

Kinetic Subspace Investigation Using Neural Network for Uncertainty Quantification in Nonpremixed Flamelets

Benjamin C. Koenig^a, Weiqi Ji^a, Sili Deng^{a,*}

^a*Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

Abstract

Propagating uncertainties in kinetic models through turbulent combustion simulations to properly quantify the uncertainties in the simulation results remains a challenging and numerically expensive problem. Efficient approaches have been proposed for certain flames in the flamelet region by reducing their uncertainty input from a high-dimensional kinetic parameter space to a one-dimensional variable. However, this one-dimensional assumption does not apply to all flamelet regimes. In the current work, we developed a systematic approach to discover low-dimensional active subspace reductions that apply to the entire mixture fraction space of the flamelet, and that function even in cases where the uncertainty response is not uniform across the entire solution domain and the one-dimensional assumption does not apply. In doing so, we are able to achieve uncertainty quantification with a tunable tradeoff between high accuracy and low computational cost through careful selection of subspace dimensionality. We facilitated computation in this method using a specifically designed deep neural network based surrogate model to compute the temperature gradients of the flamelet profile to the kinetic parameters. We presented, as a proof-of-concept, a two-stage active subspace reduction on the kinetic parameter space of a non-premixed methane flamelet. In doing so we demonstrated that its uncertainty response cannot be represented by a one-dimensional kinetic variable due to its uncorrelated behavior across the mixture fraction domain. We instead proposed a four-dimensional active subspace that captures 98% of the uncertainty response in the flame profile at largely reduced computational cost compared to the full kinetic parameter space. The tunability, generality, and reduced computational cost of this method demonstrate its potential to facilitate uncertainty quantification of complex and large-scale combustion problems.

Keywords: Uncertainty Quantification; Deep Neural Network; Active Subspace; Flamelet

* Corresponding Author.
E-mail Address: silideng@mit.edu (S. Deng)

1. Introduction

Combustion simulations for practical applications usually involve turbulent flows and hundreds of chemical species and reactions [1]. The kinetic parameters of these reactions are not precisely known, and their uncertainties can propagate through the combustion model and result in substantial uncertainties in the simulation results [2–5]. Uncertainty quantification, such as Monte Carlo sampling, aims to evaluate the effect of kinetic uncertainties on the simulation results [6]. However, due to the computational cost needed to run many turbulent combustion simulations, directly propagating high-dimensional kinetic uncertainty is not feasible [2].

In the seminal work of Mueller et al. [2], a physics-informed dimension reduction was proposed to represent high-dimensional kinetic uncertainties in a flamelet-based turbulent combustion model. This approach first propagates the kinetic uncertainty (about 300 dimensions for methane) onto the flamelets, then assumes that the temperature at different mixture fractions and scalar dissipation rates are correlated with each other such that the uncertainty across both domains can be parameterized with only one variable. If such an assumption is universally valid, huge computational savings can be expected for kinetic uncertainty quantifications in turbulent combustion simulations in the flamelet regime.

The active subspace method [3, 7, 8] is an alternative solution to derive a low-dimensional representation to high-dimensional kinetic uncertainties. This approach seeks to identify a low-rank kinetic subspace within which changes in the kinetic parameters have the greatest effect on the flame simulation results. Uncertainty quantification can be carried out in this subspace, instead of the full, high-dimensional space, to reduce computational cost. The eigenvalue-based decomposition used in the active subspace method is similar to principal component analysis (PCA) in its discovery of low-dimensional directions in a high-dimensional space. It differs in that PCA typically reduces the dimension of the output space itself, while the active subspace method uses output data and gradients in order to reduce the dimension of the input space [9]. A key difference from traditional methods such as [10] in this specific combustion case is that the active subspace method analyzes the response across the full input uncertainty space, not just at the nominal value. This method has previously been successful in propagating kinetic uncertainty to the lift-off height of the Cabra flame [3].

Previous active subspace work generally examines only a single quantity, such as liftoff height [3] [11] or ignition delay time [11], rather than the response of the entire temperature profile as is analyzed in this work. It is expected that if the aforementioned assumption of a one-dimensional flamelet response to kinetic uncertainties is valid across different mix-

ture fractions and strained conditions, then a one-dimensional active subspace could also be found for the entire flame profile. However, validating this assumption is nontrivial. The active subspace method is computationally cheaper than many other sensitivity analysis-based methods, but still requires thousands of evaluations of flamelet profile gradients with respect to the kinetic parameters to obtain good accuracy.

To evaluate the gradients of flamelet responses to kinetic uncertainties efficiently and accurately in the current work, we built a deep neural network surrogate model for flamelet simulation. Neural networks trained using the traditional backpropagation algorithm [12] are capable of highly efficient gradient computation, though are limited by their fixed grids and struggle to adapt to training data generated using PDE solvers with adaptive mesh refinement such as Cantera [13]. As will be discussed in Sec. 2, our method utilized the Non-linear Independent Dual System [14] neural network structure, which is similar to the Deep Operator Neural Network [15]. It is able to evaluate the effects of rate constant perturbations on the temperature profiles of nonpremixed flamelets on any arbitrary grid with a single training cycle due to its unique structure, while maintaining the inexpensive gradient evaluations inherent to traditional networks. With the network-supplied gradient information, we then characterized the uncertainty of the flamelet and investigated if there exists a low-dimensional kinetic subspace that can capture the uncertainty response of its entire mixture fraction profile with high accuracy. We did so by adapting the recently developed vector-valued active subspace method [16] and applying a two-stage singular-value decomposition (SVD). The first SVD was to identify local active subspaces at each fixed mixture fraction. Then, the second SVD identified the global kinetic subspace from the set of all local subspaces. The terms "local" and "global" are used here to distinguish between a local subspace at a fixed mixture fraction, and a global subspace applicable across the full mixture fraction domain. Both cases are found using the entire kinetic uncertainty space. Variable strain rates were not considered based on previous work suggesting that they do not greatly impact kinetic sensitivity directions [17], though verifying such a result in this context is a focus of our future work.

We aimed to demonstrate the skill and efficiency of this neural network surrogate model and leverage its numerically inexpensive and discretization-agnostic gradients to discover a low-dimensional kinetic subspace that applies across the entire mixture fraction domain of a nonpremixed methane flamelet. This novel and efficient approach is able to represent the flamelet uncertainty response in greatly reduced dimensionality, even in cases where the local active subspace depends on the mixture fraction. Such a method has potential to be scaled up and applied to facilitate the uncertainty quantification of large eddy turbulent

combustion simulations through the laminar flamelet model.

2. Methods

In this work, we investigated if a global kinetic subspace for a one-dimensional nonpremixed methane flamelet can be constructed using the active subspace method, facilitated by a neural network surrogate model of the flamelet. In Secs. 2.1 and 2.2, the methodologies for propagating kinetic uncertainties to flamelet temperature profiles and obtaining low-dimensional active subspaces are introduced. In Sec. 2.3, the structure, training, and validation of the neural network used to compute gradients for the subspace algorithm is discussed.

Throughout this work, a normalized version of the rate constant uncertainty ranges computed in previous work [18] is utilized,

$$x_\ell = \frac{\ln k_\ell / k_{\ell,0}}{\frac{1}{3} \ln u_\ell} \sim N(0, 1), \quad (1)$$

where x_ℓ is the ℓ th index of the normalized rate constant perturbation vector $\mathbf{x}$, k_ℓ is the perturbed value of the ℓ th rate constant, $k_{\ell,0}$ is the nominal value of the ℓ th rate constant, u_ℓ is the uncertainty factor corresponding to $k_{\ell,0}$ as reported in the literature [18], and $N(0, 1)$ denotes the standard normal distribution with zero mean and unit variance. This normalization facilitates the neural network training and subspace analysis presented below.

2.1. Local Active Subspace at Fixed Mixture Fraction

The general algorithm, methodology, and proofs for the active subspace method are presented in [7]. The following will present a basic overview of the concepts specific to our applications.

The one-dimensional flame is solved in the physical space instead of the mixture fraction space. The mixture fraction used throughout this work is then defined using the hydrogen elemental mass fraction

$$Z = \frac{W_{\text{mix}} - W_{\text{ox}}}{W_{\text{fuel}} - W_{\text{ox}}}, \quad (2)$$

where Z is the mixture fraction at a given location, and W_{mix} , W_{ox} , and W_{fuel} represent the hydrogen mass fractions of the mixture, oxidizer stream, and fuel stream, respectively.

We aim to identify, at a given location Z , an r_Z -dimensional subspace in the d -dimensional kinetic rate constant space (with $r_Z < d$) that describes the bulk of the variation of the flame temperature T_Z at location Z , such that

$$T_Z(\mathbf{x}_d) \approx T_Z(\mathbf{x}_{r_Z}), \quad (3)$$

where $\mathbf{x}_d$ is a full-rank vector of rate constant perturbations for the d reactions in the kinetic model,

while $\mathbf{x}_{r_Z}$ is the same vector expressed with rank r_Z , reduced in an information-preserving way using the r_Z -dimensional subspace. We identify this active subspace of dimension r_Z with the eigen-decomposition of the matrix $\mathbf{C}$ at the location Z , which is analogous to the SVD of the temperature gradients.

$$\begin{aligned} \mathbf{C} &= \frac{1}{M} \sum_{j=1}^M \nabla_x T_Z(\mathbf{x}_j) (\nabla_x T_Z(\mathbf{x}_j))^T \\ &= \mathbf{W} \mathbf{\Lambda} \mathbf{W}^T. \end{aligned} \quad (4)$$

Here, M is the total number of rate constant samples j that we generate from the full rate constant uncertainty space in $\mathbf{x}$, with the gradient $\nabla_x T_Z(\mathbf{x}_j)$ evaluated once per iteration j . We compute these gradient values in our specialized neural network, as detailed in Sec. 2.3. By virtue of the network-powered gradients, we do not need to treat M with as much care as in more computationally expensive methods, and can simply choose a larger M if the results do not converge properly. From the decomposition $\mathbf{W} \mathbf{\Lambda} \mathbf{W}^T$, we can extract d eigenvectors $\mathbf{w}_1$ through $\mathbf{w}_d$ of length d each, corresponding to d eigenvalues λ_1 through λ_d . Due to the symmetry of $\mathbf{C}$, these eigenvectors form an orthonormal basis that can be used to fully describe any sample $\mathbf{x}_d$. The reduced dimension r_Z is chosen by identifying the first two eigenvalues λ_r and λ_{r+1} between which there is a substantial difference of multiple orders of magnitude. This gap informs us that the primary direction(s) of variation in the flame temperature occur in the first r eigenvectors $\mathbf{w}_1$ through $\mathbf{w}_r$, the set of which is defined as the active subspace. The accuracy of the reduced dimension subspace in reconstructing the matrix $\mathbf{C}$ and thus capturing the gradient behavior is dictated by the percentage of the sum of all eigenvalues captured by the sum up to the index r . We can project a given sample $\mathbf{x}_d$ onto this subspace in order to express it as a linear combination of these r basis eigenvectors, reducing the space within which we need to perturb the kinetic model to characterize uncertainty by $d - r$.

2.2. Global Active Subspace Across A Flamelet

Performing this reduction at a given location Z generates one subspace of dimension r_Z . However, the active subspace does not necessarily remain the same or even similar at different mixture fractions, as will be shown in Sec. 3.2. To investigate if a global kinetic subspace across a flamelet can be constructed, we develop a second stage of SVD reduction. We assume for now (to be demonstrated as well in Sec. 3.2) that $r_Z = 1$ for all Z and thus the local active subspace at any given Z_i is described with the single vector $\mathbf{w}_1$, which we denote $\mathbf{w}_i$ henceforth to distinguish between local subspaces at different mixture fractions. This assumption is made here to simplify notation but is not a prerequisite for use of

this method, which is applicable to subspaces of arbitrary dimensions. The methodology presented below serves to reduce the set of one-dimensional subspaces at n mixture fractions, where n is a sampling number larger than d , down to a final global subspace basis of dimension r_u that is applicable across the entire mixture fraction domain.

We define $\mathbf{A}$ as the $n \times d$ matrix with each row i defined by $\mathbf{w}_i$. It therefore contains the local active kinetic subspaces at all n grid points in the mixture fraction domain. The singular value decomposition on $\mathbf{A}$ follows

$$\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^T, \quad (5)$$

where $\mathbf{S}$ is an $n \times d$ diagonal matrix with d singular values σ that correspond to the relative importance of the principal directions $\mathbf{V}$. The squared singular values σ^2 relate to the eigenvalues of the matrix $\mathbf{A}\mathbf{A}^T$, and can be analyzed similarly to the eigen-decomposition of $\mathbf{C}$ in Eq. (4). If we do not observe a large gap between the first two squares of singular values, as may occur in cases where a multidimensional subspace is necessary to fully characterize the uncertainty response, we can select our reduced dimension r_u such that the percentage of the sum of all squared singular values represented by the sum of those from 1 to r_u meets our accuracy needs. For example, if the sum of the first three squared singular values represents 95% of the total sum, then we expect the basis formed by the collection of principal directions up to $r_u = 3$ to capture 95% of the variance in matrix $\mathbf{A}$, and thus of the variance in the local subspaces. In this way, we can tune the accuracy versus computational cost tradeoff of the method to any specific application's needs. After selecting r_u , we obtain from $\mathbf{V}$ a reduced space $\mathbf{v}_k$ with k ranging from 1 to r_u , where the vectors $\mathbf{v}$ describe the same rate constant space as the vectors $\mathbf{w}$ but at greatly reduced dimensionality. While the first stage of our reduction method in Sec. 2.1 evaluated local subspaces from the corresponding temperature gradients at fixed mixture fractions, the second stage here combines all local results to discover a global subspace across the entire flamelet. We emphasize that in contrast to other local subspace analyses seen in the literature, our approach allows us to generate not only a local reduced-dimension subspace $\mathbf{w}_1$, but also a globally applicable and further reduced subspace basis of $\mathbf{v}_k$. We also note that the method is easily generalizable to universal cases with scopes beyond this proof of concept. In this work, we investigated the flamelet at a fixed strain rate, as previous research suggests that strain rates have minimal effects on kinetic sensitivity directions [17]. However, for future applications that require sampling of strain rate or other inputs, any local subspace may be included in the $\mathbf{A}$ matrix and then reduced further with Eq. (5).

To fully validate the global kinetic subspace, Cantera [13] was utilized to simulate temperature profiles

from perturbations in the full d -dimensional kinetic space, which were compared against those simulated using the reduced r_u -dimensional kinetic subspace. We generated, based on Eq. (1), 5,000 random perturbation vectors. To compare the subspace and full space results, we projected these samples into the r_u -dimensional subspace and then reconstructed the full-dimensional samples using these projections. In doing so, we created new samples representing the same data with greatly reduced rank that can still be used with Cantera or the neural network:

$$\mathbf{b}_{r_u,j} = \sum_{k=1}^{r_u} [(\mathbf{b}_{d,j} \cdot \mathbf{v}_k) \mathbf{v}_k]. \quad (6)$$

Here $\mathbf{b}_{d,j}$ is the j th sample in the full kinetic space d , and $\mathbf{b}_{r_u,j}$ is the same sample expressed in reduced rank using the r_u basis vectors $\mathbf{v}_k$. This operation reduces each vector $\mathbf{b}_{d,j}$ to a linear combination of the $\mathbf{v}_k$ basis directions. From this linear combination it reconstructs parameter vectors of length d but rank r_u to test in Cantera using the same boundary conditions as in the initial data generation in Sec. 2.3. This enables comparison of temperature profiles based separately on the sets of vectors $\mathbf{b}_{r_u,j}$ and $\mathbf{b}_{d,j}$, to evaluate how much of the original information is captured by the r_u -rank sample.

2.3. Neural Network Surrogate Model for Flamelets

Gradient computation to construct the matrix $\mathbf{C}$ in Eq. (4) was performed by our specifically structured neural network. Given a rate constant perturbation vector $\mathbf{x}$ as defined in Eq. (1), this network was trained to predict the full temperature response profile of a single flamelet in the mixture fraction space. We started by generating 5,000 nonpremixed methane/oxygen counterflow training data sets using the GRI-Mech 3.0 mechanism [19] with Cantera [13] at a pressure of 1 bar, strain rate of 4600 s^{-1} , inlet temperatures of 300K, and global equivalence ratio of 2, on an 18 mm wide domain in which methane enters from one boundary and oxygen from the other. We trained a neural network on this data with a structure inspired by the Non-linear Independent Dual System network [14], shown in Fig. 1. In the top (blue) parameter network depicted in Fig. 1, the j th rate constant perturbation vector $\mathbf{x}_j$ of length $d = 217$ generated from Eq. (1) is fed to the 217-node input layer. This is passed through eight layers with 16 nodes each until it reaches the final parameter layer. Similarly, the bottom (red) coordinate network has a one-node input layer for the i th mixture fraction within the grid of sample j , $(\mathbf{Z}_j)_i$. It is passed through six layers with 16 nodes each until it reaches the final coordinate layer. The inner product of the two networks' final layers provides the scalar output temperature value $T(\mathbf{x}_j)((\mathbf{Z}_j)_i)$ at that location i . By running the network in a batch over all grid points i , we can reconstruct the full temperature profile $T(\mathbf{x}_j)$.




Fig. 1: Neural network surrogate model. The parameter and coordinate branches are independent until the last layer in which their results are combined via inner product to arrive at a final temperature prediction. A batch of runs throughout the mixture fraction domain can recreate the full temperature profile. Skip connections are not shown to maintain clarity. Training occurs independently of discretization, and the network can be evaluated on any input Z in an arbitrary mesh.

We note here that the grid size is not directly specified in the network structure. The network is therefore not constrained to a fixed grid, and can solve for a temperature profile on any arbitrarily discretized Z_j . We can thus train without interpolation on simulation data from Cantera, which solves on adaptively refined grids that differ with each j . After training is complete, we can again take advantage of the network structure to generate temperature profiles on any grid of our choice, simplifying downstream calculations. We additionally remark that the isolation of the parameter and coordinate networks separates the two fundamentally different inputs: rate constants and mixture fractions. In doing so, it may yield greater performance than standard networks due to the inductive bias, or existing knowledge of the problem's structure, encoded directly into the network [15].

In order to optimize performance, we selected the hyperparameters of the network such as the number of layers, nodes per layer, and learning rate using the Ray Tune package [20], which methodically and efficiently sweeps through a provided hyperparameter input range in order to select the best values to use during training. While this tool does not provide an exactly optimal network structure, it greatly increased performance in this case when compared to initial testing. The final network is shown in Fig. 1. The learning rate and batch size used in training were 0.00076 and 64, respectively.

Certain other features in the network were selected manually. Residual networks [21] with skip connections every two layers were used for both the coordinate and parameter networks. Sigmoid-weighted Linear Units [22] were used as activation functions. The ADAM optimizer [23] was used to train for the first eight hundred epochs, after which a second-order limited memory Broy-

den-Fletcher-Goldfarb-Shanno (LBFGS) optimizer was used until the training plateaued to improve final performance while limiting overfitting. 15% of the data was withheld as out of sample testing data, and 33% of the remainder for in-sample validation.

After training, we verified the network's gradients against the Cantera solver. The network trains to minimize error in the temperature values, but this does not guarantee low error in their gradients. We justified use of the network for gradient-intensive subspace computation by comparing its gradients against those found analytically in Cantera. At each location Z in $\mathbf{Z}$ we computed gradient values from the network using the same automatic differentiation algorithm as was used in network training, and those from Cantera using built-in analytical sensitivity packages, then compared their cosine similarities,

$$CS_Z = \frac{\langle \nabla_{\mathbf{x}} T_{Z,net}, \nabla_{\mathbf{x}} T_{Z,Cantera} \rangle}{\| \nabla_{\mathbf{x}} T_{Z,net} \| \| \nabla_{\mathbf{x}} T_{Z,Cantera} \|}. \quad (7)$$

Here CS_Z is the cosine similarity at the mixture fraction Z , $\nabla_{\mathbf{x}} T_{Z,net}$ is the temperature gradient vector evaluated through the network at Z , and $\nabla_{\mathbf{x}} T_{Z,Cantera}$ is the same in Cantera. CS_Z ranges from -1 to 1 , where a value of ± 1 indicates perfect directional alignment but positive or negative correlations, respectively, and a value of 0 indicates orthogonality.

3. Results and Discussion

3.1. Surrogate Neural Network Training

We present mean squared error training results for our specialized neural network surrogate model in Fig. 2. The ADAM optimizer ran for roughly 800

training epochs, though in Fig. 2 we observe future training may be accelerated through earlier stopping. The LBFGS optimizer then ran just until the error plateaued. This second step improved accuracy, while its brevity limited overfitting. A plot comparing a sample network-generated temperature profile against the Cantera solution is shown in Fig. 3. We observe strong qualitative agreement between the network and Cantera, and a quantitative error of less than 1% within the high-temperature ($>1000\text{K}$) region.




Fig. 2: Neural network mean squared error loss at each epoch. Sharp decay at end of training from LBFGS optimizer. Training, validation, and testing error decreased similarly throughout training, suggesting minimal overfitting.




Fig. 3: Network temperature evaluation on testing data at the end of training. Error in flame is kept below 1%. Qualitative agreement with the Cantera solutions is strong.

In addition to the standard network loss, we report out-of-sample gradient validation against the Cantera data. The network loss and comparison against Cantera from Figs. 2 and 3 is based solely on the temperature values T taken directly from the final layer of the network. However, our sensitivity analysis method primarily involves use of temperature gradients $\partial T / \partial \mathbf{x}$, which require an additional auto-differentiation step. In Fig. 4, we report the cosine similarity of gradients computed through auto-differentiation in the trained network against those computed analytically in Cantera. We observe a cosine similarity

of 0.99 on average in the high-temperature flame region. As expected, the low-temperature out-of-flame regions with behavior not captured by the trained chemical kinetic model do not see useful agreement and are not relevant to the analysis at hand. With strong agreement in temperature and temperature gradient between our specialized neural network and Cantera, we proceed in the next section to the application of our neural network in kinetic sensitivity analysis and uncertainty quantification.




Fig. 4: Cosine similarity of temperature gradients across the flame. Very strong gradient agreement of 0.99 on average is seen between network predictions and Cantera solutions throughout the high-temperature in-flame region.

3.2. Kinetic Subspace Identification

We implemented the active subspace algorithm detailed in Eq. (4) and computed the eigenvalues and eigenvectors of the matrix $\mathbf{C}$. Fig. 5 depicts the eigenvalues at three representative locations in the flame: the oxidizer side, the fuel side, and near the peak temperature location. A large drop of roughly two orders of magnitude is seen between the first and second eigenvalues in all three examples. While not pictured, a similar result was also found in the rest of the mixture fraction domain. Thus, we chose $r = 1$ for the local active subspaces and defined each with $\mathbf{w}_1$ only, discarding the extra dimensions $\mathbf{w}_2$ through $\mathbf{w}_d$ as the inactive subspace.

In Fig. 6, we inspect three of these active subspaces $\mathbf{w}_i$ at the same three representative locations across the flame (oxidizer side, near temperature peak, and fuel side) in an attempt to discover generality across the domain. We do not see strong agreement in the local subspace directions, indicating that a single-dimension global subspace does not exist. We note, however, that of the 217 rate constants/reactions, the subspaces at the three locations seem to all include the same five or ten most influential reactions, with different ratios, in their active directions. This observed similarity motivates the second stage of reduction detailed in Eq. (5) to an r_u -dimensional global subspace.

We present the results of obtaining the global kinetic subspace via singular value decomposition in




Fig. 5: First ten eigenvalues from $\mathbf{C}$ matrix in Eq. (4). Sharp decline of roughly two orders of magnitude is visible between the first and second value seen in three samples spaced around the temperature peak at $Z = 0.2$, justifying truncation of active subspace to one direction per coordinate.




Fig. 6: Single-direction local active subspaces $\mathbf{w}_1$ reported at the same three mixture fractions as Fig. 5. Lack of directional agreement suggests that a single local subspace cannot describe behavior across the entire mixture fraction domain.




Fig. 7: Squares of singular values (σ^2) from the SVD decomposition of subspace matrix $\mathbf{A}$ in Eq. (5). Sharp drop-off like that in Fig. 5 is not observed after the first index, motivating the use of more than one basis to capture full behavior.

Figs 7, 8 and 9. As shown in Fig. 7, the SVD yields squared singular values that do not drop off as sharply as the original subspace eigenvalues do. This is expected from our hypothesis based on Fig. 6 that we




Fig. 8: Squares of singular values from SVD on matrix $\mathbf{A}$, expressed as the percentage of all squares of singular values in the sum up to a given index. 100% at r_u would indicate that subspace variation across $\mathbf{A}$ could be represented by the set of directions up to r_u . Here 99.5% at the fourth index justifies the truncation to $r_u = 4$.

would need $r_u > 1$ basis directions to describe the variation in subspaces across the mixture fraction domain. In Fig. 8, we take the sum of these squared singular values up to each index to inform our choice of r_u , the dimension of the final subspace basis. We see that with just two active directions we can account for over 90% of the subspace variation across the mixture fraction, with three directions over 98%, and with four over 99.5%. With one, however, we can only capture 62%. We selected the basis with $r_u = 4$, made up of the four vectors shown in Fig. 9 as our final global active subspace basis, though $r_u = 2$ or $r_u = 3$ would be valid options in applications that require lower computational cost and lower accuracy, and conversely $r_u > 4$ if extreme accuracy is needed. In our case, we are able to accurately reduce the dimensionality of the uncertainty quantification problem at hand from $d = 217$ for this methane reaction model down to $r_u = 4$. A more detailed summary of the kinetic composition of this 4-dimensional subspace is available in the supplementary material.




Fig. 9: Final four active subspace directions chosen from the 217-dimensional kinetic space, each shown in a different color.

3.3. Global Kinetic Subspace Validation

We validated the accuracy and generality of the obtained kinetic subspace by comparing results from propagating rate constant perturbations in the full d -dimensional ($d = 217$) kinetic space and in the reduced r_u -rank ($r_u = 4$) space through the Cantera solver according to Eq. 6. This simultaneously validated the network’s ability to perform the subspace reduction accurately, since the subspace was constructed by the network and is being validated in Cantera. We computed the 2.5 and 97.5 quantiles of the temperature variation around the nominal (unperturbed) value in each case to generate upper and lower temperature bounds, as reported in Fig. 10. These error bounds have excellent agreement of 98% on average between the full (d) perturbation sample and the reduced (r_u) perturbation sample, demonstrating success of the network-powered reduction from the full kinetic space of 217 dimensions in this application down to a more tractable 4-dimensional subspace.

We additionally demonstrate in Fig. 11 that temperature distributions in Cantera from perturbations in the full $d = 217$ kinetic parameter space are uncorrelated between the fuel-lean side and fuel-rich sides of the flame, thus an $r_u = 1$ dimensional subspace would not suffice to capture uncertainty responses across the entire domain. This confirms a result previously discovered using the network. We also remark that the distributions generated from our $r_u = 4$ dimensional subspace accurately represent those in the full parameter space, even across these uncorrelated mixture fractions. Future turbulent combustion work applying this information in the flamelet regime could maintain high accuracy while only needing to quantify uncertainty across four input dimensions instead of 217, indicating substantial savings in computational cost. For example, to propagate methane flamelet uncertainty to LES simulations, Mueller *et al.* found that exploration of the full kinetic space took between 2,000 and 10,000 samples (or even up to 50,000), while exploring a one-dimensional reduced space took seven or fewer [2].

This work demonstrated the feasibility of our novel dimension reduction methodology in a proof-of-concept flamelet example. We note that the methodology is designed to scale up easily to expanded input and output domains. By adding strain rate, equivalence ratio, or other parameters to the input layer of the parameter network and augmenting the training data accordingly, a subspace can be discovered that describes a more comprehensive set of flamelets. By adding neurons to the final output layer and modifying the loss function accordingly, the uncertainty response of species concentrations and other quantities of interest can be analyzed along with temperature. Further work to apply this methodology to a turbulent combustion simulation may adapt the framework of [2]. The reduced subspace would allow for a series of flamelet tables that capture the uncertainty response in an LES

simulation with greatly reduced samples compared to the full kinetic space, while maintaining a verifiable and user-selectable level of accuracy.




Fig. 10: 95% uncertainty bounds in high-temperature region. (a) Nominal temperature values and full sample uncertainty range. (b) Uncertainty ranges predicted with full and reduced subspace inputs centered on nominal temperature. Agreement of 98% on average.




Fig. 11: Joint temperature distributions that show weak correlation at fuel lean and fuel rich Z . Strong agreement between those predicted with the full rate constant space and with the reduced subspace demonstrates the accuracy and mixture fraction generality of the $r_u = 4$ subspace.

4. Conclusions

In this work, we demonstrated a novel methodology for efficient uncertainty quantification in combustion simulations. A two-stage active subspace analysis powered by a specifically designed neural network surrogate model was performed on one-dimensional

nonpremixed methane flamelets with GRI-Mech 3.0 [19]. The propagation of the kinetic uncertainty from the rate constants to the flame temperature profile was found to be more complex than can be described with a single variable (a one-dimensional active subspace) due to lack of correlation across the mixture fraction domain, thus a global kinetic subspace with a dimensionality reduction of two orders of magnitude and demonstrated accuracy was computed.

The accuracy of this reduced global subspace can be finely tuned through choice of the truncation index at each stage, giving users the option to tailor the low dimensional subspaces to meet their own accuracy or computational cost needs. Additionally, this method can be further extended to scale to augmented solution domains by adding further stages of subspace analysis, to include other target quantities such as species profiles by extending the neural network surrogate model, and to discover a universal subspace across various flame conditions by considering strained effects. The flexibility, general applicability, and low computational cost of this method make it a promising tool for efficient uncertainty quantification in large eddy simulations of turbulent combustion in the flamelet regime.

Acknowledgments

BCK is partially supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. WJ and SD would like to acknowledge the funding support by Weichai Power Co., Ltd.

Supplementary material

Supplementary figures summarizing the kinetic composition of the global subspace are available.

References

- [1] T. Lu, C. K. Law, Toward accommodating realistic fuel chemistry in large-scale computations, *Prog. Energy Combust.* 35 (2) (2009) 192–215.
- [2] M. E. Mueller, G. Iaccarino, H. Pitsch, Chemical kinetic uncertainty quantification for Large Eddy Simulation of turbulent nonpremixed combustion, *Proc. Combust. Inst.* 34 (1) (2013) 1299–1306.
- [3] W. Ji, Z. Ren, Y. Marzouk, C. K. Law, Quantifying kinetic uncertainty in turbulent combustion simulations using active subspaces, *Proc. Combust. Inst.* 37 (2) (2019) 2175–2182.
- [4] J. M. Armengol, O. Le Maitre, R. Vicquelin, Bayesian calibration of a methane-air global scheme and uncertainty propagation to flame-vortex interactions, *Combust. Flame* 234 (2021) 111642.
- [5] X. Zhao, H. Kolla, P. Zhang, B. Wu, S. Calello, H. N. Najm, A transported probability density function method to propagate chemistry uncertainty in reacting flow cfd, in: *AIAA Scitech 2019 Forum*, 2019, p. 2007.
- [6] H. Wang, D. A. Sheen, Combustion kinetic model uncertainty quantification, propagation and minimization, *Prog. Energy Combust.* 47 (2015) 1–31.
- [7] P. G. Constantine, E. Dow, Q. Wang, Active Subspace Methods in Theory and Practice: Applications to Kriging Surfaces, *SIAM J Sci Comput* 36 (4) (2014) A1500–A1524.
- [8] W. Ji, J. Wang, O. Zahm, Y. Marzouk, B. Yang, Z. Ren, C. Law, Shared Low-Dimensional Subspaces for Propagating Kinetic Uncertainty to Multiple Outputs, *Combust. Flame* 190 (2018) 146–157.
- [9] T. Lukaczyk, F. Palacios, J. Alonso, P. Constantine, Active Subspaces for Shape Optimization, in: *10th AIAA Multidisciplinary Design Optimization Specialist Conference*, 2014.
- [10] S. Vajda, P. Valko, T. Turányi, Principal component analysis of kinetic models, *Int. J. Chem. Kinet.* 17 (1) (1985) 55–81.
- [11] N. Wang, Q. Xie, X. Su, Z. Ren, Quantification of modeling uncertainties in turbulent flames through successive dimension reduction, *Combust. Flame* 222 (2020) 476–489.
- [12] D. E. Rumelhart, G. E. Hinton, R. J. Williams, Learning representations by back-propagating errors, *Nature* 323 (6088) (1986) 533–536.
- [13] D. G. Goodwin, R. L. Speth, H. K. Moffat, B. W. Weber, Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes, <https://www.cantera.org>, version 2.5.1 (2021).
- [14] J. Duvall, K. Duraisamy, S. Pan, Non-linear Independent Dual System (NIDS) for Discretization-independent Surrogate Modeling over Complex Geometries, *arXiv:2109.07018 [physics]* (2021).
- [15] L. Lu, P. Jin, G. E. Karniadakis, DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators, *Nat. Mach. Intell.* 3 (3) (2021) 218–229.
- [16] O. Zahm, P. G. Constantine, C. Prieur, Y. M. Marzouk, Gradient-Based Dimension Reduction of Multivariate Vector-Valued Functions, *SIAM J Sci Comput* 42 (1) (2020) A534–A558.
- [17] W. Ji, T. Yang, Z. Ren, S. Deng, Dependence of kinetic sensitivity direction in premixed flames, *Combust. Flame* 220 (2020) 16–22.
- [18] D. A. Sheen, X. You, H. Wang, T. Løvås, Spectral uncertainty quantification, propagation and optimization of a detailed kinetic model for ethylene combustion, *Proc. Combust. Inst.* 32 (1) (2009) 535–542.
- [19] G. P. Smith, D. M. Golden, M. Frenklach, N. W. Moriarty, B. Eiteneer, M. Goldenberg, C. T. Bowman, R. K. Hanson, S. Song, W. C. Gardiner, Jr., V. V. Lissianski, Z. Qin, Gri-mech 3.0, http://www.me.berkeley.edu/gri_mech/.
- [20] R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, I. Stoica, Tune: A research platform for distributed model selection and training, *arXiv preprint arXiv:1807.05118* (2018).
- [21] K. He, X. Zhang, S. Ren, J. Sun, Deep Residual Learning for Image Recognition, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [22] S. Elfwing, E. Uchibe, K. Doya, Sigmoid-Weighted Linear Units for Neural Network Function Approximation in Reinforcement Learning, *arXiv:1702.03118 [cs]* (2017).
- [23] D. P. Kingma, J. Ba, Adam: A Method for Stochastic Optimization, *arXiv:1412.6980 [cs]* (2017).