z depend on the fuel, however it is common to consider carbon, hydrogen and oxygen in Z . The mixture fraction is a conserved scalar, meaning that it cannot be produced or destroyed unless a phase change is present.

The mixture fraction is used for constructing the PDF in the PDF-based combustion models. The mean ($\tilde{Z}$) mixture fraction is often used to describe the PDF, however $\tilde{Z}$ by itself is not sufficient to represent the chemical source terms. Thus, the variance ($\widetilde{Z'^2}$) of the mixture fraction is also required to construct the PDF. $\tilde{Z}$ and $\widetilde{Z'^2}$ are represented by Favre-averaged transport equations, as shown below:

$$\frac{\partial}{\partial t}(\bar{\rho}\tilde{Z}) + \frac{\partial}{\partial x_i}(\bar{\rho}\tilde{u}_i\tilde{Z}) = \frac{\partial}{\partial x_i}\left(\bar{\rho}D\frac{\partial\tilde{Z}}{\partial x_i}\right) - \frac{\partial}{\partial x_i}(\bar{\rho}\widetilde{u_i''Z''}), \quad (2.37)$$

$$\frac{\partial}{\partial t}(\bar{\rho}\widetilde{Z'^2}) + \frac{\partial}{\partial x_i}(\bar{\rho}\tilde{u}_i\widetilde{Z'^2}) = \frac{\partial}{\partial x_i}\left(\bar{\rho}D\frac{\partial\widetilde{Z'^2}}{\partial x_i}\right) - \frac{\partial}{\partial x_i}\left[\bar{\rho}\widetilde{u_i''(Z'^2)''}\right] + 2\bar{\rho}D_t\left(\frac{\partial\tilde{Z}}{\partial x_i}\right) - \bar{\rho}\tilde{\chi}, \quad (2.38)$$

where D is the mixture fraction diffusivity and $\tilde{\chi}$ is the scalar dissipation rate. The terms containing u_i'' are unclosed and require modelling. The gradient diffusion hypothesis is used as shown below, where D_t is the turbulent diffusivity:

$$-\widetilde{u_i''Z''} = D_t\frac{\partial\tilde{Z}}{\partial x_i}. \quad (2.39)$$

The scalar dissipation rate is another unclosed term that requires modelling. Assuming the production rate of scalar variance equals the dissipation rate, $\tilde{\chi}$ can be represented in

terms of $\widetilde{Z''^2}$, as shown below:

$$\tilde{\chi} = 2D_\chi \frac{\partial \widetilde{Z''^2}}{\partial x_i}, \quad (2.40)$$

where D_χ is the scalar diffusivity. χ is an important parameter in non-premixed combustion, as it relates molecular diffusion to the state of turbulent mixing [63] and describes how the combustion occurs. A small value of χ implies that the turbulent mixing rate is slow, resulting in the reaction rates approaching the infinitely-fast chemistry assumption. A large value of χ implies that the turbulent mixing rate is fast, resulting in small reaction rates due to heat loss and possible extinction. A value of unity for χ at the stoichiometric mixture fraction indicates that combustion occurs at equilibrium.

It should be noted that in premixed combustion, which occurs when fuel and oxidizer are mixed before ignition, the state of reaction is described by the progress variable c . A value of zero indicates fresh reactants, while a value of unity indicates burnt gases. The progress variable can be defined multiple ways, often being constructed from the mass fractions of one or more characteristic species.

2.4.3 Beta Probability Density Function (PDF)

In methods such as CSE, CMC and MMC, the shape of the PDF describing the reacting medium is presumed rather than solved. It is common to assume a β -distribution for non-premixed combustion [64]. The β -distribution of a quantity ϕ is represented by the following equation:

$$P(\phi) = \frac{\Gamma(\beta_1 + \beta_2)}{\Gamma(\beta_1)\Gamma(\beta_2)} \phi^{\beta_1-1} (\phi - 1)^{\beta_2-1}. \quad (2.41)$$

The gamma function Γ is defined as the following for a given β :

$$\Gamma(\beta) = \int_0^\infty x^{\beta-1} e^{-x} dx, \quad (2.42)$$

while the beta-function parameters β_1 and β_2 are determined as follows:

$$\beta_1 = \tilde{Z} \left(\frac{\tilde{Z}(1 - \tilde{Z})}{\widetilde{Z''^2}} - 1 \right), \quad (2.43)$$

$$\beta_2 = (1 - \tilde{Z}) \left(\frac{\tilde{Z}(1 - \tilde{Z})}{\widetilde{Z''^2}} - 1 \right). \quad (2.44)$$

2.5 Conditional Source-term Estimation for Turbulent Non-premixed Combustion

This section describes the Conditional Source-term Estimation (CSE) combustion model in the context of turbulent non-premixed combustion. The basic principles and formulation of CSE is first discussed, followed by integral inversion techniques. Methods for chemistry representation in CSE are then presented and compared. The implementation of CSE in the CFD code is then described.

2.5.1 CSE Principles

The principles of the CSE model are identical to those in the CMC model, from which CSE is derived. The following analysis considers non-premixed combustion in the RANS framework, however the extensions to both premixed combustion and the LES framework are straightforward.

The Conditional Source-term Estimation (CSE) model was first developed by Bushe and Steiner in 1999 [9] for LES of non-premixed flames. CSE relies on the first-order closure assumption, which states that spatial and temporal fluctuations of the reactive scalars are correlated to the fluctuations of one or more characteristic variables, which are used for conditioning, and that the conditional fluctuations of the reactive scalars are much smaller than the unconditional fluctuations. The first-order closure assumption has been shown to be reasonable for attached flames without ignition or extinction [7]. Using the first-order closure assumption, the conditional source terms are evaluated directly from the conditional species mass fractions, temperature and density, as shown below:

$$\widetilde{\dot{\omega}_k|\zeta} = \dot{\omega}_k(\widetilde{Y_1|\zeta}, \widetilde{Y_2|\zeta}, \dots, \widetilde{Y_N|\zeta}, \widetilde{T|\zeta}, \overline{\rho|\zeta}), \quad (2.45)$$

where $\widetilde{\dot{\omega}_k|\zeta}$ is the conditional source term of species k , ζ is the sample space variable for Z , N is the number of species, $\overline{\rho|\zeta}$ is the conditional Reynolds average of density, and $\widetilde{T|\zeta}$ and $\widetilde{Y|\zeta}$ are the conditional Favre averages of temperature and species mass fractions, respectively. Note that the mixture fraction Z is selected as the conditioning variable, since it has been found to satisfy the first-order closure assumption for previous non-premixed combustion cases with CSE [25–33]. One conditioning variable is sufficient for most cases, however multiple conditioning variables can be included to capture additional physics [10].

The CSE model involves both Favre averaged and conditionally averaged variables. The Favre average of a given quantity ϕ can be determined from the respective conditional

average as follows:

$$\tilde{\phi}(i, t) = \int_0^1 \tilde{\phi}|\zeta(i, t) \tilde{P}(\zeta, i, t) d\zeta, \quad (2.46)$$

where $\tilde{\phi}(i, t)$ is a reactive scalar (i.e., $\tilde{\omega}_k$ or $\tilde{Y}_k$) at a given cell i and a given time t . $\tilde{P}(\zeta, i, t)$ is the PDF as described by Equation 2.41 for a given discretized mixture fraction ζ , cell i and time t .

In CSE, transport equations are solved for the Favre averages of species mass fractions (refer to Equation 2.23). To determine the unclosed chemical source terms, first the conditional species mass fractions must be determined. This is done by inverting Equation 2.46 to solve for each conditional species mass fraction, using the known Favre-averaged species mass fractions. Equation 2.45 is then applied to determine the conditional source terms from the conditional species mass fractions. The Favre-averaged source terms can then be determined from the conditional source terms by applying Equation 2.46 directly. This is in contrast to CMC, which solves additional transport equations for the conditional averages on a separate grid. Integral inversion in CSE results in reduced computational cost and avoids the need to solve for additional unclosed terms as in CMC. The following section described the integral inversion process.

2.5.2 Integral Inversion and Regularization

Integral inversion is applied in CSE to determine the conditional reactive scalars for closing the conditional source terms. Instead of defining a coarse grid as in CMC, a number of ensembles composed of CFD cells are defined in CSE, in which integral inversion is performed in each ensemble. The conditional averages are assumed to be homogeneous within a given ensemble, and are thus assumed to have no spatial dependence. This is a valid assumption due to the observation that conditional averages have low spatial variance [7, 65]. Using this assumption, Equation 2.46 can be modified to the following expression:

$$\tilde{\phi}(i, t) = \int_0^1 \tilde{\phi}|\zeta(t) \tilde{P}(\zeta, i, t) d\zeta, \quad (2.47)$$

where ϕ in this context specifically refers to the species mass fractions, temperature and density. Note that the spatial dependence in the conditional averages ($\tilde{\phi}|\zeta(t)$) has been removed in this form. Equation 2.47 is a Fredholm integral equation of the first kind with $\tilde{P}(\zeta, i, t)$ as the kernel. $\tilde{P}(\zeta, i, t)$ can be discretized into a set number of bins in the conditioning variable space (ζ in this context), allowing for Equation 2.47 to be converted

into matrix form, as shown below:

$$\mathbb{P}_{ij}\alpha_j = \varphi_i, \quad (2.48)$$

where α_j is the discrete form of $\widetilde{\phi}|\zeta$ in the j th bin in conditional space, $\mathbb{P}_{ij}$ is the discrete form of $\widetilde{P}(\zeta, i, t)$ and $\varphi_i = \widetilde{\phi}(i, t)$, i.e., the Favre averages obtained from transport equations. $\mathbb{P}_{ij}$ is determined from the following expression:

$$\mathbb{P}_{ij} = \int_{\zeta_{j-\frac{1}{2}}}^{\zeta_{j+\frac{1}{2}}} \widetilde{P}(\zeta, i, t) d\zeta. \quad (2.49)$$

$\mathbb{P}_{ij}$ is a N_{cells} by N_{bins} matrix containing the integrated PDF over a ζ interval, where N_{cells} is the number of CFD cells in a given ensemble and N_{bins} is the number of bins in the ζ space. The ζ interval is defined by the upper and lower bin limits of $j + \frac{1}{2}$ and $j - \frac{1}{2}$, respectively, which are placed at the midpoints of the discretized ζ space.

Knowing $\mathbb{P}_{ij}$ and $\widetilde{\phi}|\zeta$, it is required to solve for $\vec{\alpha}_j$. However, Equation 2.48 is ill-posed, i.e., it is sensitive to small perturbations in $\mathbb{P}_{ij}$ and $\widetilde{\phi}|\zeta$ [66]; this is because there are more knowns (i.e., unconditional mass fractions of a given species in a CSE ensemble) than unknowns (i.e., conditional mass fractions of the same species in the same CSE ensemble) in Equation 2.48. This means that direct solution methods cannot be used. Instead, a regularization method is required to obtain a unique solution of α_j that is smooth and stable. Several regularization methods exist, including Tikhonov regularization [67] and truncated singular value decomposition (SVD) [68]. Zeroth-order Tikhonov regularization is adopted in the current study due to previous successful implementations in CSE [10]. Equation 2.48 can then be transformed into the following:

$$\vec{\alpha} = \underset{\vec{\alpha}^*}{\operatorname{argmin}} \left\| \left[\frac{\mathbb{P}}{\lambda \mathbb{I}} \right] \vec{\alpha}^* - \left[\frac{\vec{\varphi}}{\lambda \vec{\alpha}^0} \right] \right\|_2^2, \quad (2.50)$$

where ‘argmin’ is the argument of the minimum of the function, $\mathbb{I}$ is the identity matrix with dimensions of N_{bins} by N_{bins} , $\vec{\alpha}^*$ is the argument of the function, and $\vec{\alpha}^0$ is *a priori* knowledge of the solution to the system. In this work, $\vec{\alpha}^0$ is obtained from the previous time step, however it is possible to use an unstrained one-dimensional laminar flame [19]. The regularization parameter, λ , is calculated as follows [27], where ‘Tr’ is the trace of the matrix:

$$\lambda = \frac{\operatorname{Tr}(\mathbb{P}^T \mathbb{P})}{\operatorname{Tr}(\mathbb{I})}. \quad (2.51)$$

Note that another option is first-order Tikhonov regularization, which has been successfully applied to CSE [32, 33]. First-order Tikhonov regularization follows the same form, but

the discrete derivative operator $\mathbb{D}$ is used in Equation 2.50 instead of the identity matrix $\mathbb{I}$. First-order Tikhonov regularization provides smoother results and relies less on *a priori* information [27, 69], however it is much more complex than zeroth-order Tikhonov regularization, which has been shown to work in CSE for various cases [10]. Thus, first-order Tikhonov regularization is not considered.

2.5.3 Chemistry Representation

CSE allows for the inclusion of detailed chemical kinetics, which can be done using various methods. Thus, it is useful to compare the benefits and implications of these methods. This section describes tabulated chemistry and direct chemistry briefly. Chemistry representation with ML methods, which is the main focus of this study, is described in Chapter 4. Note that methane fuel is considered, but the conclusions extend to other fuels.

In the direct chemistry approach, transport equations are solved for the Favre averages of each species. Integral inversion as described in Section 2.5.2 is then performed to obtain the conditional averages of each species. Each conditional source term is obtained based on all of the conditional species mass fractions. The Favre averaged source terms are then obtained directly using Equation 2.47.

For a given chemical mechanism, direct chemistry is the most accurate approach in theory. However, the computational cost of direct chemistry is prohibitive in many cases. Consider that the number of species is on the order of 10 to 100 for detailed mechanisms; GRI-Mech 3.0 [70], for example, contains 53 species. This means that 53 species transport equations and integral inversions must be solved at each time step. This is not feasible on a standard computer, considering that the time step is very small. In addition, each inversion imposes a small numerical error on the conditional conservation of mass, which can accumulate with a large number of inversions. However, conditional conservation of mass can be enforced using a variety of techniques [33, 69]. Reduced mechanisms can be applied in direct chemistry to solve for fewer species (on the order of 10), however this can compromise the accuracy of the combustion model. Sandia flame D has been successfully simulated using RANS-CSE with reduced mechanism direct chemistry [33], however the majority of CSE studies have relied on tabulated chemistry, which is discussed below.

The motivation of the tabulated chemistry approach is to include the effects of detailed chemistry with reduced computational cost. This is achieved by solving a canonical combustion problem and pre-tabulating the solution, which is then looked-up during the simulation. The majority of these approaches involve mapping a high-dimensional space onto a low-dimensional manifold. Existing methods include Intrinsic Low-Dimensional

Manifolds (ILDm) [11, 12], Trajectory-Generated Low-Dimensional Manifolds (TGLDM) [13], Flamelet Generated Manifolds [16] and In-Situ Adaptive Tabulation (ISAT) [14]. TGLDM is examined in detail since it has been used successfully in previous CSE studies [10], however the other three tabulation methods will be described briefly.

Different flamelet data generation techniques are available, leading to different types of data sets such as Flamelet Generated Manifolds (FGM) and non-premixed flamelet data in SLFM. These techniques have been used successfully in some previous CSE studies [19, 22, 27, 37], but they are not the predominant methods in CSE. In-Situ Adaptive Tabulation (ISAT) [14] is an alternative approach that operates on the accessed region of the combustion space and tabulates information as it appears. One advantage of ISAT is that the discretization in non-linear regions can be refined, whereas manifold methods have a set refinement. However, the computational cost is increased since the tables are constructed in real time, and the storage can be on the same order as other tabulation methods.

The most common chemistry tabulation method in CSE is Trajectory-Generated Low-Dimensional Manifolds (TGLDM) [13]. TGLDM is an extension of the Intrinsic Low-Dimensional Manifolds (ILDm) method [11, 12] that aims to improve convergence and simplify the manifolds. The motivation of TGLDM is to develop a general procedure for simplifying chemical kinetics that does not rely on the partial equilibrium and steady state assumptions, which limit the applicable combustion range, and instead can be specified from a mechanism and the desired number of degrees-of-freedom [11, 12]. The number of parameters can be set by the user; in the current work, the discretized ζ space is selected as one parameter, while the conditional carbon dioxide (CO_2) and water (H_2O) mass fractions ($\widetilde{Y_{\text{CO}_2}|\zeta}$ and $\widetilde{Y_{\text{H}_2\text{O}}|\zeta}$) are selected as the other two parameters due to their relatively long formation times compared to other species [25]. In order to tabulate quantities as functions of these parameters, it is assumed that the fast chemical time scales are decoupled from the turbulence scales and can therefore be removed [13].

In the TGLDM approach, a manifold is generated for each discretized ζ bin value. The parameterized $\widetilde{Y_{\text{CO}_2}|\zeta}$ and $\widetilde{Y_{\text{H}_2\text{O}}|\zeta}$ variables are used as the solution space. The manifold boundaries are determined based on the conservation of chemical elements for a given ζ . Initial points for the trajectories are seeded along the boundaries, while the initial conditions are supplied from a mass balance of the composition. The spacing of initial points ensures complete coverage of the realizable composition space. The trajectories are solved from the initial point to the chemical equilibrium point using the following ordinary

differential equation for each species k , where $\dot{\omega}_k$ is determined from the mechanism:

$$\rho \frac{dY_k}{dt} = \dot{\omega}_k. \quad (2.52)$$

Note that Equation 2.52 is derived from a constant-pressure reactor and can be solved using a stiff ODE solver. An illustration of a TGLDM manifold is provided in Figure 2.1.




Figure 2.1: Illustration of a TGLDM manifold for methane-air combustion at the stoichiometric mixture fraction ($\zeta = 0.055$)

Once Equation 2.52 is solved, redundant points are removed to obtain unique values. The reactive scalars can then be tabulated based on the parameterized ζ , $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$. Delaunay triangulation [71] is used to easily interpolate between tabulated values during the look-up process in a simulation. The same TGLDM manifold after the triangulation is shown in Figure 2.2. An example of a complete TGLDM manifold for the CO_2 reaction rate is shown in Figure 2.3.

The advantage of chemistry tabulation is that conditional source terms are no longer functions of each species mass fraction, instead being functions of only $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$. Equation 2.45 can therefore be reduced to the following:

$$\widetilde{\dot{\omega}_k|\zeta} \approx \dot{\omega}_k(\widetilde{Y_{CO_2}|\zeta}, \widetilde{Y_{H_2O}|\zeta}). \quad (2.53)$$




Figure 2.2: Illustration of a TGLDM manifold with Delaunay triangulation for pure methane-air combustion at the stoichiometric mixture fraction ($\zeta = 0.055$)

The Favre averaged source terms can then be calculated directly using Equation 2.47. Another advantage of chemistry tabulation is that integral inversion and transport equations are only required for the parameterized species (i.e., CO_2 and H_2O); the remaining Favre averaged species mass fractions can be determined using Equation 2.47 directly. Thus, tabulated chemistry approaches reduce computational cost and impose minimal errors on the conditional conservation of mass, since fewer integral inversions are performed.

The tabulated chemistry approach allows for the inclusion of detailed chemical kinetics in a computationally efficient manner. However, tabulated chemistry methods are generally limited by three factors: applicability, convergence and storage requirements. Methods such as ILDM, TGLDM and FGM rely on dimensionality reduction, which is imposed using various assumptions and simplified models. In addition, complex physics, such as radiative heat loss and diffusive effects, are not included in the base formulations and require modifications for data generation [31, 40, 44–46]. Tabulated chemistry methods are therefore limited to cases where the underlying assumptions are valid.

There is also no guaranteed convergence for the tabulated chemistry methods. This can occur if the trajectories either do not progress, or form outside of the realizable region, due to reaction rate estimates. Complex reaction mechanisms are more difficult to solve due to the coupling of various reactions. TGLDM, in particular, relies on conservation of species to set boundary conditions for the trajectories, which are more difficult to define with multi-fuel blends or a significant number of species [30]. The underlying equations are also more difficult to solve with the inclusion of complex physical effects. While tabulated




Figure 2.3: Illustration of the carbon dioxide reaction rate mapped onto the TGLDM manifold

chemistry methods work in theory for any fuel, they have only been applied to simple fuels due to these limitations.

The storage requirement is the main practical limitation of the tabulated chemistry approach. The table size increases exponentially with each reactive scalar and the amount of precision (i.e., number of trajectories, amount of points within trajectories) required. The storage size of these tables can reach the order of gigabytes (GB), which often exceeds the RAM (Random Access Memory) on a given computer. Consequently, the table size must be reduced to a tractable amount, either by removing species or reducing precision. In a given mechanism, the number of tabulated species can be reduced to the order of 10 with reasonable storage size. Generally, the major species involved in the combustion, along with a set number of minor species, are kept in the tabulation. This results in a simulation including only the reduced number of tabulated species plus the parameterized species. Although the mass fractions and reaction rates of each tabulated species are determined using the entire mechanism, the species that are excluded are non-trivial and will impact the combustion. The alternative approach to table size reduction is to keep each species but reduce the level of precision in the table. However, this makes it more difficult to interpolate non-linear quantities, such as reaction rates, and hence the combustion may not be captured accurately. It should be noted that these limitations extend to ISAT, which may create large tables, and not just manifold methods.

2.5.4 CSE Implementation

The CSE routine is performed only when there is a sufficient number of ‘active’ CFD cells in a pre-defined CSE ensemble. Here, a CFD cell is considered to be ‘active’ if $\tilde{Z} \geq 0.0015$; this is to avoid extra computation in areas with minimal combustion. The number of active cells in a given CSE ensemble must exceed the number of discretized ζ bins in order to apply Equation 2.50 and perform CSE. The CSE routine is called once and applied to all active CSE ensembles simultaneously in a given time step, rather than being called sequentially.

In CSE, the computational domain is divided into a number of CSE ensembles composed of CFD cells, and the integral inversions are performed in each CSE ensemble. There must be a sufficient number of CFD cells in each CSE ensemble to accurately represent the mixing field. However, different features of the flame will not be captured well if the ensembles are too large. In general, CSE ensembles contain hundreds or thousands of CFD cells, and 20-30 CSE ensembles are defined for a given flame. The CSE ensemble size can be made smaller to capture non-linear effects, such as ignition, however the ensembles must contain enough CFD cells to reduce numerical errors from integral inversion. Thus, the CSE ensemble size must balance between inversion accuracy and chemistry precision.

The CSE ensembles can be defined in various ways. It is possible to define the ensembles using various algorithms [72, 73], in which the CSE ensembles are dynamic and may contain different CFD cells based on the value of a given parameter, for example ζ . The resulting CSE ensemble distribution can resemble the flame structure and be discontinuous. However, it is simpler to define static CSE ensembles that contain a fixed number of CFD cells. This results in a geometric distribution of CSE ensembles. In the current study, the CSE ensembles are defined as planes in the axial direction, making use of the observation that radial fluctuations of conditional averages are low [65]. This ensemble arrangement has worked successfully in previous CSE simulations of axisymmetric flames [27, 28, 30, 35]. In addition, the CSE ensembles may be composed of either equal or uneven sizes, and may contain overlap with other CSE ensembles. The purpose of overlapping CSE ensembles is to reduce possible sharp changes in the conditional averages near the ensemble boundaries. Static CSE ensemble distributions are often designed using a grid convergence study.

The CSE combustion model is coupled with the OpenFOAM v7 [74] CFD solver. Governing equations for mass, linear momentum and energy are solved using OpenFOAM, along with transport equations for the Favre averaged CO_2 , H_2O , methane (CH_4) and oxygen (O_2) mass fractions. Additional transport equations for $\tilde{Z}$ and $\widetilde{Z''^2}$ are defined and solved using OpenFOAM. The β -PDF is constructed at each CFD cell using $\tilde{Z}$ and




Figure 2.4: Flowchart of CSE implementation with reduced chemistry in the CFD Module

$\bar{Z}''^2$. When the CSE routine is activated, first the conditional $\bar{Y}_{CO_2}|\zeta$ and $\bar{Y}_{H_2O}|\zeta$ values are determined from integral inversion using Equation 2.50. The conditional values of the other reactive scalars are then interpolated from the pre-tabulated TGLDM tables based on the values of ζ , $\bar{Y}_{CO_2}|\zeta$ and $\bar{Y}_{H_2O}|\zeta$, as described in Section 2.5.3. Finally, the Favre averages of the other reactive scalars are obtained by solving Equation 2.47 above, and are exported back to the CFD solver. Note that the remaining species mass fractions are obtained from the tables and integrated, rather than solving transport equations. This follows the approach of Wang et al. [25], where the mass fractions obtained from tabulation were shown to be reasonable to those obtained from transport equations.

The entire process is summarized in Figure 2.4. The focus of this work is to replace the chemistry representation in CSE with ANN calculations rather than searching and interpolating the pre-tabulated TGLDM solutions; this is highlighted in Figure 2.4.

2.6 Summary

In this chapter, the information required to provide a solution for turbulent reacting flows with CSE is presented. The governing equations for fluid flow are presented. The phe-

nomenon of turbulence is described, and the modelling frameworks for turbulence (DNS, LES and RANS) are discussed with an emphasis on RANS, as it is used in the current study and in many previous works. Averaging methods for representing turbulent flow with the governing equations are also presented. These processes result in unclosed turbulence and chemical terms, which must be modelled in the RANS framework. An overview of commonly used turbulence and combustion models in the context of non-premixed combustion is presented. The Conditional Source-term Estimation combustion model for non-premixed combustion is also presented here, with a focus on the chemistry representation. The following chapter presents ML models and an overview of the current state-of-the-art of ML chemistry representation for turbulent combustion.

Chapter 3

Machine Learning Methods for Turbulent Combustion

This chapter contains an introduction and review of previous ML applications to turbulent combustion chemistry. A description of two commonly-used ML models is provided first, followed by a review of previous ML studies involving various sources of training data. The review is mainly limited to non-premixed flames, however there have been many applications of ML for different turbulent combustion problems [17, 18].

3.1 Machine Learning Models

This section describes two commonly-used machine learning models for chemistry tabulation: artificial neural networks and decision trees. The related sub-models are also presented. Note that both models are considered ‘supervised learning’ approaches, in which the ML model aims to map labelled input variables to output variables. This is in contrast to ‘unsupervised learning’, in which the data are not explicitly labelled and the ML model often aims to generate outputs.

3.1.1 Decision Trees

The decision tree (DT) is a ML model that mimics the structure of a tree diagram or flowchart, where outputs are determined from the input information by continuously splitting the data based on sets of criteria. Each node or ‘leaf’ represents a conditional ar-

gument, while each connection or ‘branch’ represents an outcome from the conditional argument. Each leaf has a sole branch to a leaf in the previous layer, and may have two or more branches to leaves in the subsequent layer. There is a sole branch between any two leaves in the DT. In addition, the branches only split and do not converge at the leaves.

DTs can be used for classification or regression tasks. The purpose of classification DTs is to determine a set outcome based on information about the data. The leaves represent specific labels, and the branches represent the possible values of those labels. The output of a classification DT is a specific category in which the input data belongs. Although DTs are commonly used for classification, it is more useful to consider regression DTs in the context of chemistry representation. The purpose of a regression DT is to predict outputs from input data based on prior knowledge. In this context, data is passed to the branches based on the values of the inputs at a given leaf. The conditional arguments in which the leaves represent are determined from a training process. The outputs of a regression DT are mapped to specific variables and may take continuous values.

Various extensions of the DT model exist. The random forest (RF) model is essentially an ensemble of DTs, in which similar DTs are trained on subsets of the training data and the results are aggregated. For classification, the output is taken as the most selected outcome from the individual DTs. For regression, the output is taken as the average of the individual DT outputs. The training data is sorted to each DT with replacement, meaning that duplicate data can appear in each DT. The RF model was introduced to reduce overfitting in base DTs via aggregation; this results in lower variance, but increased bias in the outputs. RF can provide better accuracy than DT but the resulting structure is more complicated, resulting in larger computational cost and less interpretability. Another extension is the gradient boosted tree (GBT), which aims to construct a strong DT as an ensemble of weaker DTs. At each step, a weak DT is trained using the training error from the strong model rather than the training dataset itself. The resulting weak model output is then added to the strong model to improve predictions. Various algorithms exist, including XGBoost [75]. GBT models aim to reduce overfitting by minimizing the bias (and hence increasing the variance). They have been shown to outperform RF [76, 77], but GBT has the same disadvantages as RF stemming from the complexity of the model.

The main benefit of the DT model is the ease of use. DTs are analogous to graphical tree models and flowcharts, which are commonly used for planning and education. DTs are considered a ‘white-box’ model since the explanation for a given output can be easily deduced by analyzing the logic in the internal structure (i.e., constructing a path and analyzing the node and connection values/operations in the tree). This also allows for easy modification of the DT. In addition, DTs require little data preparation and can be used with large datasets. The main disadvantages of DTs, however, are the DT stiffness

and the tendency to overfit to the training data. Overfitting occurs when the output data follows too closely to the training data; this may occur with any ML model. Overfitting results in more sensitivity to noise and difficulties in extrapolating to unseen data. The DT structure is also sensitive to the training data; small changes result in vastly different trees due to different splitting criteria. The size of the DT increases significantly with many input variables, therefore DTs are more suited to cases with a small number of input variables to limit the number of nodes and connections. It will be shown in Section 3.2.5 that DT and related models are easier to train and faster to run than ANNs, however DTs require more memory than ANNs while the relative accuracy to ANNs depends on the application. It is desired to have a ML model that is compact and can generalize to unseen data; thus, DTs are not selected in the current study.

3.1.2 Artificial Neural Networks

Artificial neural networks (ANN) are a class of ML models that mimic the structure and function of the human brain. For ANNs, outputs are determined by passing input information through multiple connections and transforming the values at each node. ANNs can also be used for classification or regression problems. For classification, ANNs aim to identify patterns in the dataset and assign the outputs to a given category based on the inputs. For regression, ANNs aim to construct a function relating the output variables to the input variables. There are also multiple types of ANNs, which may be characterized by how the nodes are organized and connected, and how the nodes transform data. The specific model used in this study is the multi-layer perceptron (MLP) for regression. The MLP is one of the simplest ANN models and has been used successfully in previous ML applications for combustion chemistry [17, 18].

ANNs consist of ‘neurons’ (or nodes), which transform data (i.e., signals) based on a prescribed function, and ‘connections’, which allow for data to pass through the network. ‘Weights’ and ‘biases’ are assigned to each connection and neuron, respectively, to adjust the strength of a given signal. Each weight and bias is unique and the values are adjusted during the training process. In addition, each neuron has an ‘activation function’, which transforms the output of that neuron based on a prescribed function. Neurons are typically organized into ‘layers’, in which each layer performs a specific transformation on the data. ANNs consist of one input layer and one output layer, and may contain one or more ‘hidden’ layers between the input and output layers. The neurons in the input and output layers correspond to variables in the system under study. The transfer of data is ‘feedforward’ in an MLP, meaning that data travels from the input layer to the output layer while passing through the hidden layers. The neurons are ‘fully connected’ in an MLP, meaning that the

value of a neuron in a given layer depends on the values from all neurons in the previous layer, but not the other neurons in the same layer. The output of a given neuron y is as follows, where f is the activation function of that neuron, b is the bias of that neuron, x_i is an output of a neuron i in the previous layer, w_i is the weight for the corresponding connection, and n is the number of neurons in a given layer:

$$y = f \left(b + \sum_{i=1}^n (x_i w_i) \right). \quad (3.1)$$

ANNs are selected to represent turbulent combustion chemistry in the current study. The main advantage of ANNs is the ability to accurately represent non-linear data in a compact manner. Since ANNs involve the optimization of a series of functions, the ANNs can represent non-linear relationships by tuning these functions to the dataset. ANNs can also generalize easily since new quantities can be computed easily with the ANN functions. The compact structure makes ANNs easier to use and store versus other ML models. These characteristics are useful for turbulent combustion chemistry, which is highly non-linear and requires an intractable amount of storage space for tabulation. However, the main disadvantage of ANNs is the complexity and the amount of tuning needed. A sufficient number of training samples is needed to accurately represent the dataset. The samples must also be representative of the underlying system, otherwise the ANNs will be tuned to incorrect data. Data transformation is often required to aid the training process. In addition, the ANN architecture and hyperparameter selection is non-trivial and must be ‘optimized’ to the dataset; different settings will significantly impact the prediction accuracy (i.e., differences between ANN predictions and the outputs from the training dataset). The ANN training process is significantly longer than that of the DT due to the additional amount of training data and parameters. Finally, ANNs are considered a ‘black-box’ model, in which the internal relationships between nodes are complicated and often unknown. This makes ANNs less intuitive and more difficult to debug. Despite these shortcomings, ANNs are commonly used for ML applications [17, 18]. Chapter 4 contains a detailed presentation of the ANN development for the cases under study.

3.2 Literature Review

The following section contains a literature review of previous ML applications for turbulent combustion chemistry representation. These studies are organized primarily on the source of training data or the role of the study in the literature.

3.2.1 Early ML Studies

Christo et al. [78] were one of the first groups to investigate ANN chemistry representation, where one ANN is applied to represent a three-step $\text{H}_2\text{-CO}_2$ mechanism. They use the mixture fraction and major species molar fractions to predict changes in major species molar fractions. Following trial-and-error, the ANN consists of two hidden layers with even neurons, uses the hyperbolic tangent ('tanh') activation function, and is trained using randomized samples from a representative chemistry dataset. The learning rate is adaptive, while the momentum term is constant. The optimization of the ANN structure is found to heavily rely on the training samples. The ANN is applied to a transported PDF simulation of a turbulent jet $\text{H}_2\text{-CO}_2$ flame and compared to simulations with direct integration (DI) and a look-up table (LUT). The species and temperature from the ANN simulation are found to be accurate to the other two approaches. For this case, the CPU time was found to be similar between ANN and DI, but not as quick as the LUT approach. The memory was found to be lowest with DI and largest with LUT, with the ANN approach in the middle. The benefits of the ANN approach were more clear with larger mechanisms.

The work is extended in [79], where they compare the use of ANN and repro-modelling (RM) for representing the same flame. Both methods are found to accurately represent the chemical source term while requiring less memory than the look-up table approach and less time than direct integration. The ANN approach had five times greater computational cost than RM, but was more accurate by one order of magnitude and had better generalizability. Both models experienced issues when data outside of the model range was used. The work was extended again in [80] to use training samples from small-scale PDF Monte-Carlo simulations. The same ANN architecture is used, however it is applied to an arbitrary one-step chemical mechanism. The sampling method is shown to be feasible for the ANN training and the PDF simulation.

Blasco et al. [81] use three ANNs to predict species evolution at fixed time steps, and one ANN for temperature and density, in the context of methane-air combustion. A 4-step mechanism with seven species is considered. 4000 samples are generated from an ODE solver with CHEMKIN used for thermodynamic properties; 50% of the samples are used for training. The data is standardized using Z-score normalization (zero mean, unity variance) and mapped to a (-1,1) range. Training is completed using a second-order scaled conjugate gradient method instead of backpropagation, which is the common approach. Following trial-and-error, the architecture is either one or two hidden layers with 20 neurons each, using the sigmoid activation function. The ANN prediction errors are found to be on the order of 1-10% for each quantity. The ANN method is compared offline to DI and LUT. The ANN RAM requirement is somewhat lower than DI and much lower than LUT, while

the ANN CPU time is much lower than DI but slightly larger than LUT. The work was extended in [82] to account for arbitrary time steps. This work introduces partitioning, where ANNs are assigned for either set ζ values, set ζ and molar CH_4 values, or set ζ intervals. The ANN structure is retained for the interval case, however the time step is an additional input. 80% of the 2500 samples are used for training. The ANNs are validated using a plug-flow reactor. It is observed that ζ intervals provide better accuracy while being less restrictive than the other two partitioning methods. The accuracy, speed and memory improved relative to the previous paper. The work was extended further in [83] by introducing the self-organizing map (SOM) to cluster points based on chemical similarity, and assigning an ANN to each subdomain. This avoids the need for assuming a characteristic variable for partitioning, such as ζ . The SOM uses initial composition and time step for classification. Temperature is now calculated with the species evolution in the ANN. Various topologies are tested and applied to a partially-stirred reactor (PaSR) simulation. It is observed that relatively simple ANNs (single layer, 5-10 neurons) with moderate partitions (100) provide the most accurate results. Mean absolute errors for the species were reported in the range of 10^{-7} to 10^{-5} . The accuracy is on a similar order to the previous work, however the memory has increased slightly due to the additional networks.

A similar work by Chen et al. [84] considers fitting ISAT results to ANNs rather than using SOM. Two mechanisms (4-step and 9 scalars, 12-step and 17 scalars) are considered. Each ANN consists of one hidden layer with a matching number of neurons to the number of scalars. Training and testing data is split into 8000 to 4000 samples, respectively. The approach is also evaluated using a methane-air PaSR. It is observed that the more detailed ANN (12-step) has higher accuracy relative to ISAT, with average errors up to 0.5%, whereas the less detailed ANN (4-step) has average errors up to 3.5%. Significant memory savings are achieved with both sets of ANNs relative to ISAT, however the memory requirements are much larger than those in the previous papers. The ANN CPU time is similar to the ISAT CPU time.

The details of these studies are summarized in Table 3.1. In these studies, simple fuels with significantly reduced mechanisms are used. The ANNs often have a small structure and a small amount of training data. The ANNs are evaluated against a similar case to the training data. ANN predictions of thermochemical quantities are often accurate, with improvements being made by partitioning the solution space in some manner. Partitioning via ζ intervals [82] in particular will be useful for combining ANNs with CSE. Computational performance gains depend on the size of the mechanism. In these studies, ANNs do not always outperform DI or LUT in terms of memory or time.

Table 3.1: Summary of early ML studies for turbulent combustion chemistry

Authors	Year	Flame	Fuel	Data Source	Mechanism	Framework
Christo et al. [79]	1995	Jet flame	H ₂ -CO ₂	DI calculations	3-step	Stochastic PDF
Christo et al. [78]	1996	Piloted jet	H ₂ -CO ₂	DI calculations	3-step	Stochastic PDF
Christo et al. [80]	1996	Piloted jet	Arbitrary	PDF samples	1-step	Stochastic PDF
Blasco et al. [81]	1998	N/A	Methane	DI samples	4-step	DI
Blasco et al. [82]	1999	Plug-flow reactor	Methane	DI samples	4-step	Plug-flow calculations
Chen et al. [84]	2000	PaSR	Methane	ISAT	4- and 12-step	Stochastic PDF
Blasco et al. [83]	2000	PaSR	Methane	PaSR samples	GRI 2.11 (reduced to 5-step)	Stochastic PDF

3.2.2 ML Chemistry with Flamelet Training Data

Many applications of ML chemistry involve the use of flamelet solutions for training data. These studies are reviewed in this section. Note that the three previous studies of ML chemistry for Sandia flame D involve either traditional or augmented flamelet training data and are applied to LES.

The work by Kempf et al. [85] was one of the first applications of non-premixed flamelet data (NFD) for ANN training. One ANN each is prescribed for density, viscosity and each species mass fraction. Inputs are mean and variance of the mixture fraction, and the mean scalar dissipation rate. Each ANN consists of two hidden layers. The ANNs are applied to an LES of Sandia flame D. The experimental trends were captured, however the mass fractions of minor species were overpredicted in the fuel-rich region. The storage size relative to NFD reduced by a factor of 3, while the computational time was similar. The interpolations were observed to be smoother with ANN than NFD.

Chatzopoulos and Rigopoulos [86] apply ANNs with SOM to non-premixed methane-hydrogen flames DLR-A and DLR-B. A 415-step mechanism is reduced from 63 species to

17 species using rate-controlled constrained equilibrium (RCCE), which provides a system of differential-algebraic equations to obtain the reduced mechanism. The ANNs replace the integration in the RCCE model by determining composition at the next time step from the current composition and time. An SOM is used to partition the composition space into 400 subdomains, with one ANN per subdomain. The training data is obtained from non-premixed flamelet solutions with two artificial pilots introduced to represent ignition. Following trial-and-error, the ANN structure is (18,18), denoting two hidden layers of 18 neurons each. The ANN approach is then applied to RANS-PDF simulations of the two flames and compared with the RCCE approach. Both formulations are able to capture the experimental trends of each flame. The ANNs generally align with the RCCE solutions, however there are deviations in the minor species. The computational time is reduced by a factor of about 500 relative to RCCE. The methodology is extended by Franke et al. [87] and applied to an LES-PDF simulation of Sydney flame L, which is a non-premixed methane flame with local extinction and ignition. The mechanism is GRI-Mech 1.2 reduced to 16 species. Following testing, the same SOM topology as the previous study is used. Both ANN hidden layers now contain 30 neurons each, and the tanh activation function is used. 60% of the data is used for training, with the remainder for cross-validation. An additional 1.2 million samples are used for testing, with an average root mean squared error (RMSE) on the order of 10^{-7} for the species. Both formulations of LES are found to capture most of the experimental trends for temperature, mixture fraction and species. The ANN and RCCE approaches are often in agreement, especially towards the centerline. The CPU time was reduced by two orders of magnitude relative to RCCE.

Readshaw et al. [88] introduce a data generation procedure based on NFD to generalize the ANN approach. In this procedure, data is randomly augmented from flamelet solutions using molar ratios and the mixture fraction as constraints. The full GRI-Mech 1.2 mechanism with 31 species is used. The flamelet data are combined, resulting in 500,000 samples for training. Each ANN predicts a change in a single species from the current composition of all species. The ANNs contain a single hidden layer of 30 neurons and are trained using the Levenberg-Marquardt (LM) algorithm, which is a second-order algorithm. The tanh activation function is used. Various architectures (including with and without SOM) are tested using data extracted from LES-PDF of Sydney flame L. It is shown that the new approach provides the lowest average RMSE of 0.059%, which is about a factor of 3.5 lower than the RMSE of the architecture from [87]. The increase in computational cost is deemed to be acceptable. After a validation with laminar flamelets, the ANN approach is then used to simulate Sydney flame L. The ANN and DI solutions are in agreement for most quantities across various locations, with some deviations in the minor species. The experimental trends are generally captured by both formulations. The results are generally

improved from those in [87]. The computational speed is reported to increase by a factor of 18 with the ANN approach, while the relative cost of the reaction in LES-PDF reduced from 71% to 12% with ANN instead of DI.

This methodology is extended by Ding et al. [89] by introducing sequential ANNs for a given species. Data is passed to subsequent ANNs if predictions within a set range have large residual errors. This results in 3-5 ANNs for each species, or 97 ANNs with a storage requirement of 11 MB. The multiple ANN methodology is shown to reduce training errors by one order of magnitude. Each ANN consists of two hidden layers with 40 neurons each and the tanh activation function. Training is completed with 200,000 sets of data (out of 960,000) normalized to (-1,1). The LM algorithm is used and training is completed for 1500 epochs. In this work, the flamelet data is discarded and only the randomly generated data is used. The new methodology is validated with 1D laminar flamelets and premixed flames (different from the non-premixed training data), and is applied to LES-PDF of Sandia flames D, E and F. The ANN approach is able to capture the species mass fractions in each flame across all locations, with more prominent deviations from DI further downstream. The experimental trends are captured for each species in each flame. The ANN approach results in a speed increase of a factor of 12, whereas the relative cost of the reaction is reduced from 74% to 20%. The methodology is extended to non-premixed methane-hydrogen fuel blends in [90] and applied to LES-PDF of Sandia flame D (diluted methane) and the Sydney bluff-body swirl flame SMH1 (methane-hydrogen blend). The ANN structure and training settings are mostly identical to [89]. There are 2-3 ANNs per species, resulting in 71 total ANNs and a storage requirement of 5.5 MB. Training is completed with 100,000 samples (out of 950,000) and is ran for 1500 epochs. The ANNs are validated with the same flames as [89] and the absolute and relative errors are found to be similar to the previous case, being on the order of 10^{-5} to 10^{-2} . Additional validation cases are performed on 1D flames with fuel blends, and are found to have reasonable agreement. Species and temperature in Sandia flame D are predicted well by the ANN approach and follow experimental trends. Deviations in the minor species are observed downstream. The computational speed is reported to increase by a factor of 17 with ANNs, while the fraction of time devoted to the reaction is reduced from 78% to 14%. Similar trends are observed in Sydney flame SMH1, although there are more prominent deviations in the minor species near the base of the flame. The speed-up ratio with ANNs is reported to be 14, whereas the fraction of reaction time in the simulation is reduced from 84% to 27%. Readshaw et al. [91] extended the work to premixed methane flames. The ANN procedure is similar to the previous works, but considers premixed flamelet data (PFD) rather than NFD. Three ANNs are used for each species, for a total of 90 ANNs and 1.4 MB of storage. Contrary to the previous works, the outputs ranges for the subsequent ANNs

are determined automatically instead of manually, and data is passed from the ANN with the finest data range to the ANN with the largest data range (i.e., the reverse process of [89, 90]). After a validation with 1D premixed flames, the ANN approach is used to simulate two cases of the Cambridge/Sandia swirl burner with LES-PDF. Excellent agreement is achieved between ANN and DI for most quantities in both flames. The experimental trends are generally captured by both formulations. The speed-up with ANNs is reported to be a factor of 14, while the CPU time is reduced by about a factor of 3.5 with ANNs instead of DI. The RMSE of the species in the flames is reported to be on the order of 10^{-9} to 10^{-5} . Liu et al. [92] extend the work to non-premixed combustion of dimethyl ether (DME), using a 290-step mechanism with 55 species. The ANN implementation is retained, however the number of hidden layer neurons is increased to 50 and training uses three instances of 100,000 samples for 500 epochs, or 300,000 samples (out of 830,000 total samples) and 1500 epochs total. Validation with 1D laminar flames results in RMSE on the order of 10^{-4} to 10^{-1} for each species. The ANNs are then applied to the Sandia DME series of flames, resulting in excellent agreement with the DI solutions of most quantities in both flames. The speed-up ratio with ANNs is reported to 16, whereas the fraction of CPU time associated with the reaction is reduced from 96% to 59% with ANNs.

Bhalla et al. [93] focus on developing a training procedure for ANNs with NFD. The ANNs predict one of composition, temperature, heat release rate and chemical source term from pressure, mixture fraction and progress variable. Each ANN has 6-7 hidden layers with increasing neurons. The composition ANN features separate fully-connected networks in the final hidden layer for each species. The leaky rectified linear-unit (ReLU) activation function is used for each ANN. The inputs are normalized with Z-score and augmented using a function generator, which outputs a trigonometric or exponential function. The dataset consists of non-premixed flamelet solutions of hydrogen combustion at various pressures. Training is completed using the Adam solver with a fixed learning rate. 80% of the dataset is used for training, however ‘hard’ samples (with larger training errors) are shuffled into the training dataset such that 75% of the training dataset consists of samples with larger errors. The loss function is also weighted based on ‘hard’ samples. The standard error range is reported to be on the order of 10^{-5} to 10^{-3} for the species and 10^1 for temperature. 15-60% of these points are reported to be within the standard error range for composition and temperature. The R^2 value for temperature relative to NFD is reported to be 0.94. The ANN is shown to capture the trends in temperature at pressures excluded from the training set, however there are notable deviations in the high-pressure curve. The memory is reduced by a factor of 8 compared to LUT. A piloted flame simulation is shown briefly, with the temperature predictions being similar between ANN and LUT.

Table 3.2 summarizes the details of these studies. The typical ANN structure has evolved to contain more hidden layers and neurons, and has generally improved in both accuracy and performance relative to DI and LUT. More complex fuels have been represented with ML models using this approach. Many of the works in this area involve the generalization of the training data by using non-premixed flamelet solutions as a starting point, or the correlation of variables to improve training performance.

3.2.3 ML Chemistry with Sampled Training Data

ML models have also been trained using sampled chemistry data from various sources, mainly DNS or micro-mixing solutions. These studies are summarized in Table 3.3. Many of the same trends in ML chemistry representation are present here. These works introduce different methods for clustering data and adjusting the loss function to improve the performance of ML methods.

Sen and Menon [94] were one of the first groups to consider the use of independent problems for generating ANN data. Their first work considers samples from DNS of unsteady laminar flames with varying amounts of flame-vortex interactions, in which the ANNs are used to predict source terms from temperature and composition. The training process incorporates individual learning rate and momentum terms for the connections. The tanh activation function is used. Various ANN architectures are tested and applied to LES of turbulent premixed syngas-air flames with a 10-step mechanism. ANN structures of (10,5) and (10,5,4) are found to provide the most accurate results with similar and modest computational time, with the (10,5,4) structure following the DI results more closely.

Wan et al. [95] consider stochastic micro-mixing with heat loss effects as a general database for ANN training. Two separate ANNs are used for data with large and small burning rates (defined by $Y_{O_2} > 0.9$). Each ANN predicts all chemical source terms from the temperature and composition. Each ANN has a structure of (512,256,128,64,12) and uses the ReLU activation function. Separate ANNs for each output were tested but found to be less efficient. 17.5 million samples are generated, with 80% used for training. Adam is used as the training optimizer, with training taking place over 500 epochs with a decaying learn rate. The test dataset MSE is reported to be on the order of 10^{-7} . The ANN approach is evaluated in a DNS of a non-premixed syngas flame with side wall effects, and is compared to 2D and 3D DNS of the same flame with full and reduced chemical mechanisms. The ANNs are able to capture the scalar field and mass fractions against the full and reduced mechanisms in both simulations. More prominent deviations are observed in the centerline profiles than the radial profiles, especially for temperature and minor species. For the 2D

Table 3.2: Summary of ML turbulent combustion chemistry studies involving flamelet training data

Authors	Year	Flame	Fuel	Data Source	Mechanism	Framework
Kempf et al. [85]	2005	Sandia flame D	Methane	Steady NFD	GRI-Mech 3.0	LES
Chatzopoulos and Rigopoulos [86]	2013	DLR A and B	CH ₄ -H ₂ -N ₂	Unsteady NFD with artificial pilot	415-step (reduced to 17 species)	RANS-PDF
Franke et al. [87]	2017	Sydney flame L	Methane	Unsteady NFD with artificial pilot	GRI-Mech 1.2 (reduced to 16 species)	LES-PDF
Bhalla et al. [93]	2019	Piloted jet	Hydrogen	Flamelet	19-step	RANS
Readshaw et al. [88]	2021	Sydney flame L	Methane	Augmented NFD	GRI-Mech 1.2	LES-PDF
Ding et al. [89]	2021	Sandia flames D, E and F	Methane	Randomized NFD	GRI-Mech 1.2	LES-PDF
Ding et al. [90]	2022	Sandia flame D, Sydney SMH1	Methane and CH ₄ -H ₂	Randomized NFD	GRI-Mech 1.2	LES-PDF
Readshaw et al. [91]	2023	Cambridge Sandia swirl	Methane	Augmented PFD	GRI-Mech 1.2	LES-PDF
Liu et al. [92]	2024	Sandia DME flames D and F	DME	Randomized NFD	290-step	LES-PDF

case, the wall-clock time is reduced in the ANN approach by factors of 25 and 3 relative to the full and reduced mechanisms, respectively. The work is extended by Nguyen et al. [96] by introducing the ‘k-means’ algorithm for decomposing the thermochemical space into clusters, then applying principal component analysis (PCA) to create uncorrelated ANN input variables in the PC space. The ANN structure is retained from [95], however each ANN is assigned to a cluster instead of the reaction rates. The training process is also similar, however the number of total samples is reduced to about 7 million, and the number of epochs is reduced to 200. The ANN approach is evaluated using micro-mixing calculations involving the same flame conditions in [95] without the wall, in which excellent accuracy is achieved for temperature and species relative to the reduced mechanism. It is shown that pre-treatment of the data (via k-means and PCA) is necessary to capture the predictions after ignition. The speed-up of ANNs in the micro-mixing calculations is a factor of 10 relative to the reduced mechanism.

Chi et al. [97] consider a physics-based loss function for ANN training of premixed combustion data. Separate cases are used for 1D FGM with unity Lewis number, and 2D DNS samples with mixture-averaged diffusion. Three ANNs are used per case, with each being assigned for reaction rates, mass fractions and temperature, respectively. The FGM case considers progress variable and nitrogen mass fraction as inputs, whereas the DNS case considers progress variable and elemental mass fractions as inputs. The ANNs have 2-4 hidden layers with 12-200 neurons each, with the DNS set of ANNs being larger than the FGM set. The tanh function and standard min-max normalization are used. The new loss function is based on MSE and accounts for mass and atom conservation. Data in the flame front is weighted to account for sharp gradients. Training is completed using Stochastic Gradient Descent with momentum (SGDM) with 90% of the samples and a constant learning rate. Good training accuracy is achieved with both sets of ANNs, with absolute deviations on the order of 10^{-1} to 10^0 for temperature and species. It is shown that the physics-based loss function leads to lower training losses than the traditional MSE loss function. The FGM-ANN set is evaluated on a 2D DNS of a premixed planar flame. Excellent accuracy is achieved relative to full chemistry, with a minimum R^2 of 0.98 and absolute errors less than 1% for the species. The DNS-ANN set is evaluated on a 3D DNS of an unsteady jet flame with differential diffusion. Predictions across the flame front are quite good, with some deviations in temperature and minor species relative to both full chemistry and LUT. The R^2 for temperature is reported to be 0.976. Both sets of ANNs are found to provide significant memory savings and are 4 to 6 times faster than full chemistry. It is shown that the ANN retrieval and CPU times are slightly slower than LUT for both cases, however the chemistry is reported to be 6% of the computational cost in the 3D case.

Table 3.3: Summary of previous ANN thermochemistry studies involving samples for training data

Authors	Year	Flame	Fuel	Data Source	Mechanism	Framework
Sen and Menon [94]	2009	Premixed	Syngas	DNS samples	10-step	LES
Wan et al. [95]	2020	Wall-bounded flame	Syngas	Micro-mixing samples	GRI-Mech 3.0 (reduced to 23 steps)	DNS
Nguyen et al. [96]	2021	Wall-bounded flame	Syngas	Micro-mixing samples	GRI-Mech 3.0 (reduced to 23 steps)	Stochastic micro-mixing
Chi et al. [97]	2022	Premixed planar and jet flames	Methane	FGM and DNS samples	GRI-Mech 3.0 (Lu13)	DNS

3.2.4 ML Chemistry with Alternative Source of Tabulated Training Data

ML models for turbulent combustion chemistry have also been trained using sources of tabulated chemistry. These studies mainly use other flamelet models such as the Flamelet/Progress Variable (FPV) approach or FGM, as well as the Linear Eddy Model (LEM) approach. The details of these studies are summarized in Table 3.4, with similar trends to the previous ML studies.

Ihme et al. [98] apply a custom optimization algorithm for ANN design towards an LES of the Sydney flame SMH1, which burns a fuel blend of methane and hydrogen. Flamelet/progress variable (FPV) data is considered for ANN training, where FPV replaces the scalar dissipation rate with a progress variable defined by the mass fractions of characteristic species. Each ANN predicts one species mass fraction (CO_2 or H_2O) and one fluid property, and uses the sigmoid activation function. The optimization consists of a generalized pattern search algorithm to determine the ANN structure with the lowest errors. This resulted in networks of 3-4 hidden layers with 3-8 neurons each. For the LES, it is observed that the ANN approach provides a similar level of accuracy to FPV tables with moderate discretization, while requiring a fraction of the memory. The data retrieval time increased by a factor of 5 with the ANN approach, however the increase was justified

since the chemistry retrieval amounted to 3% of the LES computational demand.

Sen and Menon extended their previous work (see Section 3.2.3) to use the linear eddy model (LEM) for data generation of both premixed and non-premixed flames [99]. LEM solves a one-dimensional unsteady reaction-diffusion equation along the line of the strongest scalar gradient, which is across the flame front in the context of ANN training. Five million points are generated, 85% of which are used for training, and data are normalized to at least (-0.8,0.8). The ANNs are separated into temperature bins and predict all source terms from temperature and composition. Various premixed and non-premixed cases with either methane or syngas combustion are presented. The LEM-based ANN performance is found to be similar to the previous work for the premixed syngas flame. The following considers the non-premixed syngas case with LEM testing. The optimal ANN architecture (from trial-and-error) is reported to be (5,3,2), providing low errors with reasonable training time. The species mass fractions are found to have excellent agreement between DI and ANN. The CPU time decreased by a factor of up to 5 for the non-premixed case versus 11 for the premixed case, due to the use of a stiffer mechanism for the premixed case. In [100], the non-premixed case is simulated with LES, where ANNs replace LEM as the sub-grid combustion model. Results from ANN-LES are compared to a DNS of the same flame. It is found that ANN-LES is able to capture the extinction and reignition events in the flame, while being 5.5 times as quick as the DNS. Slight differences in the conditional profiles of mass fractions and the PDF of the scalar dissipation rate are observed.

Sinaei and Tabejamaat [101] use ANNs in an LES of a non-premixed planar diluted methane jet flame, with a focus on the velocity field. The ANNs predict reaction rates from composition and time. For ANN training, 4 million points are generated from LEM calculations with various Reynolds numbers. A skeletal (25-step) and reduced (2-step) mechanism are considered. The optimal ANN architecture (for the 25-step mechanism) is selected to be (10,8,6). The ANN approach is compared to LES with DI and LUT, as well as DNS results of the same flame. ANN-LES is able to capture the trends of temperature and species, with some discrepancies at the peaks. Small differences in the profiles are observed relative to both DI and LUT for both mechanisms. Species mass fraction errors are reported to be up to 5%. For the skeletal mechanism, the CPU time is reported to decrease by a factor of five relative to DI-LES, and the storage requirement is several orders of magnitude lower than DI and LUT. The reduced mechanism captures the trends but is less accurate than the skeletal mechanism.

Zhang et al. [102] trained a series of ANNs on FGM data and applied ANNs with less than 5% relative error to RANS and LES of n-heptane spray combustion. A 283-step mechanism with 68 species is used. Each ANN predicts a species mass fraction from mixture fraction and progress variable. The ANNs consist of five hidden layers with symmetric

neurons (from trial-and-error). The ReLU activation function is used. The Adam solver with scheduled learn rate decrease is used. The maximum number of epochs is set to 30,000. The ANNs are able to predict the species well, with a minimum Pearson coefficient ($=\sqrt{R^2}$) of about 0.9945. The ANN approach is first validated with RANS of Engine Combustion Network (ECH) spray flame H, and is found to capture the ignition delay and lift-off length relative to FGM at various temperatures. An LES of the same flame is performed with ANN, and it is found that the ignition delay time is captured while the lift-off length approaches the experimental value after some time. The ANN approach is reported to reduce memory by a factor of eight, although retrieval time is increased by a factor of three relative to FGM.

Owoyele et al. [103] introduce a framework to group ANN target variables by similarity. The validation case is RANS and LES of ECN spray flame A, which uses n-dodecane fuel and a mechanism consisting of 24 species. The ANNs are trained using data from an unsteady multi-flamelet model. The ANNs predict compositions from the mixture fraction, scalar dissipation rate and time. The tanh activation function is used. The LM algorithm is used for training. Various output configurations (single target, multi-target and grouped targets) are tested. The single target case is quite accurate ($R^2 \sim 0.999$) but is slower than the LUT approach, whereas the multi-target case is faster than LUT but has an R^2 of about 0.997 with scatter. The grouped target approach aims to assign one ANN to predict a number of species that are correlated to each other. In ECH spray flame A, the grouped target approach is shown to predict most species more accurately than the multi-target approach. The grouped target approach is also able to capture the ignition delay, lift-off length and heat release. The test case is LES of a compression-ignition (CI) engine with methyl decanoate, consisting of 3299 species. The ANNs are similar, however pressure is an additional ANN input. The ANN-LES is found to capture the experimental trends for pressure, heat release and lift-off length. A reduction in memory compared to LUT is observed for both cases. Speed increases of 8% and 37% are reported for the validation and test cases, respectively.

Hansinger et al. [104] consider a novel ML network for partially-premixed combustion. The network predicts composition, temperature and the progress variable source term from the mixture fraction, progress variable and the previous temperature. The network consists of blocks of traditional ANNs with two hidden layers each. Outputs of a given block are fed as inputs to the subsequent block. The ‘skip connection’ is also introduced, where the inputs are directly fed to both the first and second hidden layers in a given block. The second hidden layer is thus fully connected to the input and first hidden layers. The ReLU activation function is used. Training data consists of 95% of 2.5 million samples from a FPV table, and are normalized using both cubic root transformation and

Z-score normalization. Training is completed over 2046 epochs using the Adam optimizer. Following trial-and-error, a network consisting of 12 blocks of (300,300) ANNs is considered based on the MSE of 10^{-6} . The R^2 is reported to be 0.999 for each quantity, with relative errors up to 4%. Most of the prediction errors in temperature and progress variables source term are observed in high-gradient regions. Following a brief validation with a 2D laminar premixed methane flame, the methodology is evaluated against FPV for a 3D LES of the Sydney/Sandia piloted methane flame with inhomogeneous inlets. Results for the species, temperature, mixture fraction and velocity are similar between ANN and FPV, and generally follow the experimental trends at each location. The storage size is reported to reduce by a factor of 55 from 485 MB (FPV table) to 8.7 MB with the ANN.

Abdelwahid et al. [105] apply PCA to ANNs for high-pressure non-premixed ammonia/hydrogen combustion. Training data is obtained from 1D flamelets with a 236-step mechanism. PCA is then applied on the training dataset to obtain two PCs. The ANN predicts temperature, composition and chemical source terms using the PCs as inputs. Following testing, the ANN consists of 5 hidden layers with 500 neurons each. The ReLU activation function is used, and batch normalization is included per hidden layer. The loss function incorporates both MSE and R^2 . The ANN is validated on a 1D counterflow flame with unity Lewis number, and is found to replicate the DI results. The ANN approach is then implemented in LES of two non-premixed ammonia/hydrogen jet flames with different ammonia concentrations. The conditional averages of temperature and species are captured by the ANN in both flames relative to FPV. With the unity Lewis number assumption, the conditional predictions become less accurate downstream relative to the experimental data, especially for ammonia and hydrogen. The use of mixture-averaged diffusion in the ANN training results in better predictions of the conditional species, however the conditional temperature predictions are observed to be less accurate than the unity Lewis number ANN. The physical NO mass fraction is provided, and the trends are found to be captured with the ANN. The computational cost is reported to be reduced by 14% relative to FPV.

3.2.5 Studies Involving Comparisons of ML Methods

Certain studies have involved the comparison of two or more ML models for representing turbulent combustion chemistry. These studies often compare the use of ANNs and DTs, or similar models, for the same problem. Table 3.5 presents the details of these studies.

Mousemi et al. [106] compared a set of four ANNs and one decision tree (DT) in the prediction of four unclosed terms from the Uniform Conditional State (UCS) model, which

Table 3.4: Summary of previous ANN thermochemistry studies involving alternative sources of training data

Authors	Year	Flame	Fuel	Data Source	Mechanism	Framework
Ihme et al. [98]	2008	Sydney SMH1	CH ₄ -H ₂	Steady FPV	GRI-Mech 2.11	LES
Sen and Menon [99]	2010	Non-premixed plane jet	Syngas	LEM	21-step	LEM
Sen and Menon [100]	2010	Non-premixed plane jet	Syngas	LEM	21-step	LES
Sinaei and Tabejamaat [101]	2017	Planar jet	Methane	LEM	2- and 25-step	LES
Zhang et al. [102]	2020	ECH spray H	n-heptane	Unsteady FGM	283-step	RANS and LES
Owoyele et al. [103]	2020	ECH spray A and CI engine	n-dodecane and methyl decanoate	Unsteady Multi-flamelet	24 species and 3299 species	RANS and LES
Hansinger et al. [104]	2022	Sydney-Sandia piloted flame	Methane	FPV	GRI-Mech 3.0 (Lu13)	LES
Abdelwahid et al. [105]	2023	High-pressure jet	NH ₃ -H ₂	FPV	236-step	LES

Table 3.5: Summary of previous ML method comparison studies for thermochemistry

Authors	Year	Flame	Fuel	Data Source	Mechanism	Framework
Mousemi et al. [106]	2023	BRS pre-mixed burner	Methane	UCS calculations	GRI-Mech 3.0	Unsteady RANS
Li et al. [107]	2023	Counterflow flame	n-dodecane	Unsteady FGM	269-step	FGM

is based on the CMC model. The inputs are mean and segregation of mixture fraction and progress variable. 18.75 million samples are generated from UCS calculations, with 80% allocated for training. A single DT with all four outputs is used, whereas one ANN for each output variable is considered. The DT has a maximum depth of 25 and did not require scaling. Following optimization with a grid search method, each ANN has 4 hidden layers of 10-30 neurons and uses the ReLU activation function. The Adam solver is used for training, with the number of epochs set to 200, the learning rate set to 0.001, and the L2 regularization term set to 0.0001. For the test dataset, both techniques achieved good accuracy for most quantities (R^2 between 0.975 and 0.998), with the DT providing more accurate results. The MSE of each quantity is on the same order of magnitude between both methods, but tends to be lower with the DT. The DT took minutes to train, which is significantly lower than the ANN training time. Both ML methods are evaluated in an unsteady RANS (URANS) simulation of a premixed methane burner and compared to UCS. It is found that both DT and ANN capture the axial velocity and the general distribution of temperature and progress variable. The nitric oxide mass fraction (Y_{NO}) trends are captured with both DT and ANN, however DT tends to predict slightly lower Y_{NO} , whereas ANN tends to predict somewhat larger Y_{NO} . The DT simulation is found to be 10% quicker than URANS while saving 0.9 GB of memory, whereas the ANN simulation saves 1.9 GB of memory despite being 8% slower than URANS. It is concluded that the ANNs would have greater benefits over DT if the dataset was expanded, despite the DT having slight advantages in accuracy and speed versus the ANNs.

Li et al. [107] compared two separate ANNs with one RF and GBT for predicting the progress variable source term from the mixture fraction and the progress variable itself. 150,000 samples are generated from FGM, with 70% of this data allocated for training the DT-based models. Data are preprocessed using logarithm scaling. The optimal RF architecture is found to be 25 trees with a minimum leaf size of 2. The optimal GBT architecture is found to be 2200 trees with a minimum leaf size of 4 and a learning rate

of 0.5. Both the RF and GBT occupy about 2.56 MB each. A shallow and deep ANN are developed and compared. The shallow ANN consists of nine hidden layers with 30 neurons each and uses the tanh activation function. Training is conducted with the LM algorithm over 100 epochs. The storage requirement is reported to be 3.6 MB. The deep ANN consists of 13 hidden layers with exponentially decreasing neurons and uses the leaky ReLU activation function. Training is conducted with the Adam solver over 600 epochs, including L2 regularization and scheduled learn rate drop. From training, each model is shown to predict the progress variable source term well, with the Pearson R coefficient exceeding 0.998 in each case. Both ANN models are found to provide slightly better predictions than the ensemble models (RF and GBT), with the shallow ANN being more accurate. However, the shallow ANN took the most time to train (11 hours), whereas the ensemble models took minutes to train. For the validation case of a n-dodecane counterflow diffusion flame, each model captures the temporal evolution of temperature and species relative to FGM, with the ANN models generally being more accurate than the ensemble models. It is concluded that the shallow ANN performs the best with the given data pre-processing, despite the substantial increase in training time.

3.3 Summary

In this chapter, the current state-of-the-art of ML chemistry representation is examined for turbulent combustion. An overview of ANN and DT-based models is presented, followed by a literature review of previous studies on turbulent combustion with ML methods. With each of these studies, ANNs are often implemented to reduce computational costs related to chemical source term evaluation, either by avoiding the integration step in DI itself or alleviating the storage limitations of LUT. With this in mind, ANNs are designed to predict the chemical source term or changes in composition, however they can be used to predict other variables. The ANN approach has been evaluated in a variety of premixed and non-premixed flames, often against the LUT approach or DI in some cases. Excellent accuracy can be achieved for the flow field and compositions. Storage reductions are obtained with the ANN approach relative to both DI and LUT. The relative computational cost of the ANN approach depends on the complexity of the flame and mechanism. Speed-up is generally achieved relative to DI, while the computational effort is often on the same order as LUT. Other ML approaches have been compared with ANN in some studies. ANN is generally shown to provide reasonable accuracy with the most scalable memory requirement relative to DT and ensemble models, at the expense of significantly longer training time.

Much of the effort is focused on improving the ANN training accuracy. This is a difficult task due to the non-linearity of turbulent combustion and the difficulty to capture high-gradient regions. The ANN architecture and training settings are often determined from trial-and-error, and can influence the ANN training accuracy. Another significant factor is the quality of training data, which is often from a canonical combustion solution relevant to the problem at hand. Various efforts have been made to improve the quality of training data and the training process, using techniques such as data augmentation and scaling, applying ANNs to select quantities or to select thermochemical subdomains only, and by adjusting the loss function. Although efforts have been made to generalize the ANN approach and extend it to complex fuels, ML methods in turbulent combustion remain uncommon. Chapter 4 presents a proposed ANN framework for handling different fuels and sources of data.

Chapter 4

Machine Learning Methods for Chemistry Representation

This chapter presents and validates a methodology for chemistry representation using ML methods, specifically ANN. This methodology can be extended to any fuel, however the focus of the study is non-premixed methane-air combustion in the RANS framework. First, the development of the ANN architecture in the context of the CSE combustion model is presented in detail. The data preparation procedure is then described for data sourced from both tabulated chemistry and samples of conditional averages DI-CSE. The ANN training and testing process is then presented, followed by training results for the ANNs under study. The influence of ANN settings on the ANN performance is analyzed throughout the chapter. Finally, a brief description of the ANN implementation in CSE is provided.

A total of four sets of ANNs are developed for the works in this study. Four distinct sets of training data are used: TGLDM solutions for pure methane with GRI-Mech 3.0, TGLDM solutions for diluted methane with GRI-Mech 3.0, TGLDM solutions for diluted methane with Smooke’s mechanism, and DI-CSE samples for diluted methane with Smooke’s mechanism.

4.1 ANN Architecture

This section describes the specific ANN architecture used in the current study and how it is derived. The selection of the input and output variables, number of neurons and hidden layers, and type of activation function are discussed here.

The main motivation of ANN chemistry representation in this work is to replace the TGLDM tabulated chemistry representation in the CSE combustion model. In CSE, variables are discretized by the mixture fraction ζ for non-premixed combustion. The TGLDM data are organized following this structure, such that there is one table for each discretized ζ . Thus, for the ML approach, it is wise to construct ANNs for each discretized ζ . Considering that conditional averages are strong functions of ζ , training individual ANNs on data from one ζ allows for more accurate chemistry representation than training a single ANN on the entire ζ range. The individual ANNs can then be assigned to data from each discretized ζ in CSE.

ANNs can work with any combination of input and output nodes. To replace TGLDM, the ANN input and output nodes are matched to the quantities used in TGLDM. The ANN inputs are set to the parameterized variables in TGLDM, being $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$, with ζ being accounted for by the ANN itself. For outputs, TGLDM contains species mass fractions and source terms for all species in a given mechanism, however only certain species are considered in CSE due to storage limitations for the TGLDM tables. Currently, the conditional mass fractions of carbon monoxide, hydrogen and hydroxide ($\widetilde{Y_{CO}|\zeta}$, $\widetilde{Y_{H_2}|\zeta}$ and $\widetilde{Y_{OH}|\zeta}$), CH_4 , O_2 , CO_2 and H_2O reaction rates ($\widetilde{\dot{\omega}_{CH_4}|\zeta}$, $\widetilde{\dot{\omega}_{O_2}|\zeta}$, $\widetilde{\dot{\omega}_{CO_2}|\zeta}$ and $\widetilde{\dot{\omega}_{H_2O}|\zeta}$), temperature ($\widetilde{T|\zeta}$) and heat release rate ($\widetilde{\dot{\omega}_T|\zeta}$) are extracted from TGLDM. Thus, it is logical to have these nine quantities as the outputs from the ANNs. Note that $\widetilde{\dot{\omega}_T|\zeta}$ refers to the enthalpy source term as calculated below, where h_f is the standard enthalpy of formation in $\left[\frac{J}{kg}\right]$, and N is the number of species:

$$\widetilde{\dot{\omega}_T|\zeta} = \sum_{k=1}^N \left(h_k \widetilde{\dot{\omega}_k|\zeta} \right). \quad (4.1)$$

Optimization algorithms exist for designing the ANN architecture [98, 106], however it is more common to test various ANN architectures due to the large number of tunable parameters [78, 81, 86–88, 97, 99, 102, 104, 107]. Various test cases are shown here to justify various aspects of the ANN design process. Note that the augmented TGLDM dataset (described in Section 4.2.1) for diluted methane-air combustion with GRI-Mech 3.0 [70] chemistry representation is considered for all test cases, either using stoichiometric conditions or the entire ζ range. However, the conclusions extend to other mixture fractions, mechanisms, and sources of data. One set of cases is provided for each example, however multiple cases were tested before arriving to the present settings. It should be noted that testing all possible combinations of settings is prohibitive due to the number of tunable

parameters and the amount of data available; the aim is instead to show the general influence of these parameters for non-premixed methane-air combustion chemistry. It is also assumed that the optimized ANN architecture for GRI-Mech 3.0 TGLDM is feasible for all four cases.

Mean squared error (MSE) and coefficient of determination (R^2) are used to evaluate ANN performance. MSE is defined below, where N in this context represents the number of data points:

$$\text{MSE} = \frac{1}{9N} \sum_{i=1}^N \sum_{j=1}^9 (y_j - \hat{y}_j)^2. \quad (4.2)$$

MSE measures the average squared difference between an ANN prediction $\hat{y}_j$ and the target y_j from the dataset. R^2 measures how well the predicted data ‘fits’ the target data by comparing the average differences in normalized predicted and target quantities. The equation for R^2 is shown below:

$$R^2 = 1 - \sum_{i=1}^N \sum_{j=1}^9 \left(\frac{(\hat{y}_j - y_j)^2}{(\hat{y}_j - \bar{\hat{y}}_j)^2} \right), \quad (4.3)$$

where the numerator indicates the sum of residuals between the predictions and targets, and the denominator indicates the total sum of squares between the predictions and their mean values.

The first step in designing the ANN is to determine the number of neurons and hidden layers needed. Although hidden layers are not a requirement, it is beneficial to have hidden layers in ANNs to improve prediction accuracy. Including more hidden layers helps to differentiate non-linear data [108], which is relevant for combustion problems. The number of neurons in the hidden layers also plays an important role in the ANN prediction accuracy. Including more neurons per hidden layer increases the precision of the ANN, however this results in more weights and biases that must be tuned in the training process. Thus, the training time and the storage size increase with the number of neurons. The distribution of neurons also impacts the flow of information; it is possible to create a bottleneck if there are too few neurons in a hidden layer. Often, the neuron distribution in the hidden layers is equal or descending to aid in the flow of information.

The following example shows the influence of the neurons on the ANN performance, i.e., evaluating MSE, R^2 , training time and storage size. The ANN has two inputs ($\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$) and nine outputs ($\widetilde{Y_{CO}|\zeta}$, $\widetilde{Y_{H_2}|\zeta}$, $\widetilde{Y_{OH}|\zeta}$, $\widetilde{\dot{\omega}_{CH_4}|\zeta}$, $\widetilde{\dot{\omega}_{O_2}|\zeta}$, $\widetilde{\dot{\omega}_{CO_2}|\zeta}$, $\widetilde{\dot{\omega}_{H_2O}|\zeta}$, $\widetilde{T|\zeta}$, $\widetilde{\dot{\omega}_T|\zeta}$). The ReLU activation function is selected for the single hidden layer. The stoichiometric

mixture fraction is considered since it is the most relevant case. Data are split 80/20 into training and testing datasets, and Z-score normalization is applied. Note that the datasets are consistent for each case. The SGDM solver with default hyperparameters is used. A mini-batch size (i.e., number of samples used to update weights and biases) of 250 samples is selected, the validation frequency is set to 250 iterations, and the number of epochs (i.e. how many times to go through the entire training dataset) is chosen to be 100. Results for the test dataset MSE and coefficient of R^2 for each quantity, along with total ANN training time and storage size, are presented in Figure 4.1 for ANNs with an increasing number of neurons. The ANN accuracy (i.e., lower MSE, R^2 approaching unity) generally increases with additional neurons, however these changes are less significant with more neurons. The training time increases gradually, while the storage size increases linearly with the number of neurons. The cost-accuracy trade off occurs around 50 neurons, thus 50 neurons are selected.




Figure 4.1: a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for single-layer ANNs with increasing number of neurons. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates 50 neurons

The next test case considers the use of multiple hidden layers to improve the ANN accuracy. The number of neurons per hidden layer is set to 50. All other settings are identical to the previous case. Results for the same quantities are presented in Figure 4.2 for ANNs with an increasing number of hidden layers. The number of hidden layers has a more significant impact on training time and storage size than the number of neurons. The inclusion of hidden layers also allows for greater accuracy than adding more neurons

to a single-layer ANN. The training time and storage size increase linearly with additional hidden layers, however there are minimal changes in MSE and R^2 after five hidden layers. Thus, ANNs with five hidden layers are selected for the current study. It should be noted that the current architecture is consistent with other chemistry tabulation studies [86, 87, 89–92, 97, 102, 106], where the amount of hidden layers is typically 2-5 and the number of hidden layer neurons is typically 10-50.




Figure 4.2: a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for 50-neuron ANNs with increasing number of hidden layers. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates 5 hidden layers

Additional tests are performed to study the influence of the neuron distribution in the hidden layers. First, ANNs consisting of five hidden layers of either 25, 50 or 100 neurons each are compared with identical settings to validate the chosen architecture. The accuracy and training time increase marginally with more neurons while the storage size increases exponentially, as expected. The results are summarized in Table 4.1.

Next, two ANNs with five hidden layers are considered; the first ANN consists of 50 neurons in each hidden layer, while the second ANN consists of a decreasing number of neurons (128,64,32,16,10) for the same total of 250 neurons. All other settings are identical to the previous cases. Results for the same quantities are presented in Table 4.2 for these two ANNs. It can be seen that the training time is reduced with a decreasing number of neurons, however the storage size increases slightly. MSE and R^2 are nearly identical between the two cases, with a slight advantage for the even neuron distribution. The even

Table 4.1: MSE, R^2 , training time and storage size for ANNs with various even neuron distributions. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The 5x50 distribution is selected

	5x25	5x50	5x100
Mean squared error (MSE)	8.0827E-3	7.0150E-3	6.4028E-3
Coefficient of determination (R^2)	0.9919	0.9930	0.9936
Training time [s]	281.8	320.0	338.5
ANN storage requirement [kB]	31.5	120.0	473.6

architecture is selected due to better accuracy and lower storage size.

Table 4.2: MSE, R^2 , training time and storage size for ANNs with even and decreasing neuron distributions. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The even distribution is selected

	Even neurons (5x50)	Decreasing neurons
Mean squared error (MSE)	7.0150E-3	7.0185E-3
Coefficient of determination (R^2)	0.992992	0.992990
Training time [s]	320.0	300.2
ANN storage requirement [kB]	120.0	129.6

Representing the outputs with ANNs is open-ended. It is possible to predict all nine outputs from a single ANN for each discretized ζ , or assign one ANN to each output, resulting in nine ANNs for each discretized ζ . The following example illustrates the influence of the ANN output arrangement on the training accuracy and the storage size. The first arrangement consists of 50 ANNs (one per ζ) with two inputs ($\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$) and nine outputs ($\widetilde{Y_{CO}|\zeta}$, $\widetilde{Y_{H_2}|\zeta}$, $\widetilde{Y_{OH}|\zeta}$, $\widetilde{\dot{\omega}_{CH_4}|\zeta}$, $\widetilde{\dot{\omega}_{O_2}|\zeta}$, $\widetilde{\dot{\omega}_{CO_2}|\zeta}$, $\widetilde{\dot{\omega}_{H_2O}|\zeta}$, $\widetilde{T|\zeta}$, $\widetilde{\dot{\omega}_T|\zeta}$). The second arrangement consists of 450 ANNs (nine per ζ) with the same inputs but each ANN has a single output corresponding to a given quantity above. The ANNs in the first arrangement consist of five hidden layers with 50 neurons each, while the ANNs in the second arrangement consist of two hidden layers with 14 neurons each. This results in a similar total number of neurons; there are a total of 12,500 hidden neurons in the first arrangement and 12,600 hidden neurons in the second arrangement. All other settings are identical to the previous test cases, however the entire ζ range is used here. Note that the training and testing datasets are identical for each arrangement. Results for MSE, R^2 and training time (full ζ range and zoomed-in ζ range) are presented in Figure 4.3 for both arrangements. Note that the averages of MSE and R^2 have been taken with the 450x1 ANN case, however

the quoted training time is a sum of the individual ANN training times (which were similar for each quantity at a given ζ). The accuracy of both arrangements are similar for a given quantity, with the 50x9 ANN case performing slightly better in general. One advantage of the 450x1 ANN case is the reduced storage size of 1.2 MB compared to 6.0 MB for the 50x9 ANN case, despite having three times as many files. However, the training time is significantly larger with 450 ANNs. The 50x9 ANN case is selected for better accuracy and lower training time.




Figure 4.3: a) MSE, b) R^2 , and c) training time for ANNs with multiple outputs (50x9, blue) or single outputs (450x1, red). The bottom row is a reproduction of the top row in a reduced ζ range. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 across all ζ . The vertical dashed line indicates the stoichiometric ζ of 0.353

The activation function is the final component of the ANN architecture. The activation function turns neurons ‘off and on’ by transforming the output of a given neuron based on a prescribed function. Various activation functions exist, including ReLU [109], leaky ReLU, hyperbolic tangent (tanh), exponential linear unit (ELU), Gaussian error linear unit (GELU), sigmoid/swish, and many others. ReLU is a piecewise function that transforms negative outputs to zero, while positive outputs remain unchanged. ReLU is an efficient function and has been applied successfully to previous ML combustion chemistry cases [95, 102, 104–106], However, training with ReLU can be difficult since it is unbounded

and non-differentiable at zero. It is also possible to end up with neurons that are inactive (i.e., the neuron output is consistently zero) if the learning rate is large; this is known as the ‘dying ReLU’ problem [110]. Functions such as leaky ReLU address the dying ReLU problem, at the potential cost of performance [110]. Continuous functions, such as the other examples provided above, are derived from the ReLU function and aim to improve smoothness. Note that the activation function is applied to the hidden layers only; the output layer uses a linear function to capture negative values.

The next test case considers the choice of activation function. The same ANN architecture is used for each case: five hidden layers with 50 neurons each. Each ANN has nine outputs described above. All other settings are identical to previous cases. Results for the same quantities are presented in Table 4.3, however the storage size was found to be consistent and is removed from this table. The training accuracy is fairly consistent (aside from the sigmoid function), however the training time is reduced significantly with both the ReLU and leaky ReLU activation functions. In addition, the ReLU activation function is the most simple to code, which is expected to aid the computational performance. Thus, the ReLU activation function is selected.

Table 4.3: MSE, R^2 , and training time for ANNs with various activation functions. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The ‘ReLU’ activation function is selected

	ReLU	Leaky	ELU	GELU	tanh	swish	sigmoid
MSE	7.015E-3	7.147E-3	8.021E-3	7.456E-3	7.464E-3	7.976E-3	1.466E-2
R^2	0.99299	0.99286	0.99199	0.99255	0.99254	0.99203	0.98535
Time [s]	320.0	331.5	464.6	371.9	438.5	434.6	433.2

4.2 ANN Data Pre-processing

The data preparation procedure for ANN development is described in this section. Two sources of data are considered: tabulated chemistry via TGLDM, and direct chemistry via sampling of conditional averages from CSE. The pre-processing procedure is similar for both sources of data and is analogous to any fuel type or mechanism.

4.2.1 Tabulated Chemistry Data Preparation

The TGLDM dataset described in Section 2.5.3 is used to train and test each ANN for the applications in Chapter 5 and parts of Chapter 6. However, it was found from preliminary testing that using the TGLDM data itself for ANN training led to instabilities when paired with CSE. It was thus required to refine and prepare the TGLDM dataset for ANN training. First, non-unique data points are removed from the TGLDM tables. These occur in fuel-rich manifolds where the trajectories may intersect near the chemical equilibrium point, which creates unrealistically large values for conditional minor species mass fractions. Similar observations have occurred in different tabulation methods [111, 112]. Next, additional points are added outside of the TGLDM data range to account for cases in which the ANN reaction rate predictions lead to new estimates in $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$ that fall outside of the standard range. It was assumed that $\widetilde{Y_{H_2}|\zeta}$ and $\widetilde{Y_{OH}|\zeta}$ were zero in the outside range due to these quantities being zero at the TGLDM boundaries. $\widetilde{Y_{CO}|\zeta}$, however, is non-zero at the TGLDM boundaries, thus the $\widetilde{Y_{CO}|\zeta}$ value in the outside range was calculated from the balance of CO_2 , H_2O , CH_4 and O_2 conditional mass fractions. It should be stressed that elemental conservation is satisfied in the outside border and TGLDM itself. The conditional temperature is then calculated using the standard enthalpy of formation and specific heat capacity for each species. The four reaction rates and the enthalpy source term are assigned a value of zero. Finally, Delaunay triangulation [71] is applied three times to generate additional points via interpolation. This results in greater coverage of the input space, especially in the regions with large output gradients, as shown in Figure 4.4. This process is completed for pure and diluted methane-air combustion mechanisms, which in Chapter 5 are both represented by GRI-Mech 3.0 [70] containing 53 species and 325 reactions, and in Chapter 6 for diluted methane with Smooke’s mechanism [113] containing 16 species and 35 reactions. The number of input-output pairs increases by a factor of about 130 from the base TGLDM dataset to the augmented TGLDM datasets for each case. The number of total samples is about 17.6 million and 11.7 million for pure and diluted methane with GRI-Mech 3.0, and 8.5 million for diluted methane with Smooke’s mechanism; this is on the order of other training datasets in similar ANN applications [95, 96, 98, 102, 103, 106]. The ζ distribution is logarithmic for GRI-Mech 3.0 and linear for Smooke’s mechanism. Note that most of the data are concentrated near the stoichiometric mixture fraction for each case.




Figure 4.4: Illustration of TGLDM data augmentation for ANN training, considering pure methane-air combustion at $\zeta = 0.2$: a) base TGLDM data, b) overlapping trajectory is removed, c) additional points are added near the boundaries as constraints, d) Delaunay triangulation is used to generate additional points

4.2.2 Direct Chemistry Data Preparation

A dataset constructed from samples of CSE with direct chemistry (DI-CSE) is used to train and test one set of ANNs for the application in Chapter 6. Results from [33] for Sandia flame D are used as the starting point. The chemistry is represented by Smooke’s mechanism [113]. To generate a sufficient range of conditional averages, the main fuel jet velocity is prescribed as a cosine function with a set amplitude and frequency. The function is defined such that the minimum velocity is equivalent to the Sandia flame D jet velocity of $49.6 \left[\frac{m}{s}\right]$. Various amplitudes and frequencies are tested, in addition to the sampling rate (i.e., how often data are saved). Following various tests, an amplitude of $40 \left[\frac{m}{s}\right]$, a frequency of $100 [Hz]$ and a sampling rate of 250 time steps over a period of 8 cycles are selected. This results in conditional averages that are representative of Sandia flames D to F. An illustrative example of the influence of these parameters on the conditional averages is shown in Figure 4.5, considering a sample range of 10 seconds from steady-state conditions.

Sandia flame D with the specified periodic velocity boundary condition is simulated using DI-CSE from the steady-state solution, which arises from the set velocity boundary condition. Conditional averages in each ensemble are extracted from the simulation at the specified sample rate, restructured into tables for each discretized mixture fraction, and saved as text files. Samples exceeding the maximum $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$ (from TGLDM) are removed. Delaunay triangulation [71] is then applied three times consecutively, as above. Note that no additional points are manually set with the DI data. The number of raw input-output pairs is 76,799, however this increases by a factor of 44 for a total of about 3.4 million input-output pairs; this is on the order of other training datasets in similar ANN applications [98, 102–104, 106]. Note that in this approach, the mixture fraction distribution is linear and the number of base samples for each discretized mixture fraction is identical.

4.3 ANN Training and Testing Process

This section describes the ANN training and testing process and the specific ANN training/testing settings used in the current study. This includes the selection of the loss function and the iterative method to optimize the prescribed loss function, and the determination of the training settings (i.e., mini-batch size, epochs, hyperparameters). The MATLAB Deep Learning Toolbox is used for training and testing the ANNs. All cases are ran on a CPU with a 5 GHz processor.




Figure 4.5: Influence of velocity amplitude (columns) and frequency (rows) on $\widetilde{Y_{CO}}|\zeta$ at $\zeta = 0.353$ for Sandia flame D with DI samples. Data is saved every 250 time steps (top two rows) or 25 time steps (bottom two rows)

In this study, the augmented TGLDM and DI datasets for each case, as described in Section 4.2 above, are utilized for training separate sets of ANNs. The dataset is traditionally split into training and testing datasets. It is common practice [78, 97, 104, 106] to split 80% of the data into the training dataset, and the remaining 20% into the testing dataset. In order to obtain adequate coverage of the composition space, data are sorted randomly into these datasets for each ANN. The purpose of the training process is to optimize the weight and bias values to the training dataset. These values are obtained using ‘backpropagation’, in which the gradient of a prescribed loss function is passed from the output layer to the input layer to update the weight and bias values based on an iterative method. In contrast, the testing dataset is completely independent from the training dataset and instead provides a measure of accuracy for ‘unseen’ data. The testing process may be performed concurrently with the training process, however the testing dataset does not influence the ANN weights and biases. Note that it is possible to define a third ‘validation’ dataset that is independent of both the training and testing datasets, however this is not consistently used in the ML combustion literature. The validation dataset would be used to tune the ANN hyperparameters during the training process, however the testing completed in this chapter is used to determine the optimal ANN hyperparameters before the training process instead.

It is common to normalize each dataset before training/testing the ANNs. The purpose of normalization is to enforce identical statistical properties for each variable, thus reducing non-linearities and improving the speed and accuracy of the ANN training and testing process [108]. There are multiply ways to normalize the datasets. Common methods are min-max and Z-score (normal distribution) [108]. The min-max function scales the range of each input/output to (0,1), whereas the Z-score function scales each input/output to a normal distribution, such that the mean is zero and the variance is unity. Note that it is possible to normalize the ANN weights and the data within the mini-batch, however both techniques were found to be unnecessary.

The next tests involve these normalization schemes for the same ANN architecture (five hidden layers with 50 neurons each) with the ReLU activation function. All other settings are identical to previous cases. Results for the same quantities (aside from storage size, which is consistent) are presented in Table 4.4. It is observed that the normalized training error is one order of magnitude larger in the Z-score case due to the wider data ranges; min-max scales to (0,1), whereas Z-score typically scales to (-3,3). However, there are fewer errors when rescaling data with Z-score versus min-max, as can be seen in Table 4.5, due to the Z-score data being zero-centered. Thus, Z-score normalization is selected in the current study.

The loss function is a measure of accuracy between the ANN predictions and the ‘tar-

Table 4.4: MSE, R^2 , training time and storage size for ANNs with min-max and Z-score normalization (selected). Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta=0.353$

	min-max	Z-score
Mean squared error (MSE)	5.9340E-4	7.0150E-3
Coefficient of determination (R^2)	0.9923	0.9930
Training time [s]	311.2	320.0

Table 4.5: MSE (top) and R^2 (bottom) of rescaled variables for ANNs with min-max and Z-score normalization (selected). Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta=0.353$

	$\widetilde{Y_{CO}} \zeta$	$\widetilde{Y_{H_2}} \zeta$	$\widetilde{Y_{OH}} \zeta$	$\widetilde{\dot{\omega}_{CO_2}} \zeta$	$\widetilde{\dot{\omega}_{H_2O}} \zeta$	$\widetilde{\dot{\omega}_{CH_2}} \zeta$	$\widetilde{\dot{\omega}_{O_2}} \zeta$	$\widetilde{\dot{\omega}_T} \zeta$
min-max	1.29E-6	4.79E-10	2.06E-8	923.1	1578.0	38.0	4109.8	1.89E+17
Z-score	2.50E-7	2.09E-10	1.17E-8	423.5	939.9	14.2	2634.0	1.12E+17
min-max	0.9979	0.9980	0.9936	0.9964	0.9570	0.9863	0.9748	0.9895
Z-score	0.9996	0.9991	0.9963	0.9984	0.9744	0.9949	0.9838	0.9937

gets', i.e., the outputs from the training dataset. The objective is to minimize the loss function during the training process by updating the ANN weights and biases. MSE (defined in Section 4.1) is commonly used as a loss function for regression problems and can provide accurate results [108]. Note that other loss functions, including physics-based loss functions [97], exist but are not within the MATLAB library; instead, a manual training loop must be defined to use these loss functions. MSE with mass conservation was tested but found to produce less accurate predictions.

An iterative method is required to optimize the loss function. Common algorithms include Stochastic Gradient Descent with Momentum (SGDM) [108] and Adaptive Moment Estimation (ADAM) [114]. The SGDM algorithm considers the optimization of a parameter (which, in this context, are the ANN weights and biases) that minimizes the loss function. Rather than computing the actual loss function gradient of the entire dataset, SGDM calculates the loss function gradient of the mini-batch and updates the weights and biases using the results from each mini-batch. The loss function gradient is multiplied by a 'learning rate' (i.e., step size), which may be fixed or variable. The momentum term is included to smooth the changes to the weights and biases by using a linear combination of the results from the current and previous steps. SGDM allows for efficient ANN training with large datasets at the expense of convergence [115]. The ADAM solver is a recent

Table 4.6: MSE, R^2 , training time and storage size for ANNs with the SGDM and Adam solvers. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The SGDM solver is selected

	SGDM	Adam
Mean squared error (MSE)	7.0150E-3	6.7696E-3
Coefficient of determination (R^2)	0.9930	0.9932
Training time [s]	320.0	1154.3
ANN storage requirement [kB]	120.0	128.2

development derived from the Root Mean Square Propagation (RMSProp) model, which considers a running average of the gradient magnitudes to adjust the learning rate, but is extended to include the momentum term.

Two solvers (SGDM and Adam) are applied to the same ANN architecture (five hidden layers with 50 neurons each) with Z-score normalization and the ReLU activation function. All other settings are identical to previous cases. Note that RMSProp is not considered here, as Adam is an extension to RMSProp that includes momentum. Results for the same quantities are presented in Table 4.6. It is observed that both solvers produce accurate results, with Adam performing better than SGDM. However, the Adam solver is almost four times slower than SGDM. In addition, many of the weights are set to very small numbers with the Adam solver, which results in more precision and storage needed. For these reasons, SGDM is selected for the current study.

The training process involves the tuning of the ANN weights and biases by reducing the loss function between the targets (i.e outputs from the training dataset) and the predictions (i.e., outputs calculated from the training dataset inputs). These predictions are obtained using a ‘mini-batch’, which consists of a subset of unique training samples. The weights and biases are updated using the predictions from the mini-batch, and the calculations are repeated for the next mini-batch of data. The mini-batch size (i.e., number of samples) impacts the accuracy and speed of the training process; having few samples results in high sensitivity and longer training time, while having too many samples results in lower training time but the underlying patterns in the data may not be captured as well. The calculation of weights and biases from the mini-batch is considered as one iteration. An ‘epoch’ is defined as the time period when the entire training dataset has been used in a mini-batch. After each epoch, the training dataset is randomly shuffled such that the dataset within each mini-batch is ‘new’; this allows for more robust ANNs. The training process continues for a set number of epochs. Note that the testing process is performed concurrently at a frequency of 250 iterations, i.e., every 250 mini-batch calculations.

The training settings (i.e., mini-batch size and number of epochs) are investigated for the same ANN architecture using the SGDM solver. All other settings are identical to previous cases. Results for the same quantities are presented in Figure 4.6 for various mini-batch sizes with a fixed number of 100 epochs. The mini-batch size has a strong influence on the training accuracy. It is desired to have a small number of points to calculate the weights accurately. However, making the mini-batch size too small results in overfitting and exponentially increased training time. A mini-batch size of 250 is selected; this value does not yield the optimal training accuracy, however the training time is feasible. Following this, the number of epochs is adjusted with a fixed mini-batch size of 250 samples. These results are presented in Figure 4.7. The training accuracy improves in an exponentially decaying manner with more epochs, however the training time increases linearly with the number of epochs. There is a small increase in storage size with more epochs. The number of epochs is set to 100 based on the point where the increase in training time outweighs the increase in training accuracy.




Figure 4.6: a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing batch sizes. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates a mini-batch size of 250

The SGDM solver contains a number of hyperparameters that must be tuned. These include the ‘learning rate’, the ‘momentum’, and the ‘L2 regularization’ term. The learning rate is the weight of the loss function gradient. It is effectively the ‘step size’ in SGDM, i.e., the influence of new information on the update of the current information. A large learning rate may cause overshoot and lead to instability, whereas a small learning rate




Figure 4.7: a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing number of epochs. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates 100 epochs

takes longer to converge and may be trapped in a local minimum. The learning rate can be set to a constant value, be set to a schedule based on the amount of training progress, or be adaptable to prevent these issues. The momentum term is an extension to standard SGD that aims to avoid local minima and reduce training time by increasing the learning rate. The momentum term is a factor between zero and unity that is applied to the change in the weight or bias from the previous update. SGDM then considers a linear combination of the gradient and the change from the previous iteration. Finally, the L2 regularization term is included to smooth the solution and reduce overfitting; it is analogous to the Tikhonov matrix in Equation 2.50 above. A small value makes the predictions more sensitive to noise, while a large value results in too much smoothing and may not capture non-linear effects.

The hyperparameters (i.e., learning rate, momentum term and L2 regularization term) are investigated for the same ANN architecture using the training settings described above. The learning rate is investigated using various fixed quantities, as well as various schedules. Results for the same quantities are presented in Figure 4.8 for various constant learning rates. The ANN accuracy improves with a larger learning rate, while the training time is generally reduced slightly. Instability occurs when the learning rate is too large, which in this case occurs around 0.1. A constant learning rate of 0.01 is selected. Following this, the

momentum term is varied. Results for the same quantities are presented in Figure 4.9. The results are fairly consistent with improvements around 0.9, but the results become unstable as the momentum term approaches unity. A value of 0.9 is selected for the momentum term. Finally, the L2 regularization term is varied. Results for the same quantities are presented in Figure 4.10. Reducing the L2 regularization term (i.e., no penalty to the loss function) leads to improved accuracy, however this result could be due to overfitting. A value of 0.0001 is selected for the L2 regularization term based on sufficient accuracy without overfitting. Note that these correspond to the default MATLAB settings, however these default settings have been selected due to their success in various cases.




Figure 4.8: a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing learning rate. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates a learning rate of 0.01

The previous cases consider a fixed learning rate of 0.01. A series of tests were conducted where the learning rate was specified as a piecewise function, starting from a value of 0.01. The learning rate decay factor (i.e., multiplier for the initial learning rate) is selected from 0.01 to 0.9, while the learning rate drop period (i.e., number of epochs until the learning rate changes) is selected from 5 to 75 epochs. The results for MSE and R^2 are shown in Figure 4.11, with dashed lines indicating results from a fixed learning rate of 0.01. Slight accuracy improvements from the fixed learning rate case can be achieved with a proper learning rate schedule. The accuracy tends to improve when the learning rate factor and period increase, however this is not always the case. The storage size increases marginally




Figure 4.9: a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing momentum. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates a momentum factor of 0.9




Figure 4.10: a) MSE (blue) and R^2 (red), and b) training time (green) and storage size (black) for ANNs with increasing L2 regularization. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. The vertical dashed line indicates an L2 regularization factor of 0.0001

with learning rate factor and period, and is not shown for brevity. No significant trends were observed in the training time, however the times were slightly larger than the constant learning rate case.

For this work, the constant learning rate of 0.01 is selected for the TGLDM cases for simplicity and to reduce training time. For the DI case, however, it was necessary to reduce the learning rate to achieve convergence across all ζ . Thus, a constant learning rate of 0.025 is selected for the DI case. Note that all other settings remain identical between ANNs with TGLDM and DI training data.




Figure 4.11: a) MSE and b) R^2 for ANNs with various learning rate drop factors and periods. Training data is augmented TGLDM for diluted methane-air with GRI-Mech 3.0 at $\zeta = 0.353$. Horizontal dashed lines indicate results with a constant learning rate

4.4 ANN Training Results

This sections contains results from the ANN training and testing process for each set of ANNs in the study. The results are organized by the chemical mechanism, and hence the subsequent chapters in this work.

4.4.1 TGLDM with GRI-Mech 3.0 for Pure and Diluted Methane

The MSE and R^2 are determined from the testing dataset for both fuels across all mixture fractions. Figure 4.12 shows the MSE and R^2 for each ANN in the pure methane (left) and diluted methane (right) cases. Most of the ANNs have excellent accuracy for both fuels, indicated by smaller MSE (approaching zero) and larger R^2 (approaching unity). There are outliers towards pure air conditions (ζ approaching zero) for both fuels, as well as pure fuel conditions (ζ approaching unity) for the pure methane case. It should be noted that there is no combustion in these regions, thus the errors are expected to have a minimal impact on the flame. There are also fewer points in these regions compared to the stoichiometric mixture fraction, hence training is more difficult towards pure fuel and pure air.

The average MSE is approximately 7.44×10^{-2} for pure methane and 2.13×10^{-2} for diluted methane, however the median MSE is approximately 1.04×10^{-2} for pure methane and 4.22×10^{-3} for diluted methane. Note that the MSE results are scaled to the normalized data, which tend to have ranges of $(-3,3)$ for each quantity. This indicates that the MSE is often two to three orders of magnitude lower than the data range of a given variable when the data is rescaled, i.e., species mass fractions will have errors on the order of 10^{-3} , and species reaction rates will have errors on the order of $10^1 \frac{1}{s}$.

The average R^2 is approximately 0.918 for pure methane and 0.975 for diluted methane (due to outliers), however the median R^2 is approximately 0.990 for pure methane and 0.996 for diluted methane. This indicates that the ANN predictions often have good agreement to the test dataset, stemming from TGLDM with GRI-Mech 3.0. R^2 is reduced towards pure air and pure fuel due to the data points being similar to each other, resulting in a smaller denominator in Equation 4.3.

An additional test was performed to analyze the influence of data volume on the ANN training results, specifically in the fuel-rich regions. The training data was increased by a factor of four via an additional Delaunay interpolation. Slight improvements in MSE and R^2 were achieved, however the training time increased by up to a factor of nine. It was concluded that three Delaunay interpolations was sufficient. It should also be noted that larger ANNs were also tested for the fuel-rich regions, which resulted in minimal accuracy improvements while significantly increasing the training time and storage.

Note that the training of certain fuel-lean and fuel-rich ANNs in each case was iterated up to ten times, with the best-performing ANN being selected for each ζ . With this method, the total training time is about 16 hours for the pure methane dataset, and 7 hours for the diluted methane dataset, on a single core of a 5 GHz processor. The training time is reasonable given the volume of data. The training time for each case also does not




Figure 4.12: MSE and R^2 for all 50 ANNs using GRI-Mech 3.0 with a) pure methane and b) diluted methane chemistry. The vertical line indicates the stoichiometric mixture fraction

exceed the time savings from ANN-CSE relative to CSE with tabulated chemistry, as will be seen in Section 5.3.5.

The storage requirements of both sets of ANNs are compared to the TGLDM datasets in Table 4.7 below. The augmented TGLDM dataset is included to show the direct influence of ANNs, however this dataset is not used or stored in CSE. Two observations can be made from Table 4.7: the ANN storage requirement is much smaller than the TGLDM storage requirement for both fuels, and the ANN storage requirement is independent of the data source. The direct savings (comparing ANN to augmented TGLDM) are about 99% for both fuels, however the actual savings (comparing ANN to base TGLDM) are about 78% and 66% for pure and diluted methane, respectively. The storage size is due to the amount of information to be saved. TGLDM-CSE uses only 50 text files (one per ζ), however the TGLDM storage is greater since 11 variables need to be saved for each location in the ζ , $\widetilde{Y_{CO_2}}|\zeta$ and $\widetilde{Y_{H_2O}}|\zeta$ space, which results in many data points required. In contrast, ANN-CSE requires 600 files, or six weight and six bias files for each mixture fraction, however these files are relatively small. There are a total of 527,500 weights and 12,950 biases, or 10,550 weights and 259 biases per mixture fraction, which constitute the data points for the ANN method. It should be noted that both formulations of CSE require additional files which have negligible sizes.

Table 4.7: Amount of data and storage size for the original TGLDM dataset, augmented TGLDM dataset, and the weights and biases of the TGLDM-based ANNs for both fuels

	Base TGLDM	Augmented TGLDM	ANN
Pure methane-air	1,470,095 points 27.9 MB	193,640,634 points 3500 MB	540,450 points 6.0 MB
Diluted methane-air	929,962 points 17.4 MB	129,097,254 points 2200 MB	540,450 points 6.0 MB

In addition, the ANN-CSE storage requirement only depends on the architecture and not on the dataset, whereas the TGLDM storage requirement depends on the level of discretization. (This can be seen from base to augmented TGLDM for either fuel.) Adding more species is also more flexible with ANNs than TGLDM: for each additional species (or output), there would be 2550 additional points with the ANN method (i.e., 2500 more weights and 50 more biases), or there would be 84,500 to 17.6 million more points with the TGLDM method, depending on the dataset. This shows that data scaling with TGLDM is about 33 to 6900 times greater than ANNs. While 3.4 GB of storage is manageable, adding more detail to the TGLDM datasets (via more species or more trajectories) would make them infeasible for CSE. This shows the potential of ANNs for representation of chemical kinetics with more complex fuels, which may contain hundreds of species.

4.4.2 TGLDM and DI Samples with Smooke’s Mechanism for Diluted Methane

The same analysis is now provided for the TGLDM and DI cases involving Smooke’s mechanism. Figure 4.13 shows the MSE and R^2 for each mixture fraction, with the TGLDM source on the left and the DI source on the right. The TGLDM case has excellent accuracy, with R^2 consistently above 0.9 and MSE consistently on the order of 10^{-2} . The stoichiometric region has the best performance, while fuel-lean and fuel-rich ANNs have slightly worse performance as expected due to the relative lack of data points in the manifolds. The DI case, on the other hand, has similar trends but the accuracy changes rapidly with each ANN. The accuracy is sufficient around stoichiometry, but is significantly worse in the fuel-lean and fuel-rich regions. This behaviour is attributed to the DI sample space being less structured than TGLDM, with all target variable being highly non-linear.

The average MSE and R^2 are approximately 3.40×10^{-2} and 0.966, respectively, for ANNs trained on the TGLDM-based dataset, indicating good accuracy. This is further




Figure 4.13: MSE and R^2 for all 50 ANNs using Smooke’s mechanism with a) TGLDM solutions and b) DI samples. The vertical line indicates the stoichiometric mixture fraction

supported by median values of 3.86×10^{-2} and 0.962 for MSE and R^2 , respectively, which are very similar to the averages. In contrast, ANNs trained on the DI samples have an average MSE and R^2 of approximately 1.75×10^{-1} and 0.825, respectively. Both values are considerably worse than the TGLDM cases, however the R^2 is large enough to indicate that the general trends of the output variables are captured by the ANNs. For comparison, the median values of MSE and R^2 are slightly improved at 1.68×10^{-1} and 0.864, respectively, indicating that outliers have a small effect on the DI-based ANNs.

The total training time is about 2.5 hours for the TGLDM case and 3 hours for the DI case on a single core of a 5 GHz processor. Note that iteration was performed for certain fuel-lean mixture fractions in TGLDM, but not for DI. Iterating the DI case may benefit the accuracy, however the entire ζ range would have to be investigated, and the additional time may negate the time savings in CSE.

The storage requirements of the TGLDM and DI cases with Smooke’s mechanism are presented in Table 4.8. As observed in Section 4.4.1, the ANN storage requirement is small and independent of the source of data; the small differences in storage are attributed to data precision. The TGLDM and DI samples initially have similar storage requirements, however the trimming and augmentation process results in much more TGLDM samples relative to DI. The direct savings (comparing ANN to augmented) are over 99% for each case, however the actual savings (comparing ANN to base) are about 55.5% for the TGLDM case and 51% for the DI case. Increasing the number of augmented DI samples to the same

Table 4.8: Amount of data and storage size for the original and augmented TGLDM and DI datasets for diluted methane-air combustion with Smooke’s mechanism, and the weights and biases of the ANNs for each case

	Base	Augmented	ANN
TGLDM samples	721,171 points 13.3 MB	93,488,813 points 1600 MB	540,450 points 5.9 MB
DI samples	844,789 points 11.9 MB	37,205,542 points 554.2 MB	540,450 points 5.8 MB

order as TGLDM may improve training results (refer to Figure 4.13) and would result in greater storage savings with ANNs. The TGLDM storage requirement is modest, but the potential of ANN for representing detailed chemical kinetics is demonstrated. It should be noted that the DI storage comes solely from the training sample generation; DI-CSE does not require the storage of chemical information. This is an advantage compared to ANN, however the purpose of ANN is to replace the tabulated chemistry approach, not DI.

4.5 CSE Implementation with ANN Chemistry

The majority of the CSE implementation with ANN chemistry representation is identical to the implementation of CSE with tabulated chemistry, as described in Section 2.5.3. The difference is how the remaining conditional averages are determined. In the previous approach (TGLDM-CSE), the TGLDM tables are loaded initially, and a search/interpolation method is used at each time step to determine the values of the other species mass fractions and source terms based on $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$. These two steps are replaced by ANNs in the current approach (ANN-CSE). First, the weights and biases of all 50 ANNs (for a given fuel) are saved as text files. These are loaded once at the beginning of the simulation and saved as matrices. Then, given a series of inputs $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$, the nine outputs ($\widetilde{Y_{CO}|\zeta}$, $\widetilde{Y_{H_2}|\zeta}$, $\widetilde{Y_{OH}|\zeta}$, $\widetilde{\dot{\omega}_{CH_4}|\zeta}$, $\widetilde{\dot{\omega}_{O_2}|\zeta}$, $\widetilde{\dot{\omega}_{CO_2}|\zeta}$, $\widetilde{\dot{\omega}_{H_2O}|\zeta}$, $\widetilde{T|\zeta}$ and $\widetilde{\dot{\omega}_T|\zeta}$) can be calculated by applying Equation 3.1 sequentially to each hidden layer plus the output layer for each ζ . This process replaces the search procedure in TGLDM-CSE and avoids the need for interpolation. The inversion process and the calculation of Favre averages remains the same between CSE formulations.

4.6 Summary

In this chapter, a methodology for representing turbulent combustion chemistry with ANNs in the context of the CSE combustion model is presented and validated. The ANN architecture and design decisions are thoroughly described. The data pre-processing procedure is described for both sets of data in the study. The ANN training and testing process is then described in detail, followed by a brief description on the implementation of ANNs into CSE. The ANN training results show that the ANN storage requirement is smaller than either data source and is independent of the number of data points. The next chapter provides the application of TGLDM-based ANN-CSE with GRI-Mech 3.0 to two turbulent non-premixed methane jet flames.

Chapter 5

ANN-CSE for Turbulent Non-premixed Jet Flames with Pure and Diluted Methane

The objective of this chapter is to assess the influence of chemistry representation in CSE in the context of turbulent non-premixed jet flames. This chapter provides the development and application of ANN-CSE to two well-known turbulent non-premixed methane jet flames: one with pure methane, and one with diluted methane. Tabulated chemistry via TGLDM using GRI-Mech 3.0 is considered for ANN development in this chapter. ANN-CSE is used to simulate both flames. The experimental methods, ANN development, and computational settings are described for both flames. Results for the mixing field, conditional and Favre averages, and computational cost from ANN-CSE are then analyzed for both flames and compared to results using CSE with tabulated chemistry (TGLDM-CSE). Three locations are considered for each flame, however the conclusions apply to other locations in each flame. A summary of findings and recommendations for this work is provided at the end of this chapter.

5.1 Case 1: Brookes-Moss Flame

The pure methane jet flame experimentally studied by Brookes and Moss [47] is the first flame under consideration in this chapter. This flame has been previously simulated using CSE with tabulated chemistry [28, 31]. This section describes the experimental setup of

the Brookes-Moss flame, the ANN development for pure methane, and the computational settings for the Brookes-Moss flame.

5.1.1 Experimental Setup

The Brookes-Moss flame is a confined turbulent methane-air flame. The case with atmospheric pressure is considered. The experimental setup consists of a cylindrical 4.07 mm diameter nozzle for the main fuel flow and an annular premixed pilot nozzle with a width of 160 μm for stabilization. A co-flow of air is present around the main fuel nozzle. Two circular perforated metal plates are mounted to the burner to flatten the co-flow velocity profile and prevent recirculation near the base of the flame. The flame is confined by a 155 mm diameter Pyrex tube, which is split at the midpoint for probing and separated by an 8 mm thick stainless steel spacer. The co-flow of air has a mass flow rate of 708 g/min. The main fuel stream consists of pure methane and has a mass flow rate of 10.3 g/min, a mean velocity of 20.3 m/s, and an exit Reynolds number of 5000. The co-flow air and main fuel enter the combustion chamber at 290 K. The pilot flame is a premixed methane-air flame with a methane flow rate of less than 2% of the main fuel flow rate. The calculated mean experimental error for mixture fraction is reported to be 6% [47]. There is no experimental error provided on the temperature measurements.

5.1.2 ANN Preparation

Tabulated chemistry via TGLDM [13] with GRI-Mech 3.0 [70] for pure methane is used to develop ANNs for the Brookes-Moss flame. The training data is augmented following the procedure outlined in Section 4.2.1. The ANN weights and biases are determined from the training and testing procedure with optimized settings, as described in Section 4.3. The combined procedure took about 16 hours for all 50 ANNs on a single core of a 5 GHz processor. The weights and biases of each ANN are saved as text files to be used in ANN-CSE, which is described in Section 4.5.

5.1.3 Computational Details

The Brookes-Moss flame is simulated with both ANN-CSE and TGLDM-CSE using OpenFOAM v7 [74]. For consistency, many of the computational details are taken from previous CSE simulations of the Brookes-Moss flame [28, 31]. The numerical settings are identical

for ANN-CSE and TGLDM-CSE, including the time step, turbulence settings, computational grids and boundary conditions. Combustion is triggered by setting a small amount of $\widetilde{Y_{CO_2}}$ and $\widetilde{Y_{H_2O}}$ (0.06 and 0.1, respectively) in the domain, with a balance of air. The same numerical settings for both CSE implementations ensures that potential differences in the results can only be attributed to differences in the chemistry representations.

The computational domain is cylindrical with a length of 1000 mm and a radius of 77.5 mm. Following a previous CSE study [28], the geometry is defined as a small angle wedge ($\theta = 10^\circ$) with one cell in the θ direction to create a two-dimensional domain. This is justified since the dominant motion of the flame is in the axial direction, and there is no swirl in the flame. The CFD mesh consists of 58,800 non-uniform rectilinear cells that are concentrated near the main fuel nozzle in both the axial and radial directions. A total of 25 uniform CSE ensembles are defined in the axial direction and this arrangement is shown to provide CSE ensemble independent results [28, 31]. The mixture fraction sample space is discretized into 50 bins, with more bins concentrated near the stoichiometric mixture fraction of 0.055. The co-flow and main fuel jet velocities and temperatures are defined as Dirichlet boundary conditions using the respective values from the experiment. Following previous studies [28, 31], the k - ϵ turbulence model [54] with a modified coefficient of $C_{\epsilon 1} = 1.47$ is used. The maximum Courant number is set to 0.9, and the flame is simulated to 10 s with data saved every 0.01 s. The Favre-averaged transport equations for mass, momentum, species mass fractions (CO_2 and H_2O), enthalpy, and mean and variance of the mixture fraction [50] as described in Chapter 2 are solved using OpenFOAM. All spatial terms are discretized using second-order Gaussian integration, while the temporal terms are discretized using an implicit Euler scheme. Both CSE cases are run on a single core of a 5 GHz processor.

An illustration of the computational domain is provided in Figure 5.1. The outlet is located at the top of Figure 5.1 a). A table of boundary conditions for each quantity and surface, using OpenFOAM definitions, is provided in Table 5.1. Note that all quantities along the axis are specified as “empty”, and all quantities at the front and back faces are specified as “wedge”. All “fixed” boundary conditions are uniform.

5.2 Case 2: Sandia Flame D

The well-known diluted methane-air Sandia flame D is the second flame under consideration in this chapter. Sandia flame D has been previously simulated with both CSE [30, 32, 33] and ML techniques [85, 89, 90]. This section describes the experimental setup of Sandia




Figure 5.1: Illustration of computational domain for Case 1 (Brookes-Moss flame): a) temperature contour with CSE ensemble boundaries shown in black and axial locations of results shown in red, b) mesh details near the main fuel inlet

flame D, the ANN development for diluted methane, and the computational settings for Sandia flame D.

5.2.1 Experimental Setup

Sandia flame D is an unconfined turbulent methane-air jet flame at atmospheric conditions. The experimental setup consists of a cylindrical 7.2 mm nozzle for the main fuel flow and an annular pilot nozzle with inner and outer diameters of 7.7 mm and 18.2 mm, respectively. The burner is positioned 15 cm above a vertical wind tunnel, which produces a co-flow of air with a mean velocity of 0.9 m/s. The main fuel nozzle consists of 25% methane and 75% air by volume, which leaves the nozzle at a mean velocity of 49.6 m/s and a Reynolds

Table 5.1: Boundary conditions for Case 1 (Brookes-Moss flame), split by variable and boundary name

	Jet (patch)	Co-flow (patch)	Outlet (patch)	Wall (wall)
Velocity	Fixed ($20.3 \frac{m}{s}$)	Fixed ($0.514 \frac{m}{s}$)	Inlet-outlet ($0 \frac{m}{s}$)	No-slip
Pressure	Zero-gradient	Zero-gradient	Fixed (100,000 Pa)	Zero-gradient
Temperature	Fixed (290 K)	Fixed (290 K)	Zero-gradient	Zero-gradient
Species	Fixed (pure methane)	Fixed (pure air)	Zero-gradient	Zero-gradient
TKE	Fixed ($0.515 \frac{m^2}{s^2}$)	Fixed ($0.0003 \frac{m^2}{s^2}$)	Zero-gradient	kqR wall function
Dissipation of TKE	Fixed ($30 \frac{m^2}{s^3}$)	Fixed ($0.0004 \frac{m^2}{s^3}$)	Zero-gradient	epsilon wall function
Turbulent viscosity	Calculated	Calculated	Calculated	nutk wall function
Mean mixture fraction	Fixed (1)	Fixed (0)	Zero-gradient	Zero-gradient
Variance of mixture fraction	Fixed (0)	Fixed (0)	Zero-gradient	Zero-gradient

number of 22,400. The pilot consists of a burnt mixture of ethyne, hydrogen, air, carbon dioxide and nitrogen (N_2) with the same nominal enthalpy and mixture composition as the main fuel at an equivalence ratio of 0.77 ($Z = 0.27$). The mean pilot velocity is 11.4 m/s. The fuel, pilot, and co-flow streams exit the burner at temperatures of 294 K, 1880 K, and 291 K, respectively. The experimental uncertainties for $\tilde{T}$, $\widetilde{Y_{N_2}}$, $\widetilde{Y_{H_2O}}$, $\widetilde{Y_{CO_2}}$, $\widetilde{Y_{H_2}}$, $\widetilde{Y_{CO}}$, and $\widetilde{Y_{OH}}$ are 3%, 3%, 4%, 4%, 6–12%, 10–20%, and 10%, respectively [48, 49]. The measurement uncertainty for the calculated mixture fraction is reported to be about 15% [48].

5.2.2 ANN Preparation

Tabulated chemistry via TGLDM [13] with GRI-Mech 3.0 [70] for diluted methane is used to develop ANNs for Sandia flame D. The training data is augmented following the

procedure outlined in Section 4.2.1. The ANN weights and biases are determined from the training and testing procedure with optimized settings, as described in Section 4.3. The combined procedure took about 7 hours for all 50 ANNs on a single core of a 5 GHz processor. The weights and biases of each ANN are saved as text files to be used in ANN-CSE, which is described in Section 4.5.

5.2.3 Computational Details

Sandia flame D is simulated using both ANN-CSE and TGLDM-CSE using OpenFOAM v7 [74], with many of the computational details being identical to those in a previous CSE study [30]. The initialization for Sandia flame D (Case 2) is slightly different: ANN-CSE requires a small amount of $\widetilde{Y_{CO_2}}$ and $\widetilde{Y_{H_2O}}$ to initiate combustion, as in the Brookes-Moss flame (Case 1), however this was not required for TGLDM-CSE of Case 2. All other numerical settings are identical between ANN-CSE and TGLDM-CSE for Case 2, indicating that potential differences are still due to differences in the chemistry.

The computational domain is also cylindrical with a length of 720 mm and a radius of 151 mm. Following previous studies [30, 32, 33], the geometry is defined as a small angle wedge ($\theta = 5^\circ$) with one cell in the θ direction. The two-dimensional domain is justified by the fact that the velocity is axially dominated, and that there is no swirl in the flame. The mesh includes 51,000 non-uniform rectilinear cells that are concentrated near the main fuel nozzle in both the axial and radial directions. A total of 19 non-uniform CSE ensembles are defined in the axial direction. This arrangement has been used successfully in a previous CSE study of Sandia flame D [30]. Similar to Case 1, the mixture fraction space is discretized into 50 bins, with more bins concentrated near the stoichiometric mixture fraction of 0.353. The co-flow, pilot, and main fuel jet temperatures, along with the co-flow velocity, are defined as Dirichlet boundary conditions using the respective experimental values. Following previous studies [30, 32, 33], the k - ϵ turbulence model [54] with a modified coefficient of $C_{\epsilon 1} = 1.52$ is used for better agreement with the experimental data. The maximum Courant number is set to 0.9, and the flame is simulated to 2 s with data saved every 0.01 s. The spatial and temporal discretizations of each term are identical to Case 1, and each CSE case is run on a single core of a 5 GHz processor.

An illustration of the computational domain is provided in Figure 5.2. The outlet is located at the top of Figure 5.2a). A table of boundary conditions for each quantity and surface is provided in Table 5.2. As in Case 1, all quantities along the axis are specified as “empty”, and all quantities at the front and back faces are specified as “wedge”.

Table 5.2: Boundary conditions for Case 2 (Sandia flame D), split by variable and boundary name

	Inlet (patch)	Pilot (patch)	Co-flow (patch)	Outlet (patch)	Cylinder (patch)
Velocity	Fixed non-uniform list	Fixed non-uniform list	Fixed (0.9 $\frac{m}{s}$)	Inlet-outlet (0 $\frac{m}{s}$)	Zero-gradient
Pressure	Zero-gradient	Zero-gradient	Zero-gradient	Fixed (100,616 Pa)	Zero-gradient
Temperature	Fixed (294 K)	Fixed (1880 K)	Fixed (291 K)	Zero-gradient	Zero-gradient
Species	Fixed (25% CH ₄ , 75% air by volume)	Fixed (73% N ₂ , 10% CO ₂ , 10% H ₂ O, 5% O ₂ by mass fraction)	Fixed (pure air)	Zero-gradient	Zero-gradient
TKE	Fixed non-uniform list	Fixed non-uniform list	Fixed (0.00304 $\frac{m^2}{s^2}$)	Zero-gradient	Zero-gradient
Dissipation of TKE	Fixed non-uniform list	Fixed non-uniform list	Fixed (0.00764 $\frac{m^2}{s^3}$)	Zero-gradient	Zero-gradient
Turbulent viscosity	Calculated	Calculated	Calculated	Calculated	Calculated
Mean mixture fraction	Fixed (1)	Fixed (0.27)	Fixed (0)	Zero-gradient	Zero-gradient
Variance of mixture fraction	Fixed (0)	Fixed (0)	Fixed (0)	Zero-gradient	Zero-gradient




Figure 5.2: Illustration of computational domain for Case 2 (Sandia flame D): a) temperature contour with CSE ensemble boundaries shown in black and axial locations of results shown in red, b) mesh details near the main fuel inlet

5.3 Results

In this section, results for the mixing field, conditional and Favre averages, and computational speed and memory are presented and analyzed from both ANN-CSE and TGLDM-CSE for both fuels. Results are examined at three axial locations in each flame, characterized by the normalized distance $\frac{x}{D}$, where x is the absolute distance from the main fuel inlet and D is the diameter of the main fuel inlet.

5.3.1 Mixing Field

The mean mixture fraction ($\tilde{Z}$) and its variance ($\widetilde{Z''^2}$) are solved at each CFD cell. These quantities are used in the CSE model to construct the PDF for each cell. The PDF is then used to determine the Favre averages of the conditional quantities by applying Equation 2.47 to each cell. Thus, it is important that the turbulent mixing field is well represented




Figure 5.3: Radial profiles of the mean (top) and RMS (bottom) of the mixture fraction at various locations in the Brookes-Moss flame from ANN-CSE (solid line), TGLDM-CSE (dashed line), and experimental data (symbols)

by ANN-CSE to ensure that the conditional averages are accurate, in particular $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$, as these are input quantities for the present ANNs.

With this in mind, the radial $\tilde{Z}$ and $Z''_{rms} (= \sqrt{\widetilde{Z''^2}})$ profiles are examined in Figures 5.3 and 5.4 for Cases 1 and 2, respectively. As can be seen, the $\tilde{Z}$ and Z''_{rms} profiles from ANN-CSE and TGLDM-CSE are very similar at each location; the maximum absolute deviations are on the order of 10^{-4} to 10^{-3} for both $\tilde{Z}$ and Z''_{rms} . This indicates that the resulting PDFs are similar for a given CFD cell within each CSE ensemble. Therefore, any major differences between ANN-CSE and TGLDM-CSE can be attributed to the chemistry representation. In addition, the experimental trends for $\tilde{Z}$ are captured with both formulations of CSE, however $\tilde{Z}$ is consistently underpredicted in the Brookes-Moss flame and overpredicted in Sandia flame D. In addition, the experimental trend for Z''_{rms} is captured in Sandia flame D, with some discrepancies in the shape of the profile. These deviations have been observed in previous CSE studies [28, 30–33].




Figure 5.4: Radial profiles of the mean (top) and RMS (bottom) of the mixture fraction at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line), and experimental data (symbols)

5.3.2 ANN Input Variables: Conditional CO_2 and H_2O Mass Fractions

The $\widetilde{Y_{\text{CO}_2}|\zeta}$ and $\widetilde{Y_{\text{H}_2\text{O}}|\zeta}$ profiles at three locations are shown in Figures 5.5 and 5.6 for both flames respectively. The stoichiometric mixture fraction is indicated by a vertical dashed line for each plot. The $\widetilde{Y_{\text{CO}_2}|\zeta}$ and $\widetilde{Y_{\text{H}_2\text{O}}|\zeta}$ profiles from ANN-CSE and TGLDM-CSE are very close to each other, overlapping for most parts. The maximum absolute deviations are on the order of 10^{-4} to 10^{-2} for both $\widetilde{Y_{\text{CO}_2}|\zeta}$ and $\widetilde{Y_{\text{H}_2\text{O}}|\zeta}$ in both flames. The largest deviations occur in the fuel-rich region at each location, with the peak absolute error generally increasing downstream. However, the PDF is either low or near-zero in the fuel-rich region, as indicated by the $\widetilde{Z}$ and Z'_{rms} profiles in Figures 5.3 and 5.4. Further, the differences between ANN-CSE and TGLDM-CSE are less than 2% in the Brookes-Moss flame and 4% in Sandia flame D near the stoichiometric mixture fraction. The experimental trends in Sandia flame D are captured from both CSE formulations. However, CSE tends to underpredict the experimental $\widetilde{Y_{\text{CO}_2}|\zeta}$ and $\widetilde{Y_{\text{H}_2\text{O}}|\zeta}$ profiles, with the exception of $\widetilde{Y_{\text{H}_2\text{O}}|\zeta}$




Figure 5.5: Conditional CO_2 (top) and H_2O (bottom) mass fraction profiles at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line). The stoichiometric mixture fraction is indicated by a vertical dashed line

in the fuel-rich region at $\frac{x}{d} = 30$. Similar observations are obtained in previous CSE studies of Sandia flame D [30, 32, 33].

5.3.3 Other Conditional Averages

The conditional minor species mass fractions ($\widetilde{Y_{CO}}|\zeta$, $\widetilde{Y_{H_2}}|\zeta$, and $\widetilde{Y_{OH}}|\zeta$) are displayed in Figures 5.7 and 5.8 for both flames respectively. The minor species mass fractions are generally well represented by ANN-CSE at each location, which is indicated by similar trends to TGLDM-CSE. The maximum absolute errors in both flames are on the order of 10^{-3} to 10^{-2} for $\widetilde{Y_{CO}}|\zeta$, and 10^{-4} to 10^{-3} for $\widetilde{Y_{H_2}}|\zeta$ and $\widetilde{Y_{OH}}|\zeta$, with the deviations generally increasing downstream. Most of the deviations from TGLDM-CSE occur at regions where $\widetilde{Y_{CO_2}}|\zeta$ and $\widetilde{Y_{H_2O}}|\zeta$ differ (i.e., in the fuel-rich region at each location). An exception to this is $\widetilde{Y_{OH}}|\zeta$, which is consistently larger with ANN-CSE in the fuel-lean region. This may be due to the sensitivity of $\widetilde{Y_{OH}}$, however a minimal impact is expected




Figure 5.6: Conditional CO_2 (top) and H_2O (bottom) mass fraction profiles at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols). The stoichiometric mixture fraction is indicated by a vertical dashed line

on the flame behaviour as the overall magnitude of $\widetilde{Y_{OH}|\zeta}$ is small. Larger differences with TGLDM-CSE are observed for all minor species at $\frac{x}{d} = 87.5$ in the Brookes-Moss flame and $\frac{x}{d} = 60$ in Sandia flame D. The experimental trends in Sandia flame D are captured by both CSE formulations, however the profiles are noticeably underpredicted at $\frac{x}{d} = 7.5$ and overpredicted at $\frac{x}{d} = 30$. Similar trends are observed in previous CSE studies [30, 32, 33] and are attributed to interpolation in the TGLDM tables [30].

The conditional major species source terms ($\widetilde{\dot{\omega}_{\text{CO}_2}|\zeta}$, $\widetilde{\dot{\omega}_{\text{H}_2\text{O}}|\zeta}$, $\widetilde{\dot{\omega}_{\text{CH}_4}|\zeta}$, $\widetilde{\dot{\omega}_{\text{O}_2}|\zeta}$) and enthalpy source term ($\widetilde{\dot{\omega}_T|\zeta}$) are presented in Figures 5.9 and 5.10 for both flames respectively. The conditional reaction rates are more difficult to predict than the conditional species mass fractions due to the high non-linearity. However, the trends of each quantity are captured by ANN-CSE. The maximum absolute deviations for the Brookes-Moss flame are consistently on the order of 10^0 for $\widetilde{\dot{\omega}_{\text{CH}_4}|\zeta}$, 10^1 for $\widetilde{\dot{\omega}_{\text{CO}_2}|\zeta}$, $\widetilde{\dot{\omega}_{\text{H}_2\text{O}}|\zeta}$ and $\widetilde{\dot{\omega}_{\text{O}_2}|\zeta}$, and 10^8 for $\widetilde{\dot{\omega}_T|\zeta}$. The deviations in Sandia flame D are larger, being on the order of 10^1 to 10^2 for each species and 10^8 to 10^9 for $\widetilde{\dot{\omega}_T|\zeta}$. As expected, the largest differences between




Figure 5.7: Conditional CO (top), H₂ (middle) and OH (bottom) mass fraction profiles at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line). The stoichiometric mixture fraction is indicated by a vertical dashed line

ANN-CSE and TGLDM-CSE occur in the fuel-rich region, where $\widetilde{Y_{CO_2}}|\zeta$ and $\widetilde{Y_{H_2O}}|\zeta$ between ANN-CSE and TGLDM-CSE differ the most. However, there are more prominent deviations towards the fuel-lean region for each quantity at each location. The conditional reaction rate profiles from ANN-CSE become less accurate and less smooth downstream, however the reaction rate magnitudes decrease downstream as the fuel is consumed. The error magnitudes are generally consistent in the Brookes-Moss flame, however they decrease downstream in Sandia flame D for each source term.

The ANN error is expected to accumulate due to prediction errors at each time step. To analyze this, the conditional averages from ANN-CSE and TGLDM-CSE in both cases were saved every 250 time steps, and the absolute errors between ANN-CSE and TGLDM-




Figure 5.8: Conditional CO (top), H₂ (middle) and OH (bottom) mass fraction profiles at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols). The stoichiometric mixture fraction is indicated by a vertical dashed line

CSE were calculated over time for each variable. After a brief initialization period, it was observed that the absolute errors approach a steady value on the same order of magnitude as the maximum error observed for a given variable, with some fluctuations about this mean value. The results are consistent with mixture fraction and ensemble locations. This shows that the error accumulation in ANN-CSE is stable. Note that additional errors from integral inversion are present in both formulations of CSE. Another test was performed outside of CSE, in which the conditional averages were evolved from time zero to unity using ANN and TGLDM with an assumed fixed time step of 10^{-5} seconds for both cases. The inputs for the subsequent time step are calculated using the predicted reaction rates (CO₂ and H₂O) and the time step value. The ANN inputs were set to those from TGLDM




Figure 5.9: Conditional CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate profiles at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line). The stoichiometric mixture fraction is indicated by a vertical dashed line




Figure 5.10: Conditional CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate profiles at various locations in Sandia flame D from ANN-CSE (solid line) and TGLDM-CSE (dashed line). The stoichiometric mixture fraction is indicated by a vertical dashed line

to allow for a pure evaluation of the prediction error. In this test, the results in either ANN or TGLDM tend to approach the equilibrium point in the manifold. The absolute error between ANN and TGLDM tends to start near zero and approach an asymptotic limit; this limit is on the same order of magnitude as observed for the maximum absolute errors observed in CSE. For some ensembles and mixture fractions, however, the absolute errors remain small and consistent. These results show that the ANN prediction errors accumulate to a limit, which is observed in the conditional profiles at steady state.

5.3.4 Favre Averages

The radial profiles of $\widetilde{Y_{CO_2}}$, $\widetilde{Y_{H_2O}}$ and $\widetilde{T}$ are shown in Figures 5.11 and 5.12 for both flames respectively. These quantities are solved by OpenFOAM, however they are influenced by the ANN predictions. The general trends are similar between ANN-CSE and TGLDM-CSE at each location, with absolute errors on the order of 10^{-4} to 10^{-3} for $\widetilde{Y_{CO_2}}$ and $\widetilde{Y_{H_2O}}$ in both flames. These results impact the $\widetilde{T}$ profile, which has an absolute error on the order of 10^1 for both flames. Peak differences between ANN-CSE and TGLDM-CSE occur towards the centerline at each location, where the flame is expected to be fuel-rich. An exception is Sandia flame D at $\frac{x}{D} = 60$, which occurs in a fuel-lean region. The Favre absolute errors are often lower than the respective conditional absolute errors, however the peak absolute error tends to increase downstream in both flames. The experimental trends are captured with both formulations of CSE. The Brookes-Moss $\widetilde{T}$ profiles are underpredicted at $\frac{x}{d} = 87.5$ and wider at $\frac{x}{d} = 37.5$ and $\frac{x}{d} = 62.5$, as in previous CSE studies [28, 31]. In Sandia flame D, the experimental peaks of these quantities are underpredicted by ANN-CSE near the base of the flame, however the entire profile becomes overpredicted by ANN-CSE for these quantities farther downstream. These trends are also observed in previous CSE calculations [30, 32, 33].

The radial profiles of the minor species mass fractions are shown in Figures 5.13 and 5.14 for both flames respectively. The general trends of each quantity are captured by ANN-CSE at each location. In both flames, the peak deviations of $\widetilde{Y_{CO}}$ and $\widetilde{Y_{H_2}}$ occur near the centerline, where the combustion is fuel-rich, and are on the order of 10^{-4} to 10^{-3} and 10^{-5} to 10^{-4} , respectively. The peak value of $\widetilde{Y_{OH}}$ is overestimated by 10^{-5} to 10^{-4} in both flames, however this occurs in regions where the combustion is fuel-lean. These differences are expected based on the trends in the respective conditional averages, however the differences are reduced when integrated with the PDF. In general, the minor species mass fractions are more accurate towards the base of the flame, while the fit between ANN-CSE and TGLDM-CSE becomes less accurate downstream. This is especially noticeable




Figure 5.11: Radial Favre averages of the CO_2 (top) and H_2O (middle) mass fractions and temperature (bottom) at various locations in the Brookes-Moss flame from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols)

for Sandia flame D at $\frac{x}{d} = 60$. However, the magnitude of the minor species mass fractions generally decreases downstream, so the absolute errors remain small. The experimental trends are also captured well in Sandia flame D, with the aforementioned differences at $\frac{x}{D} = 60$. These trends are consistent with previous CSE studies [30, 32, 33].

The radial Favre averages of the source terms are seen in Figures 5.15 and 5.16 for each flame respectively. The trends are generally well represented by ANN-CSE across all locations. The exception is $\widetilde{\omega}_{\text{CH}_4}$ downstream, where minor amounts of production are predicted by ANN-CSE in both flames. However, $\widetilde{Y}_{\text{CH}_4}$ is small at these locations (not shown for brevity) since most of the fuel is consumed, therefore this has minimal impact on the flames. The species reaction rate deviations are on the order of 10^{-1} in the Brookes-




Figure 5.12: Radial Favre averages of the CO_2 (top) and H_2O (middle) mass fractions and temperature (bottom) at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols)

Moss flame, and 10^{-1} to 10^0 for Sandia flame D. The maximum $\widetilde{\omega_T}|\zeta$ deviations are on the order of 10^6 for the Brookes-Moss flame, and 10^7 for Sandia flame D. These generally occur at the peaks of each source term, which occur in the fuel-rich region. For both flames, the Favre reaction rate discrepancies with those from TGLDM-CSE are lower than what is seen in conditional space. Again, larger differences are seen farther downstream.

A significant number of deviations occur in the fuel-rich regions for most quantities in both flames. This is likely due to the fuel-rich manifolds being less smooth and ‘orderly’ than the fuel-lean and stoichiometric manifolds, which is inherent in TGLDM and other manifold methods. However, the PDF for these fuel-rich mixtures is small or zero, therefore creating a minimal effect on the Favre averaged quantities.




Figure 5.13: Radial Favre averages of the CO (top), H₂ (middle) and OH (bottom) mass fractions at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line)

The deviations in the Favre averages between ANN-CSE and TGLDM-CSE occur in regions where there is a corresponding deviation in the conditional averages for a given quantity. Ultimately, better predictions of the reaction rates at each time step would minimize differences between ANN-CSE and TGLDM-CSE, as these predictions impact the Favre averages solved by OpenFOAM and hence the inputs to the ANNs (i.e., $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$). This could be achieved by refining the ζ grid, potentially leading to smoother profiles.

The downstream results are less accurate than those near the base of the flame in terms of the fit between the ANN-CSE and TGLDM-CSE profiles. This may be due to the difference in inputs, arising from previous reaction rate errors. However, the magnitudes of




Figure 5.14: Radial Favre averages of the CO (top), H₂ (middle) and OH (bottom) mass fractions at various locations in Sandia flame D from ANN-CSE (solid line), TGLDM-CSE (dashed line) and experimental data (symbols)

the predicted quantities are lower downstream, resulting in low absolute errors. It should also be noted that there are fewer CFD cells downstream. Increasing the number of CFD cells or reducing the size of the CSE ensembles downstream could lead to more accurate results for ANN-CSE. However, since the same computational settings are used between CSE formulations, these deviations are due to the chemistry implementation.

The majority of CFD cells have low values of $\tilde{Z}$, where the conditional averages from ANN-CSE are similar to those from TGLDM-CSE for most quantities at each location. This can be seen in the Favre averaged profiles for ANN-CSE, as the general trends for each quantity are similar between ANN-CSE and TGLDM-CSE at each location. This indicates that the errors in the conditional averages are often minimized in the Favre averages.




Figure 5.15: Radial Favre averages of the CO₂, H₂O, CH₄, O₂ reaction rate and heat release rate at various locations in the Brookes-Moss flame from ANN-CSE (solid line) and TGLDM-CSE (dashed line)




Figure 5.16: Radial Favre averages of the CO₂, H₂O, CH₄, O₂ reaction rate and heat release rate at various locations in Sandia flame D from ANN-CSE (solid line) and TGLDM-CSE (dashed line)

Table 5.3: Memory usage, CSE routine execution time, and simulation wall-clock time for ANN-CSE and TGLDM-CSE for both flames

	Brookes-Moss flame		Sandia flame D	
	ANN-CSE	TGLDM-CSE	ANN-CSE	TGLDM-CSE
Memory [MB]	216	285	183	242
Execution time [s]	0.80	2.10	0.60	1.55
Wall-clock time [hh:mm:ss]	108:46:24	262:05:19	26:04:36	73:05:30

5.3.5 Computational Efficiency

The computational efficiency of ANN-CSE and TGLDM-CSE is assessed by examining memory usage, execution time of the CSE routine, and overall wall-clock time. ANN-CSE requires less memory than TGLDM-CSE and is faster in terms of execution time and wall-clock time for both flames, as shown in Table 5.3. The execution time is defined as the time to run one time step in each flame. The quoted times are taken as an average of the last 20 time steps. At this stage, both flames are steady and the CSE routine is performed in all ensembles. The memory is fairly consistent during the simulation in each case and is quoted as a rounded average from the system monitor. The wall-clock time is quoted as the clock time listed at the end of each simulation from OpenFOAM.

ANN-CSE requires approximately 24% less memory than TGLDM-CSE for both flames. Similarly, the execution time is reduced by approximately 62% for both flames. The memory and time savings are due to the ANN calculations being performed faster than searching through and interpolating the TGLDM tables. These savings are based on the CSE solver and are independent of the computational settings, however the Brookes-Moss flame simulations require more memory and time than the respective Sandia flame D simulations since the Brookes-Moss flame has more CFD cells. Finally, the wall-clock time reductions are about 58% for the Brookes-Moss flame and about 64% for Sandia flame D. The absolute time wall-clock time savings are greater for the Brookes-Moss flame, which runs longer than Sandia flame D, due to the execution time savings in each time step. It should be noted that the CSE routine itself is not computationally expensive, requiring only 4% of the processor power in each case. However, the relative computational requirement of the chemistry aspect in CSE is reduced using ANN instead of TGLDM.

5.4 Summary

The objective of this chapter was to assess the influence of chemistry representation in CSE in the context of turbulent non-premixed jet flames. In this chapter, ANN-CSE is applied to two turbulent non-premixed jet flames with pure and diluted methane. Both flames have been previously simulated using CSE with tabulated chemistry. The experimental and computational details, along with the ANN development procedure, are described for both flames. Simulation results from ANN-CSE and TGLDM-CSE are then presented and analyzed. Since the computational details are identical for each CSE formulation, the results in a given flame are wholly dependent on the chemistry representation in CSE.

The conditional averages for each quantity are predicted fairly well by ANN-CSE in both flames at various locations. Most of the deviations occur due to different values of $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$ in ANN-CSE compared to TGLDM-CSE, and the accuracy tends to decrease downstream where the temperature becomes large. However, the error magnitude is low throughout the domain, and the majority of these errors are minimized when integrated with the PDF. This can be seen in the trends of the mixing field and the Favre averages for each quantity, as the profiles from ANN-CSE are nearly identical to those from TGLDM-CSE for both flames at various locations. ANN-CSE is also shown to require about 24% less memory and is at least 58% faster than TGLDM-CSE for both flames.

The results in this chapter indicate that ANN-CSE can accurately simulate different turbulent non-premixed jet flames while requiring less computational cost than TGLDM-CSE. However, it is expected that the results from ANN-CSE are similar since the tabulated chemistry data used in TGLDM-CSE is used to train the ANNs in ANN-CSE. The general accuracy of ANN-CSE is also dependent on the quality of the training data used. The next chapter focuses on ANN-CSE simulations of turbulent non-premixed jet flames using data obtained from an alternative source.

Chapter 6

ANN-CSE for Sandia Flame D Using Multiple Sources of Data

The objective of this chapter is to assess the influence of the source of data in the ANN development process and the simulation of Sandia flame D using ANN-CSE. ANN development via tabulated chemistry and sampling of CSE with direct chemistry (DI-CSE) is considered. Since a reduced mechanism is used for the direct chemistry samples, the TGLDM-CSE reference case is redone with the new mechanism. The ANN development process is described for both data sources. Results for the mixing field, conditional and Favre averages, and computational cost are then presented and compared to those from TGLDM-CSE with the reduced mechanism. Note that the DI-CSE case is run on a separate code and is not plotted for clarity. A summary of findings and recommendations for this work are provided at the end of this chapter.

Note that the experimental setup of Sandia flame D is described in Section [5.2.1](#).

6.1 ANN Preparation

The ANN preparation procedure for TGLDM is similar to that described in Section [5.2.2](#), however Smooke’s mechanism is used instead of GRI-Mech 3.0 and the total training time is about 2.5 hours on a single core of a 5 GHz processor. To prepare the ANNs with DI samples, first a simulation of Sandia flame D with DI-CSE is carried out from steady state. This simulation uses Smooke’s mechanism for chemistry and is based on the work by Mahdipour and Devaud [[33](#)]. The input velocity is then perturbed by a cosine wave of 40

$\left[\frac{m}{s}\right]$ amplitude and 100 [Hz] frequency. Conditional averages are extracted every 250 time steps. These samples are augmented following the procedure outlined in Section 4.2.2. The training and testing process is similar, taking a total of about 3 hours on a single core of a 5 GHz processor. The usage of ANNs in CSE is identical for both cases, and is described in Section 4.5.

6.2 Computational Details

The DI case comes from a previous publication [33], which uses different computational settings than the cases in Chapter 5. Since the DI samples are generated from the previous work, the computational settings in TGLDM-CSE and both ANN-CSE cases are modified to match the previous work as closely as possible. However, it was not possible to use identical settings to DI-CSE due to the differences in the codes.

OpenFOAM v7 [74] is used to simulate each case. In both ANN-CSE cases, the domain is initialized with a small amount of $\widetilde{Y_{CO_2}}$ (0.06) and $\widetilde{Y_{H_2O}}$ (0.1), with a balance of N_2 . The TGLDM-CSE and DI-CSE cases do not require initialization; the entire domain is filled with N_2 . All four cases use second-order Gaussian integration for discretizing the spatial terms. The temporal terms in TGLDM-CSE and both ANN-CSE cases are discretized using an implicit Euler scheme. However, the DI-CSE case uses the Crank-Nicolson scheme with a blend factor of 0.1, making it similar but not identical to the Euler scheme. In addition, TGLDM-CSE and both ANN-CSE cases are ran on a single core of a 5 GHz processor, whereas the DI-CSE case is ran on eight cores of the same processor. These differences are expected to have minimal impact on the chemistry.

The computational domain is cylindrical with a length of 700 mm and a radius of 150 mm. The geometry is defined as a small angle wedge ($\theta = 5^\circ$) with one cell in the θ direction, for the same reasons as described in Chapter 5. The CFD mesh consists of 43,580 non-uniform rectangular cells that are concentrated near the fuel inlet in the axial and radial directions. A total of 25 non-uniform CSE ensembles are defined in the axial direction, being concentrated in the middle of the domain where the flame is expected to burn. Note that the parallel processing in DI-CSE results in a total of 32 CSE ensembles due to partitioning the domain. The mixture fraction space is discretized into 50 bins as before, however the ζ distribution in all four cases is linear, ranging from $\zeta = (0.01, 0.99)$ with intervals of $\zeta = 0.02$. The co-flow, pilot, and main fuel jet temperatures, along with the co-flow velocity, are defined as Dirichlet boundary conditions using the respective experimental values. The $k-\epsilon$ turbulence model is used for all four cases, with a modified




Figure 6.1: Illustration of computational domain for Sandia flame D in Chapter 6: a) temperature contour with CSE ensemble boundaries shown in black, axial locations of results shown in red, and parallel processing domain boundaries in grey, b) mesh details near the main fuel inlet

coefficient of $C_{\epsilon 1} = 1.52$ being used. For all four cases, the maximum Courant number is set to 0.6, and the flame is simulated to 2.64 s with data saved every 0.02 s.

An illustration of the computational domain is provided in Figure 6.1. The outlet is located at the top of Figure 6.1a). Note that to match the geometry in [33], the inlet and pilot tubes are physically modelled and have walls at the outer edges, which are not annotated in Figure 6.1b). A table of boundary conditions for each quantity and surface is provided in Table 6.1. As in Chapter 5, all quantities along the axis are specified as “empty”, and all quantities at the front and back faces are specified as “wedge”.

6.3 Results

Results from the ANN-CSE simulations of Sandia flame D using either TGLDM or DI training data are provided in this section. The mixing field, conditional and Favre averages,

Table 6.1: Boundary conditions for Sandia flame D (Chapter 6), split by variable and boundary name

	Inlet (patch)	Pilot (patch)	Co-flow (patch)	Outlet (patch)	Walls (wall)
Velocity	Fixed (49.6 $\frac{m}{s}$)	Fixed (11.4 $\frac{m}{s}$)	Fixed (0.9 $\frac{m}{s}$)	pressure-Inlet-Outlet-Velocity	Zero-gradient (outer), no-slip (tube)
Pressure	Zero-gradient	Zero-gradient	Zero-gradient	Total pressure	Zero-gradient
Temperature	Fixed (294 K)	Fixed (1880 K)	Fixed (291 K)	Zero-gradient	Zero-gradient
Species	Fixed (25% CH ₄ , 75% air by volume)	Fixed (73% N ₂ , 10% CO ₂ , 10% H ₂ O, 5% O ₂ by mass fraction)	Fixed (pure air)	Zero-gradient	Zero-gradient
TKE	TKE-inlet (5% intensity)	TKE-inlet (5% intensity)	TKE-inlet (5% intensity)	Zero-gradient	kqR wall function
Dissipation of TKE	Inlet-outlet, mixing length of 0.0003 m	Inlet-outlet, mixing length of 0.000274 m	Inlet-outlet, mixing length of 0.0015 m	Zero-gradient	epsilon wall function
Turbulent viscosity	Calculated	Calculated	Calculated	Calculated	nutk wall function
Mean mixture fraction	Fixed (1)	Fixed (0.27)	Fixed (0)	Zero-gradient	Zero-gradient
Variance of mixture fraction	Fixed (0.0004)	Fixed (0)	Fixed (0)	Zero-gradient	Zero-gradient

and computational speed and memory are presented and analyzed relative to TGLDM-CSE. Experimental data are also included when available. Note that the same three locations from Chapter 5 are considered here.

6.3.1 Mixing Field

The mean and root mean square (RMS) of the mixture fraction are analyzed in Figure 6.2 at the three locations in Sandia flame D. In each case, the experimental trends of $\tilde{Z}$ and Z''_{rms} are captured at each location, with both quantities being generally overpredicted as seen in Chapter 5 and [30, 32, 33]. The ANN-CSE cases have good agreement to TGLDM-CSE at each location, with maximum absolute errors on the order of 10^{-2} for $\tilde{Z}$ and 10^{-3} for Z''_{rms} , indicating that ANN-CSE can replicate the general PDF shapes. These errors increase downstream and are larger for DI-ANN-CSE. The two ANN-CSE cases are also generally similar to each other, indicating that ANN-CSE can represent the mixing field with either data source. Notable deviations between the two ANN-CSE cases occur downstream at $\frac{x}{D} = 60$, with DI-ANN-CSE being closer to the experimental results. This may be due to differences in the species mass fractions arising in the two ANN-CSE cases, itself caused by different reaction rate predictions, which will be observed in Section 6.3.4.

6.3.2 ANN Input Variables: Conditional CO₂ and H₂O Mass Fractions

The conditional CO₂ and H₂O mass fraction profiles are now assessed at the same three locations for each case. The profiles are shown in Figure 6.3. The experimental trends are again captured by each formulation of CSE at each location, however the conditional quantities are generally underpredicted by the TGLDM-based cases and overpredicted by DI-ANN-CSE. Notable deviations from TGLDM-CSE occur in the fuel-rich region at $\frac{x}{D} = 60$, however the PDF is near-zero at this location, thus mitigating this error. The fit between the two ANN-CSE cases and TGLDM-CSE is quite good everywhere else. The maximum absolute errors are on the order of 10^{-2} to 10^{-1} for both quantities in TGLDM-ANN-CSE, and are consistently on the order of 10^{-1} for both quantities in DI-ANN-CSE. Differences between the two TGLDM-based cases increase downstream and occur near the peaks and in the fuel-rich region for both quantities, as observed in Chapter 5. Differences between DI-ANN-CSE and TGLDM-CSE also increase, however they occur in the fuel-lean region for $\widetilde{Y_{CO_2}|\zeta}$ and the fuel-rich region for both quantities. The DI-ANN-CSE profiles tend to follow the trends of the TGLDM-ANN-CSE profiles, however the profile for $\widetilde{Y_{CO_2}|\zeta}$




Figure 6.2: Radial profiles of the mean (top) and RMS (bottom) of the mixture fraction at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue), with experimental data (symbols)

at $\frac{x}{D} = 60$ becomes non-linear in the fuel-rich region. This is attributed to incomplete statistics for the integral inversion, as there are very few cells past $\zeta = 0.3$ at this location (refer to Figure 6.2).

6.3.3 Other Conditional Averages

The conditional profiles of the minor species (CO, H₂, OH) mass fractions from each CSE formulation are assessed in Figure 6.4 at the same three locations. The experimental trends are generally captured for each quantity, but are less accurate than the ANN input variables. TGLDM-ANN-CSE generally captures the peaks but underpredicts the fuel-lean and fuel-rich regions, whereas DI-ANN-CSE is more broad but underpredicts the quantities everywhere except at $\frac{x}{D} = 60$. DI-ANN-CSE is generally less smooth than TGLDM-ANN-CSE, especially in the fuel-rich region. This is due to the high non-linearity of the output variables with $\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$ in the DI training data. There is a prominent ‘spike’ around $\zeta = 0.7$ at $\frac{x}{D} = 30$ that affects all three minor species; this is likely due to an




Figure 6.3: Conditional CO_2 (top) and H_2O (bottom) mass fraction profiles at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue) with experimental data (symbols). The stoichiometric mixture fraction is indicated by a vertical dashed line

error in preparing the samples, as the inputs and ζ field are smooth at this location. The conditional minor species mass fractions aside from $\widetilde{Y_{OH}|\zeta}$ are also much smaller in DI; this comes directly from the samples and follows what is observed in [33]. Despite these shortcomings, the two ANN-CSE cases follow the trends in TGLDM-CSE, with maximum absolute errors on the order of 10^{-3} to 10^{-2} in $\widetilde{Y_{CO}|\zeta}$ and 10^{-4} to 10^{-3} in $\widetilde{Y_{H_2}|\zeta}$ and $\widetilde{Y_{OH}|\zeta}$. Once again, the absolute errors increase downstream for each species, and are larger in DI-ANN-CSE than TGLDM-ANN-CSE.

The conditional major species reaction rates (CO_2 , H_2O , CH_4 and O_2), along with the enthalpy source term ($\widetilde{\dot{\omega}_T|\zeta}$), are assessed in Figure 6.5. The reaction rates are quite non-linear, as expected from the observations in Chapter 5, however the general trends are captured by both ANN-CSE cases. The maximum absolute errors are on the order of 10^1 to 10^2 for each species in both cases. $\widetilde{\dot{\omega}_T|\zeta}$ has maximum absolute errors on the order of 10^8 in TGLDM-ANN-CSE and 10^9 in DI-ANN-CSE. Interestingly, the maximum absolute errors decrease downstream for each quantity in DI-ANN-CSE, whereas the errors are somewhat




Figure 6.4: Conditional CO (top), H₂ (middle) and OH (bottom) mass fraction profiles at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue), with experimental data (symbols). The stoichiometric mixture fraction is indicated by a vertical dashed line

consistent in TGLDM-ANN-CSE. DI-ANN-CSE tends to predict larger reaction rates in the fuel-lean region than TGLDM-ANN-CSE, especially at $\frac{x}{D} = 7.5$. This is likely caused by the differences in $\widetilde{Y_{CO_2}}|\zeta$ and $\widetilde{Y_{H_2O}}|\zeta$ in the fuel-lean region, which are larger in DI-ANN-CSE (refer to Figure 6.3). It is expected that these differences will create some discrepancies between the Favre averages in the two ANN-CSE cases, as the PDF is non-zero in fuel-lean conditions.




Figure 6.5: Conditional CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate profiles at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue). The stoichiometric mixture fraction is indicated by a vertical dashed line

6.3.4 Favre Averages

The Favre averages of temperature and the major species (CO_2 and H_2O) mass fractions are now assessed in Figure 6.6. These quantities are solved by OpenFOAM, however they represent the flame structure and impact the ANN inputs. It can be seen that the experimental trends of each quantity are captured by each CSE formulations, however the quantities tend to be underpredicted near the inlet and overpredicted downstream, similar to what is observed in Chapter 5. The ANN-CSE cases are nearly identical to TGLDM-CSE. Maximum absolute errors for $\widetilde{Y_{\text{CO}_2}}$ and $\widetilde{Y_{\text{H}_2\text{O}}}$ are on the order of 10^{-3} to 10^{-2} , with the errors generally increasing downstream. The errors in the Favre averages are generally reduced by a factor of at least 2 relative to the conditional averages, similar to what was observed in Section 5.3. These errors impact the temperature profiles, with maximum absolute errors on the order of 10^1 for TGLDM-ANN-CSE and 10^2 for DI-ANN-CSE. The errors increase downstream but remain well below the respective conditional errors. $\widetilde{Y_{\text{CO}_2}}$ and $\widetilde{Y_{\text{H}_2\text{O}}}$ from DI-ANN-CSE tend to have larger peaks and wider profiles than TGLDM-ANN-CSE, especially at $\frac{x}{D} = 30$. This is directly caused by the larger reaction rate predictions from DI-ANN-CSE in the fuel-lean region (refer to Figure 6.5). The $\widetilde{T}$ profiles, however, generally have better agreement since the lower $\widetilde{Y_{\text{CO}}}$ predictions from DI-ANN-CSE, which will be observed in Section 6.7, counteract the larger $\widetilde{Y_{\text{CO}_2}}$ and $\widetilde{Y_{\text{H}_2\text{O}}}$ predictions. One exception is $\widetilde{T}$ at $\frac{x}{D} = 60$, which is significantly lower in DI-ANN-CSE because of the mass fraction values.

The Favre averages of the minor species mass fractions are analyzed for the same three locations in Figure 6.7. The experimental trends are captured with each CSE formulation, however each quantity tends to be underpredicted aside from the $\frac{x}{D} = 60$ location, where each quantity is generally overpredicted. These trends follow what is observed in Chapter 5 and in [30, 32, 33]. One exception is the $\widetilde{Y_{\text{OH}}}$ profile from DI-ANN-CSE at $\frac{x}{D} = 30$, which is unrealistically large and increasing towards the centerline. This is caused by the ‘spike’ observed in the fuel-rich region of $\widetilde{Y_{\text{OH}}|\zeta}$ (refer to Figure 6.4), which stems from an error in the sample data. The two ANN-CSE cases differ significantly for the minor species, contrary to what was observed in the major species. $\widetilde{Y_{\text{CO}}}$ and $\widetilde{Y_{\text{H}_2}}$ are generally lower with the DI-based cases, whereas $\widetilde{Y_{\text{OH}}}$ is larger instead. These trends follow what is observed in the conditional averages and is attributed to the differences of these quantities in the TGLDM and DI sample spaces. The TGLDM-based cases are similar to each other, with maximum absolute errors on the order of 10^{-3} for $\widetilde{Y_{\text{CO}}}$ and 10^{-4} to 10^{-3} for $\widetilde{Y_{\text{H}_2}}$ and $\widetilde{Y_{\text{OH}}}$. Larger differences are seen with the DI-ANN-CSE results, with maximum errors on the order of 10^{-2} for $\widetilde{Y_{\text{CO}}}$ and 10^{-4} to 10^{-3} for $\widetilde{Y_{\text{H}_2}}$ and $\widetilde{Y_{\text{OH}}}$. In both ANN-CSE cases,




Figure 6.6: Radial Favre averages of the CO₂ (top) and H₂O (middle) mass fractions and temperature (bottom) at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue), with experimental data (symbols)

the absolute errors generally increase downstream and are reduced in the Favre averages, as expected from Chapter 5.3.

The Favre averages of the conditional source terms are now presented in Figure 6.8. The fit between both ANN-CSE formulations and TGLDM-CSE is quite good, with maximum absolute deviations on the order of 10^{-1} to 10^1 for each species and 10^7 to 10^8 for $\widetilde{\omega_T|\zeta}$. The Favre-averaged source term profiles are much smoother and have lower errors than what is seen in conditional space. The trends are predicted correctly for each quantity at each location, contrary to what was observed in Chapter 5 with GRI-Mech 3.0. The errors between both ANN-CSE cases and TGLDM-CSE generally decrease downstream, which is also in contrast to the trends in Chapter 5 and for the species in this case.




Figure 6.7: Radial Favre averages of the CO (top), H₂ (middle) and OH (bottom) mass fractions at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue), with experimental data (symbols)

The two ANN-CSE cases are generally similar to each other, with DI-ANN-CSE correctly predicting the general location of the peaks for each quantity, although the peak value itself is underestimated near the inlet and overestimated downstream. The differences in the reaction rates lead to slightly different flame shapes in the two ANN-CSE cases.

For most quantities at each location, it is observed that both ANN-CSE formulations capture the general trends of Sandia flame D, despite the use of different sources of data. The differences observed in both ANN-CSE formulations arise from the differences in the training data. The minor species in DI come from the results in [33] and are generally smaller than what is seen in TGLDM. The reaction rates are also different between both ANN-CSE cases, which may be caused by differences in the data or the accumulation of




Figure 6.8: Radial Favre averages of the CO_2 , H_2O , CH_4 , O_2 reaction rate and heat release rate at various locations in Sandia flame D from TGLDM-ANN-CSE (solid blue), DI-ANN-CSE (solid red) and TGLDM-CSE (dashed blue)

prediction errors. These differences are not necessarily bad, however, as DI and TGLDM rely on different assumptions. The results show that the given ANN structure is capable of representing different sources of data effectively, however the results also highlight the importance of selecting high-quality training data in order to represent the flame properly.

Many of the same trends observed in Chapter 5 are present here. The conditional errors have error magnitudes on the order of the scaled MSE from the ANN training, and are reduced when integrated to obtain the Favre averages. The maximum absolute errors tend to increase downstream for each quantity, however the errors for the reaction rates tend to decrease downstream instead. Most of the errors occur in the fuel-rich region, however there are more notable differences in the fuel-lean region with the DI-based results. These errors likely come from the DI sample space, which is much more ‘patchy’ than the TGLDM manifolds. Smoothing the DI sample space would lead to improved ANN training accuracy and ANN-CSE performance. In addition, refining ζ could benefit the DI-ANN-CSE case in particular. It is also possible to use a finer CFD mesh, as this mesh contained fewer cells than the mesh in Chapter 5.

It was generally observed that DI-ANN-CSE had larger absolute errors than TGLDM-ANN-CSE. This was expected based on the ANN training results (refer to Section 4.4.2), as the DI case had much lower accuracy than the TGLDM case. Greater ANN accuracy would improve DI-ANN-CSE, which could be achieved in three ways. Increasing the network size may lead to more benefits than the TGLDM case, as the DI samples are much more non-linear, however the training time and storage size would increase accordingly. The training settings could also be modified further, however this was attempted and did not result in significant improvement. The third and most useful option is to improve the quality of training data, which is not straightforward. Ideally, the training data should cover a wider composition space and be smoother, but it is unlikely that these factors can be directly controlled. The ANN performance is highly sensitive to the DI samples; DI-ANN-CSE cases with different velocity perturbations to generate data were unsuccessful. The generation of high-quality DI data could benefit the generalization of ANN-CSE.

It is useful to compare the two sources of data for the purposes of ANN representation. The TGLDM samples are fairly smooth, which allows for efficient and accurate ANN training. TGLDM can also incorporate reduced or full mechanisms. However, TGLDM is known to have issues with complex hydrocarbon fuels, such as heptane, where the boundaries are unclear and the trajectories may overlap. It is also difficult to represent complex combustion modes with TGLDM. An alternative data source would be needed to simulate these cases with ANN-CSE. The motivation of using DI samples is that DI could work in theory for any mechanism and provide results that rely on fewer assumptions about the chemistry. However, DI is limited to simple combustion modes and reduced mechanisms

due to computational cost. It is also necessary to simulate a successful DI case before generating the data. Although this would mitigate the time savings in ANN-CSE, the purpose of this is to generate training data for cases where CSE with tabulated chemistry is not possible, rather than replicating the DI case with ANN-CSE. The quality of the samples in either case restricts the benefits of ANN-CSE. Other tabulation methods (such as FGM) or sources such as DNS samples may benefit the ANN-CSE approach; these methods have been used in other ANN approaches [17, 18].

It should be noted that the DI samples were assumed to be parameterized by $\widetilde{Y_{CO_2}}|\zeta$ and $\widetilde{Y_{H_2O}}|\zeta$ for the purposes of ANN-CSE. An additional test was conducted where the minor species mass fractions were added as inputs (for a total of five inputs), while the outputs considered the reaction rates of the minor species mass fractions instead (for a total of nine outputs). This configuration led to significantly improved ANN training results, with MSE and R^2 being on the same order as the TGLDM results. (A similar configuration led to the same trends for the cases in Chapter 5.) However, this configuration was unstable when paired with CSE, as the ANN was much more sensitive with five inputs rather than two. Thus, the two-input ANN structure was retained.

Table 6.2: Memory usage, CSE routine execution time, and simulation wall-clock time for ANN-CSE and DI-CSE for both flames

	Sandia flame D, TGLDM		Sandia flame D, DI	
	ANN-CSE	TGLDM-CSE	ANN-CSE	DI-CSE
Memory [MB]	185	240	185	620
Execution time [s]	0.45	1:35	0.45	0.45
Wall-clock time [hh:mm:ss]	39:58:14	115:48:29	32:46:37	34:46:06

6.3.5 Computational Efficiency

The memory, execution time and wall-clock time of each simulation is compared in Table 6.2, using the same definitions as Chapter 5.3. Unsurprisingly, the two ANN-CSE cases have similar performance and have similar trends to Chapter 5.3. ANN-CSE with either data source is the quickest CSE formulation while requiring the least memory. Memory savings of about 23% are achieved relative to TGLDM-CSE, while the execution time is reduced by about 67% for each ANN-CSE case. The wall-clock time is reduced by about 65.5% with TGLDM-ANN-CSE and 72% with DI-ANN-CSE. These findings are similar to

what was observed in Chapter 5 with GRI-Mech 3.0 instead of Smooke’s mechanism. DI-CSE requires less memory than both ANN-CSE and TGLDM-CSE, requiring about 77.5 MB per core or 620 MB total. Running DI-CSE in parallel allows for similar execution and wall-clock times to both ANN-CSE cases, which are run on a single core. It is expected that the DI-CSE times would increase by a factor of about eight if DI-CSE was ran on a single core instead, making DI-CSE at least twice as slow as TGLDM-CSE. However, the purpose of DI-CSE in this context is to provide training data for ANN-CSE, rather than to be replaced by ANN-CSE. It should also be noted that the combined ANN-CSE and DI-CSE times do not exceed the TGLDM-CSE times, thus there are no computational trade-offs with using DI-CSE to generate training data. The results show that the ANN-CSE computational performance is independent of the source of data, as expected.

6.4 Summary

The objective of this chapter was to assess the influence of the source of data in the ANN development process and the simulation of Sandia flame D using ANN-CSE. In this chapter, ANN training data via TGLDM calculations (as before) and samples from direct integration are considered for diluted methane. Both methods have previously been applied to CSE of Sandia flame D [33], and are investigated here to determine if the ANN-CSE methodology is suitable for different sources of data.

The analysis closely follows that of Chapter 5, where the conditional and Favre averages of the reactive scalars are analyzed at three locations. It is found that either formulation of ANN-CSE can generally capture the trends of each quantity at each location. The TGLDM-based cases match better to each other than the DI-based cases as expected from the training results, while there are slight differences between ANN-CSE with TGLDM and DI training data due to differences in the training data. Many of the same trends from Chapter 5 are present here, namely the accuracy decreasing downstream and being worse in the fuel-rich region, however the errors are once again reduced from the conditional averages from the Favre averages. The computational performance is also assessed. ANN-CSE is found to have slight advantages over TGLDM-CSE and DI-CSE, but not to the same extent as in Chapter 5 due to differences in the codes.

The use of DI samples is promising in ANN-CSE for representing fuels that are not possible with TGLDM, however the accuracy is not as good with DI samples instead of TGLDM training data. Improvements could be made by refining the DI samples to be smoother and cover a broader region of the thermochemical space. Refining the ζ grid could also improve the DI-ANN-CSE case in particular. Finally, DI in general is

limited to reduced mechanisms, whereas tabulated chemistry methods such as TGLDM can incorporate detailed chemical kinetics. DNS samples and FGM solutions have worked in previous ML combustion chemistry applications [17, 18], and should be investigated for ANN-CSE.

Chapter 7

Conclusions

The objectives of this work were to develop a methodology for representing a given thermochemistry dataset with ML methods in the CSE combustion model, and to apply this methodology to the simulation of turbulent non-premixed methane jet flames in RANS. This work represents the first application of ML methods to CSE. Datasets consisting of tabulated chemistry (TGLDM) solutions or direct integration (DI) samples are considered, using either full or reduced mechanisms, as well as pure or diluted methane. The ANN structure and settings are optimized to the chemistry data. The ANN calculations are then integrated into the CSE code and applied to two different flames. The results show that ANN-CSE can produce results of similar quantity to the CSE cases with the reference chemistry, while generally being quicker than both DI and TGLDM. ANN-CSE is a promising approach for representing complex fuel chemistry, which so far has not been accomplished in CSE due to limitations in both tabulated and direct chemistry approaches. The main findings from this study are summarized below.

7.1 Summary of Main Findings

A review of current methods for ML turbulent combustion chemistry representation is presented in Chapter 3. ANN is generally used over other methods such as the DT-based models. Various sources of data, including tabulated chemistry, DNS samples and solutions from other models, have been considered for training. A wide variety of flames have also been investigated, including premixed and non-premixed flames and fuels such as methane, heptane and more complex fuels. The ML approach generally has excellent agreement to its source data, with computational improvements generally observed depending on the

chemistry. Much of the focus is on the improvement of the ANN training, which is a difficult task. Tasks such as normalizing and augmenting the data, partitioning the data or the ANN itself, and enforcing mass conservation in the loss function have been attempted with successful results. However, ML methods remain uncommon in turbulent combustion due to the difficulties in finding suitable data and preparing this data for training. The current work represents the first coupling of ML with CSE, by using ML techniques to represent detailed turbulent combustion chemistry in the CSE combustion model.

Chapter 4 considers the development of the ANN chemistry framework for CSE. This is an important process since the prediction accuracy highly depends on both the ANN structure and settings, as well as the quality of training data. First, ANNs are selected over other ML models, such as tree-based models, since they are more compact and can represent data to the same order of accuracy (summarized in Chapter 3). The ANN structure is designed to mimic the tabulated chemistry setup via TGLDM, being applied to predict conditional species mass fractions and chemical source terms from conditional CO_2 and H_2O mass fractions. The remaining aspects of the ANNs are determined from extensive testing. Increasing the number of neurons and hidden layers is shown to improve the ANN accuracy (measured by MSE and R^2) at the cost of increased training time and storage. ANNs consisting of five hidden layers of 50 neurons each are selected based on sufficient accuracy while minimizing training time and storage. Various activation functions are tested, with the ReLU function being selected based on simplicity and having the least training time. Both the TGLDM and DI-based datasets are augmented via Delaunay interpolation, which provides better coverage of the thermochemical space by interpolating the data. In the TGLDM cases, it was necessary to include additional data outside of the TGLDM bounds to provide stability in CSE. This data satisfies mass conservation and acts as a boundary condition in the ANN training. Z-score normalization is applied to the training data, which provides more accurate scaling than min-max normalization. The SGDM solver is selected over the Adam solver due to training time reductions of a factor of 4. The training settings are optimized to a mini-batch size of 250 samples, training length of 100 epochs, and hyperparameters of 0.01 for the constant learning rate, 0.9 for momentum, and 0.0001 for L2 regularization. Each setting is identical for all ANNs across each ζ and between datasets, however it was necessary to reduce the learning rate to 0.025 for the DI case. Training results for each case show that the ANN accuracy is generally quite good across all ζ , with notable outliers towards pure fuel and pure air due to the behaviour of the manifolds. The average R^2 for each TGLDM case exceeds 0.92 and is often around 0.98, whereas the average MSE is on the order of 10^{-2} . The DI case is less accurate, with an average R^2 of about 0.825 and average MSE on the order of 10^{-1} . Increasing both the ANN size and the volume of training data led to insignificant improvements relative to

the training time increases. ANN chemistry results in storage reductions of over 50% for each case. The ANN size is observed to be independent of the mechanism and data source, which is important for characterizing complex fuels with many species or high refinement.

In Chapter 5, ANN-CSE is applied to the Brookes-Moss flame [47] and Sandia flame D [48, 49] using TGLDM training data with GRI-Mech 3.0 chemistry. These flames burn either pure or diluted methane, resulting in two unique sets of ANNs. Results for conditional and Favre averages of the predicted quantities are analyzed at three location in each flame. First, the mixing field (represented by mean and RMS of mixture fraction) is found to be nearly identical between CSE formulations in both flames, indicating that the PDFs (used to convert between conditional and Favre averages) are similar. The ANN inputs ($\widetilde{Y_{CO_2}|\zeta}$ and $\widetilde{Y_{H_2O}|\zeta}$) are also shown to be similar in both flames, however some deviations are observed downstream and in the fuel-rich regions in both flames. Consequently, the ANN output variables are predicted well by ANN-CSE, with deviations manifesting in areas where the inputs differ. The reaction rate predictions are less smooth with ζ due to the high non-linearity. Predictions of $\widetilde{Y_{OH}|\zeta}$ were notably less accurate in the fuel-lean regions, in contrast to the other findings. Absolute errors up to the order of 10^{-3} are observed for the species mass fractions and around 10^1 for the species reaction rates. However, these errors are reduced by at least one order of magnitude when integrated with the PDF to obtain the Favre averages. Better agreement between CSE formulations for each variable is obtained for the Favre averages in both flames. The accuracy however tends to decrease downstream in both flames, as observed in the conditional averages. Errors in the fuel-rich region are attributed to the manifold mapping, while further refinement of the ζ grid and CSE ensembles could improve predictions. The computational performance of both CSE formulations is also assessed. For the cases under study, ANN-CSE is observed to require 24% less memory and 62% less execution time, while reducing wall clock time by at least 58%. The results in Chapter 5 show that ANN-CSE can represent tabulated chemistry with different fuels in turbulent non-premixed jet flames.

In Chapter 6, ANN-CSE is applied to Sandia flame D [48, 49] using either TGLDM or DI sample training data with the reduced Smooke’s mechanism for diluted methane. The focus of this chapter is to assess whether or not ANN-CSE can successfully represent data from sources aside from TGLDM, which is limited to simple fuels. The same quantities and locations as Chapter 5 are assessed here. The mixing field is captured by both ANN-CSE cases with some deviations downstream, indicating that the PDFs are similar but not exact to the reference TGLDM-CSE case. The ANN input variables are also captured by both ANN-CSE cases with minor deviations in the fuel-rich region. DI-ANN-CSE, however, has deviations in the fuel-lean regions, and $\widetilde{Y_{CO_2}|\zeta}$ is notably unstable in the fuel-rich region at

$\frac{x}{D} = 60$. The conditional minor species differ significantly in the TGLDM and DI cases due to differences in the training data, with DI underestimating the $\widetilde{Y_{CO}|\zeta}$ and $\widetilde{Y_{H_2}|\zeta}$ values. The conditional reaction rates are captured well despite the non-linearity, although the peaks are shifted towards the fuel-lean region in DI-ANN-CSE. The conditional errors are once again reduced in the Favre averages, while the errors generally increase downstream for each quantity as observed in Chapter 5. The fit between the TGLDM-based cases is quite good for the Favre averages. The trends between the two ANN-CSE cases are similar, aside from the Favre-averaged minor species mass fractions. The computational performance is also assessed, with ANN-CSE requiring about 23% less memory, 67% less execution time, and at least 65.5% less wall-clock time than TGLDM-CSE. Fewer benefits are observed relative to DI-CSE due to parallel processing. The results show that the ANN-CSE methodology is suitable for different sources of data, however the performance of ANN-CSE highly depends on the quality of training data.

7.2 Summary of Accomplishments

The main outcomes of this thesis are the following:

- Development of an ANN structure that is sufficient across all ζ for various methane mechanisms and sources of data
- Extension of CSE to include ML chemistry representation for the first time
- Implementation of a new data augmentation procedure for ANN training that results in improved training accuracy and better stability when coupled with CSE
- Incorporation of ANN calculations into the CSE code
- Successful application of ANN-CSE to two different turbulent non-premixed methane jet flames in RANS, using either tabulated chemistry or DI samples for training

7.3 Recommendations and Future Work

ANN-CSE was shown to be suitable for the cases under study. However, each case involves methane, which is a simple fuel that has been successfully simulated with either tabulated chemistry or direct integration. It is recommended to investigate fuels such as methanol

or ethanol, which have been simulated with TGLDM-CSE, as a next step in validating the ANN-CSE approach. The full benefits of ANN-CSE would be realized by simulating complex fuels, such as jet fuel or heptane, that are not possible with TGLDM-CSE or other tabulated chemistry methods.

Chapter 6 demonstrated that alternative sources of data could be used for ANN training in CSE. However, the results with DI samples were less accurate than those using TGLDM solutions due to the DI samples being more non-linear and covering a small portion of the thermochemical space. Future work should continue to refine ANN-CSE for DI samples to improve the accuracy. The use of laminar flamelet training data could also be investigated further, as these sources have been used in other ML applications. These alternative methods may be necessary for representing fuels that do not work with TGLDM or DI.

The ANN training results could be improved by using more sophisticated networks, training methods or data preparation methods. The automation of the data preparation procedure would also be beneficial, as the current method is ad-hoc and requires additional points for acceptable results. ANN-CSE is also limited in the fuel-rich regions due to issues with chemistry tabulation in fuel-rich manifolds, which are not unique to TGLDM. Better data preparation (sources, volume, scaling) could alleviate these issues.

The computational settings used throughout the study were optimized to previous CSE cases (with either TGLDM or DI) and were not modified for the cases under study. It would be useful to perform convergence studies on the cases in the current work to quantify the impact of the computational grids in ANN-CSE. This would include the refinement of the CSE ensembles and the ζ grid, and possibly the CFD mesh. The choice of turbulence models could also be investigated, including the tuning of coefficients in the standard $k-\epsilon$ model as used here.

Finally, ANN-CSE in its current form could be extended to premixed combustion using data sorted by progress variable rather than mixture fraction. Extensions to more complex combustion modes, such as MILD, could also be beneficial. For the current work, the addition of enthalpy as a third characterizing variable (or third input to the ANNs) to capture radiation may be investigated for the flames under study, however this has been completed with TGLDM-CSE. The extension to LES could also be made with ANN-CSE.

References

- [1] International Energy Agency, Key world energy statistics, Tech. rep., Organisation for Economic Co-operation and Development (2021).
- [2] D. Spalding, Mixing and chemical reaction in steady confined turbulent flames, in: Symp. (Int.) Combust., Vol. 13, Elsevier, 1971, pp. 649–657.
- [3] D. B. Spalding, Development of the eddy-break-up model of turbulent combustion, in: Symp. (Int.) Combust., Vol. 16, Elsevier, 1977, pp. 1657–1663.
- [4] B. F. Magnussen, B. H. Hjertager, On mathematical modeling of turbulent combustion with special emphasis on soot formation and combustion, in: Symp. (Int.) Combust., Vol. 16, Elsevier, 1977, pp. 719–729.
- [5] B. F. Magnussen, The eddy dissipation concept: A bridge between science and technology, in: ECCOMAS Themat. Conf. on Comput. Combust., Vol. 21, Lisbon, Portugal, 2005, p. 24.
- [6] S. B. Pope, Pdf methods for turbulent reactive flows, Prog. Energy Combust. Sci. 11 (2) (1985) 119–192.
- [7] A. Y. Klimenko, R. W. Bilger, Conditional moment closure for turbulent combustion, Prog. Energy Combust. Sci. 25 (6) (1999) 595–687.
- [8] A. Y. Klimenko, S. B. Pope, The modeling of turbulent reactive flows based on multiple mapping conditioning, Phys. Fluids 15 (7) (2003) 1907–1925.
- [9] W. K. Bushe, H. Steiner, Conditional moment closure for large eddy simulation of nonpremixed turbulent reacting flows, Phys. Fluids 11 (7) (1999) 1896–1906.
- [10] M. M. Salehi, C. Devaud, W. K. Bushe, Advances and challenges of the conditional source-term estimation model for turbulent reacting flows, Prog. Energy Combust. Sci. 104 (2024) 101172.

- [11] U. Maas, S. B. Pope, Simplifying chemical kinetics: intrinsic low-dimensional manifolds in composition space, *Combust. Flame* 88 (3-4) (1992) 239–264.
- [12] U. Maas, S. B. Pope, Implementation of simplified chemical kinetics based on intrinsic low-dimensional manifolds, in: *Symp. (Int.) Combust.*, Vol. 24, Elsevier, 1992, pp. 103–112.
- [13] S. Pope, U. Maas, Simplifying chemical kinetics: Trajectory-generated low-dimensional manifolds, *Mechanical and Aerospace Engineering Report: FDA 11* (1993).
- [14] S. B. Pope, Computationally efficient implementation of combustion chemistry using in situ adaptive tabulation, *Combust. Theor. Model.* 1 (1) (1997) 41–63.
- [15] N. Peters, Laminar diffusion flamelet models in non-premixed turbulent combustion, *Prog. Energy Combust. Sci.* 10 (3) (1984) 319–339.
- [16] J. Van Oijen, L. De Goey, Modelling of premixed laminar flames using flamelet-generated manifolds, *Combust. Sci. Technol.* 161 (1) (2000) 113–137.
- [17] T. Echekki, A. Farooq, M. Ihme, S. Sarathy, Machine learning for combustion chemistry, in: *Machine learning and its application to reacting flows: ML and combustion*, Springer International Publishing Cham, 2023, pp. 117–147.
- [18] L. Zhou, Y. Song, W. Ji, H. Wei, Machine learning for combustion, *Energy AI* 7 (2022) 100128.
- [19] B. Jin, R. Grout, W. Bushe, Conditional source-term estimation as a method for chemical closure in premixed turbulent reacting flow, *Flow Turbul. Combust.* 81 (2008) 563–582.
- [20] M. Salehi, W. Bushe, K. Daun, Application of the conditional source-term estimation model for turbulence–chemistry interactions in a premixed flame, *Combust. Theor. Model.* 16 (2) (2012) 301–320.
- [21] D. Dovizio, M. M. Salehi, C. B. Devaud, Rans simulation of a turbulent premixed bluff body flame using conditional source-term estimation, *Combust. Theor. Model.* 17 (5) (2013) 935–959.
- [22] N. Shahbazian, M. M. Salehi, C. P. Groth, Ö. L. Gülder, W. K. Bushe, Performance of conditional source-term estimation model for les of turbulent premixed flames in thin reaction zones regime, *Proc. Combust. Inst.* 35 (2) (2015) 1367–1375.

- [23] D. Dovizio, A. Debbagh, C. Devaud, Rans simulations of a series of turbulent v-shaped flames using conditional source-term estimation, *Flow Turbul. Combust.* 96 (4) (2016) 891–919.
- [24] D. Dovizio, C. Devaud, Doubly conditional source-term estimation (dcse) for the modelling of turbulent stratified v-shaped flame, *Combust. Flame* 172 (2016) 79–93.
- [25] M. Wang, J. Huang, W. Bushe, Simulation of a turbulent non-premixed flame using conditional source-term estimation with trajectory generated low-dimensional manifold, *Proc. Combust. Inst.* 31 (2) (2007) 1701–1709.
- [26] J. Huang, W. Bushe, Simulation of transient turbulent methane jet ignition and combustion under engine-relevant conditions using conditional source-term estimation with detailed chemistry, *Combust. Theor. Model.* 11 (6) (2007) 977–1008.
- [27] R. Grout, W. K. Bushe, C. Blair, Predicting the ignition delay of turbulent methane jets using conditional source-term estimation, *Combust. Theor. Model.* 11 (6) (2007) 1009–1028.
- [28] J. W. Labahn, C. B. Devaud, Investigation of conditional source-term estimation applied to a non-premixed turbulent flame, *Combust. Theor. Model.* 17 (5) (2013) 960–982.
- [29] S. Lee, C. Devaud, Application of conditional source-term estimation to two turbulent non-premixed methanol flames, *Combust. Theor. Model.* 20 (5) (2016) 765–797.
- [30] J. W. Labahn, I. Stanković, C. B. Devaud, B. Merci, Comparative study between conditional moment closure (cmc) and conditional source-term estimation (cse) applied to piloted jet flames, *Combust. Flame* 181 (2017) 172–187.
- [31] S. M. Ashrafizadeh, C. Devaud, Investigation of conditional source-term estimation coupled with a semi-empirical model for soot predictions in two turbulent flames, *Combust. Theor. Model.* 26 (5) (2022) 856–878.
- [32] A. H. Mahdipour, C. Devaud, Different conditional source-term estimation formulations applied to turbulent nonpremixed jet flames with varying levels of extinction, *Combust. Theor. Model.* 28 (3) (2024) 237–261.
- [33] A. H. Mahdipour, C. Devaud, Investigation of a new method for direct chemistry integration in conditional source-term estimation, *Combust. Theor. Model.* (2024) 1–27.

- [34] D. Dovizio, J. W. Labahn, C. B. Devaud, Doubly conditional source-term estimation (dcse) applied to a series of lifted turbulent jet flames in cold air, *Combust. Flame* 162 (5) (2015) 1976–1986.
- [35] M. Mortada, C. Devaud, Large eddy simulation of lifted turbulent flame in cold air using doubly conditional source-term estimation, *Combust. Flame* 208 (2019) 420–435.
- [36] M. P. Sitte, E. Mastorakos, Large eddy simulation of a spray jet flame using doubly conditional moment closure, *Combust. Flame* 199 (2019) 309–323.
- [37] X. H. Fang, R. Ismail, K. Bushe, M. Davy, Simulation of ecn diesel spray a using conditional source-term estimation, *Combust. Theor. Model.* 24 (4) (2020) 725–760.
- [38] A. Hussien, C. B. Devaud, Simulations of turbulent acetone spray flames using the conditional source term estimation (cse) approach, *Combust. Theor. Model.* 25 (2) (2021) 269–292.
- [39] A. Hussien, C. Devaud, Simulations of partially premixed turbulent ethanol spray flames using doubly conditional source term estimation (dcse), *Combust. Flame* 239 (2022) 111651.
- [40] A. Hussien, C. Devaud, Les study of turbulent ethanol spray flames using cse coupled with non-adiabatic chemistry tables, *Proc. Combust. Inst.* 39 (2) (2023) 2379–2388.
- [41] J. Labahn, D. Dovizio, C. Devaud, Numerical simulation of the delft-jet-in-hot-coflow (djhc) flame using conditional source-term estimation, *Proc. Combust. Inst.* 35 (3) (2015) 3547–3555.
- [42] J. Labahn, C. Devaud, Large eddy simulations (les) including conditional source-term estimation (cse) applied to two delft-jet-in-hot-coflow (djhc) flames, *Combust. Flame* 164 (2016) 68–84.
- [43] J. W. Labahn, C. Devaud, Species and temperature predictions in a semi-industrial mild furnace using a non-adiabatic conditional source-term estimation formulation, *Combust. Theor. Model.* 21 (3) (2017) 466–486.
- [44] A. M. Abdalhamid, C. Devaud, Large eddy simulations of a buoyant turbulent line flame using conditional source-term estimation (cse), *Combust. Theor. Model.* 28 (3) (2024) 344–365.

- [45] A. M. Abdalhamid, C. Devaud, E. Weckman, Development of the conditional source-term estimation (cse) framework including finite-rate kinetics and non-prescribed radiation applied to a methanol pool fire, *Fire Saf. J.* (2024) 104333.
- [46] A. M. Abdalhamid, C. Devaud, E. Weckman, Large eddy simulations of a medium-scale ethanol pool fire using conditional source-term estimation including coupled soot radiation effects, *Combust. Sci. Technol.* (2025) 1–29.
- [47] S. Brookes, J. Moss, Measurements of soot production and thermal radiation from confined turbulent jet diffusion flames of methane, *Combust. Flame* 116 (1-2) (1999) 49–61.
- [48] R. S. Barlow, J. H. Frank, Effects of turbulence on species mass fractions in methane/air jet flames, *Symp. (Int.) Combust.* 27 (1) (1998) 1087–1095.
- [49] R. Barlow, J. Frank, A. Karpetis, J.-Y. Chen, Piloted methane/air jet flames: Transport effects and aspects of scalar structure, *Combust. Flame* 143 (4) (2005) 433–449.
- [50] T. Poinso, D. Veynante, *Theoretical and numerical combustion*, R.T. Edwards, 2005.
- [51] S. B. Pope, *Turbulent Flows*, Cambridge University Press, 2000.
- [52] A. Favre, Statistical equations of turbulent gases, in: *Probl. Hydrodyn. Continuum Mech.*, SIAM Philadelphia, 1969, pp. 231–266.
- [53] L. Prandtl, 7. bericht über untersuchungen zur ausgebildeten turbulenz, *ZAMM-J. Appl. Math. Mech./Z. Angew. Math. Mech.* 5 (2) (1925) 136–139.
- [54] B. Launder, D. Spalding, The numerical computation of turbulent flows, *Comput. Methods Appl. Mech. Eng.* 3 (2) (1974) 269–289.
- [55] V. Yakhot, S. A. Orszag, Renormalization group analysis of turbulence. i. basic theory, *J. Sci. Comput.* 1 (1) (1986) 3–51.
- [56] T.-H. Shih, W. W. Liou, A. Shabbir, Z. Yang, J. Zhu, A new k-epsilon eddy viscosity model for high reynolds number turbulent flows: Model development and validation, *Tech. rep.*, NASA Lewis Research Center (1994).
- [57] D. C. Wilcox, Reassessment of the scale-determining equation for advanced turbulence models, *AIAA journal* 26 (11) (1988) 1299–1310.

- [58] D. C. Wilcox, Formulation of the kw turbulence model revisited, *AIAA journal* 46 (11) (2008) 2823–2838.
- [59] F. R. Menter, Two-equation eddy-viscosity turbulence models for engineering applications, *AIAA journal* 32 (8) (1994) 1598–1605.
- [60] C. Dopazo, E. E. O’Brien, An approach to the autoignition of a turbulent mixture, *Acta Astronaut.* 1 (9-10) (1974) 1239–1266.
- [61] C. Dopazo, E. E. O’Brien, Functional formulation of nonisothermal turbulent reactive flows, *Phys. Fluids* 17 (11) (1974) 1968–1975.
- [62] R. Bilger, The structure of diffusion flames, *Combust. Sci. Technol.* 13 (1-6) (1976) 155–170.
- [63] K. Bray, M. Champion, P. A. Libby, N. Swaminathan, Scalar dissipation and mean reaction rates in premixed turbulent combustion, *Combust. Flame* 158 (10) (2011) 2017–2022.
- [64] S. Girimaji, Assumed β -pdf model for turbulent mixing: Validation and extension to multiple scalar mixing, *Combust. Sci. Technol.* 78 (4-6) (1991) 177–196.
- [65] W. Bushe, Spatial gradients of conditional averages in turbulent flames, *Combust. Flame* 192 (2018) 314–339.
- [66] P. C. Hansen, Numerical tools for analysis and solution of fredholm integral equations of the first kind, *Inverse Probl.* 8 (6) (1992) 849.
- [67] A. N. Tikhonov, Solution of incorrectly formulated problems and the regularization method., *Sov. Dok.* 4 (1963) 1035–1038.
- [68] P. C. Hansen, The truncated svd as a method for regularization, *BIT Numer. Math.* 27 (1987) 534–553.
- [69] J. W. Labahn, C. B. Devaud, T. A. Sipkens, K. J. Daun, Inverse analysis and regularisation in conditional source-term estimation modelling, *Combust. Theor. Model.* 18 (3) (2014) 474–499.
- [70] G. P. Smith, Gri-mech 3.0, <http://combustion.berkeley.edu/gri-mech/version30/text30.html>.

- [71] R. J. Renka, Algorithm 751: Tripack: a constrained two-dimensional delaunay triangulation package, *ACM Trans. Math. Softw. (TOMS)* 22 (1) (1996) 1–8.
- [72] G. V. Nivarti, M. M. Salehi, W. K. Bushe, A mesh partitioning algorithm for preserving spatial locality in arbitrary geometries, *J. Comput. Phys.* 281 (2015) 352–364.
- [73] H. P. Tsui, W. K. Bushe, Conditional source-term estimation using dynamic ensemble selection and parallel iterative solution, *Combust. Theor. Model.* 20 (5) (2016) 812–833.
- [74] *OpenFOAM*, <https://openfoam.org/>.
- [75] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, 2016, pp. 785–794.
- [76] T. Hastie, R. Tibshirani, J. Friedman, *Additive Models, Trees, and Related Methods*, Springer New York, New York, NY, 2009, pp. 295–336. doi:10.1007/978-0-387-84858-7_9. URL https://doi.org/10.1007/978-0-387-84858-7_9
- [77] J. H. Friedman, Stochastic gradient boosting, *Comput. Stat. Data Anal.* 38 (4) (2002) 367–378.
- [78] F. Christo, A. Masri, E. Nebot, Artificial neural network implementation of chemistry with pdf simulation of h₂/co₂ flames, *Combust. Flame* 106 (4) (1996) 406–427.
- [79] F. Christo, A. Masri, E. Nebot, T. Turányi, Utilising artificial neural network and repro-modelling in turbulent combustion, in: *Proc. ICNN’95 Int. Conf. on Neural Netw.*, Vol. 2, IEEE, 1995, pp. 911–916.
- [80] F. C. Christo, A. R. Masri, E. M. Nebot, S. B. Pope, An integrated pdf/neural network approach for simulating turbulent reacting systems, in: *Symp. (Int.) Combust.*, Vol. 26, Elsevier, 1996, pp. 43–48.
- [81] J. Blasco, N. Fueyo, C. Dopazo, J. Ballester, Modelling the temporal evolution of a reduced combustion chemical system with an artificial neural network, *Combust. Flame*. 113 (1-2) (1998) 38–52.
- [82] J. A. Blasco, N. Fueyo, J. Larroya, C. Dopazo, Y.-J. Chen, A single-step time-integrator of a methane–air chemical system using artificial neural networks, *Comput. Chem. Eng.* 23 (9) (1999) 1127–1133.

- [83] J. A. Blasco, N. Fueyo, C. Dopazo, J.-Y. Chen, A self-organizing-map approach to chemistry representation in combustion applications, *Combust. Theor. Model.* 4 (2000) 61–76.
- [84] J.-Y. Chen, J. Blasco, N. Fueyo, C. Dopazo, An economical strategy for storage of chemical kinetics: Fitting in situ adaptive tabulation with artificial neural networks, *Proc. Combust. Inst.* 28 (1) (2000) 115–121.
- [85] A. Kempf, F. Flemming, J. Janicka, Investigation of lengthscales, scalar dissipation, and flame orientation in a piloted diffusion flame by les, *Proc. Combust. Inst.* 30 (1) (2005) 557–565.
- [86] A. Chatzopoulos, S. Rigopoulos, A chemistry tabulation approach via rate-controlled constrained equilibrium (rcce) and artificial neural networks (anns), with application to turbulent non-premixed $\text{CH}_4/\text{H}_2/\text{N}_2$ flames, *Proc. Combust. Inst.* 34 (1) (2013) 1465–1473.
- [87] L. L. Franke, A. K. Chatzopoulos, S. Rigopoulos, Tabulation of combustion chemistry via artificial neural networks (anns): Methodology and application to les-pdf simulation of sydney flame 1, *Combust. Flame* 185 (2017) 245–260.
- [88] T. Readshaw, T. Ding, S. Rigopoulos, W. Jones, Modeling of turbulent flames with the large eddy simulation–probability density function (les–pdf) approach, stochastic fields, and artificial neural networks, *Phys. Fluids* 33 (3) (2021).
- [89] T. Ding, T. Readshaw, S. Rigopoulos, W. Jones, Machine learning tabulation of thermochemistry in turbulent combustion: An approach based on hybrid flamelet/random data and multiple multilayer perceptrons, *Combust. Flame* 231 (2021) 111493.
- [90] T. Ding, S. Rigopoulos, W. Jones, Machine learning tabulation of thermochemistry of fuel blends, *Appl. Energy Combust. Sci.* 12 (2022) 100086.
- [91] T. Readshaw, L. L. Franke, W. Jones, S. Rigopoulos, Simulation of turbulent premixed flames with machine learning-tabulated thermochemistry, *Combust. Flame* 258 (2023) 113058.
- [92] A. Liu, T. Ding, R. Liu, S. Rigopoulos, K. Luo, Machine learning tabulation of thermochemistry for turbulent dimethyl ether (dme) flames, *Fuel* 360 (2024) 130338.

- [93] S. Bhalla, M. Yao, J.-P. Hickey, M. Crowley, Compact representation of a multi-dimensional combustion manifold using deep neural networks, in: *Jt. Eur. Conf. Mach. Learn. Knowl. Discov. Databases*, Springer, 2019, pp. 602–617.
- [94] B. A. Sen, S. Menon, Turbulent premixed flame modeling using artificial neural networks based chemical kinetics, *Proc. Combust. Inst.* 32 (1) (2009) 1605–1611.
- [95] K. Wan, C. Barnaud, L. Vervisch, P. Domingo, Chemistry reduction using machine learning trained from non-premixed micro-mixing modeling: Application to dns of a syngas turbulent oxy-flame with side-wall effects, *Combust. Flame.* 220 (2020) 119–129.
- [96] H.-T. Nguyen, P. Domingo, L. Vervisch, P.-D. Nguyen, Machine learning for integrating combustion chemistry in numerical simulations, *Energy AI* 5 (2021) 100082.
- [97] C. Chi, X. Xu, D. Thévenin, Efficient premixed turbulent combustion simulations using flamelet manifold neural networks: A priori and a posteriori assessment, *Combust. Flame* 245 (2022) 112325.
- [98] M. Ihme, C. Schmitt, H. Pitsch, Optimal artificial neural networks and tabulation methods for chemistry representation in les of a bluff-body swirl-stabilized flame, *Proc. Combust. Inst.* 32 (1) (2009) 1527–1535.
- [99] B. A. Sen, S. Menon, Linear eddy mixing based tabulation and artificial neural networks for large eddy simulations of turbulent flames, *Combust. Flame* 157 (2010) 62–74.
- [100] B. A. Sen, E. R. Hawkes, S. Menon, Large eddy simulation of extinction and reignition with artificial neural networks based chemical kinetics, *Combust. Flame.* 157 (3) (2010) 566–578.
- [101] P. Sinaei, S. Tabejamaat, Large eddy simulation of methane diffusion jet flame with representation of chemical kinetics using artificial neural network, *Proc. Inst. Mech. Eng., Part E: J. Process Mech. Eng.* 231 (2) (2017) 147–163.
- [102] Y. Zhang, S. Xu, S. Zhong, X.-S. Bai, H. Wang, M. Yao, Large eddy simulation of spray combustion using flamelet generated manifolds combined with artificial neural networks, *Energy AI* 2 (2020) 100021.
- [103] O. Owoyele, P. Kundu, M. M. Ameen, T. Echekeki, S. Som, Application of deep artificial neural networks to multi-dimensional flamelet libraries and spray flames, *Int. J. Engine Res.* 21 (1) (2020) 151–168.

- [104] M. Hansinger, Y. Ge, M. Pfitzner, Deep residual networks for flamelet/progress variable tabulation with application to a piloted flame with inhomogeneous inlet, *Combust. Sci. Technol.* 194 (8) (2022) 1587–1613.
- [105] S. Abdelwahid, M. R. Malik, H. A. A. K. Hammoud, F. E. Hernández-Pérez, B. Ghanem, H. G. Im, Large eddy simulations of ammonia-hydrogen jet flames at elevated pressure using principal component analysis and deep neural networks, *Combust. Flame.* 253 (2023) 112781.
- [106] A. Mousemi, M. Jadidi, S. B. Dworkin, W. K. Bushe, Application of machine learning in low-order manifold representation of chemistry in turbulent flames, *Combust. Theor. Model.* 27 (1) (2023) 83–102.
- [107] K. Li, P. Rahnama, R. Novella, B. Somers, Combining flamelet-generated manifold and machine learning models in simulation of a non-premixed diffusion flame, *Energy AI* 14 (2023) 100266.
- [108] C. M. Bishop, *Neural networks for pattern recognition*, Oxford university press, 1995.
- [109] K. Fukushima, Visual feature extraction by a multilayered network of analog threshold elements, *IEEE Trans. Syst. Sci. Cybern.* 5 (4) (1969) 322–333.
- [110] A. L. Maas, A. Y. Hannun, A. Y. Ng, et al., Rectifier nonlinearities improve neural network acoustic models, in: *Proc. Int. Conf. Mach. Learn.*, Vol. 30, Atlanta, GA, 2013, p. 3.
- [111] C. D. Pierce, P. Moin, Progress-variable approach for large-eddy simulation of non-premixed turbulent combustion, *J. Fluid Mech.* 504 (2004) 73–97.
- [112] X. Wen, X.-S. Bai, K. Luo, H. Wang, Y. Luo, J. Fan, A generalized flamelet tabulation method for partially premixed combustion, *Combust. Flame* 198 (2018) 54–68.
- [113] M. D. Smooke, *Reduced kinetic mechanisms and asymptotic approximations for methane-air flames: a topical volume*, Springer, 1991.
- [114] D. P. Kingma, J. L. Ba, Adam: A method for stochastic optimization, in: *3rd Int. Conf. Learn. Represent.*, 2015.
- [115] L. Bottou, Large-scale machine learning with stochastic gradient descent, in: *Proc. COMPSTAT’2010: 19th Int. Conf. Comput. Stat.*, Springer, 2010, pp. 177–186.