



Contents lists available at SciVerse ScienceDirect

Combustion and Flame

journal homepage: www.elsevier.com/locate/combustflame

Principal component analysis of turbulent combustion data: Data pre-processing and manifold sensitivity

Alessandro Parente^{a,*}, James C. Sutherland^b^a Service d'Aéro-Thermo-Mécanique, Université, Libre de Bruxelles, Bruxelles, Belgium^b Department of Chemical Engineering, University of Utah, Salt Lake City, UT 84112, USA

ARTICLE INFO

Article history:

Received 7 February 2012

Received in revised form 10 July 2012

Accepted 29 September 2012

Available online xxx

Keywords:

Combustion

Low-dimensional manifolds

Outliers

Pre-processing

Principal component analysis

Scaling

ABSTRACT

Principal component analysis has demonstrated promise in its ability to identify low-dimensional chemical manifolds in turbulent reacting systems by providing a basis for the *a priori* parameterization of such systems based on a reduced number of parameterizing variables. Previous studies on PCA have only mentioned the importance of data pre-processing and scaling on the PCA analysis, without detailed consideration. This paper assesses the influence of data-preprocessing techniques on the size-reduction process accomplished through PCA. In particular, a methodology is proposed to identify and remove outlier observations from the datasets on which PCA is performed. Moreover, the effect of centering and scaling techniques on the PCA manifold is assessed and discussed in detail, to investigate how different scalings affect the size of the manifold and the accuracy in the reconstruction of the state-space. Finally, the sensitivity of the chemical manifold to flow characteristics is considered, to investigate its invariance with respect to the Reynolds number. Several high-fidelity experimental datasets from the TNF workshop database are considered in the present work to demonstrate the effectiveness of the proposed methodologies.

© 2012 The Combustion Institute. Published by Elsevier Inc. All rights reserved.

1. Introduction

Recently, principal component analysis (PCA) was introduced as a method of identifying manifolds in turbulent combustion [1]. PCA has also been used by others to analyze combustion data [2–4], but for different purposes – see [1] for a discussion. The merits of PCA in the context of modeling turbulent reacting flows have been demonstrated for identifying low-dimensional manifolds underlying the thermo-chemical state [1,5] and toward the development of PCA-based combustion models [6,7]. A particularly noteworthy feature of PCA-based models is the possibility of obtaining low-dimensional parameterizations satisfying well-defined error bounds. Previous studies on PCA [1,5] have mentioned the importance of pre-processing data prior to applying PCA, but the effects of pre-processing strategies have not been assessed in detail. In particular, the effect of potential outlier observations as well as the role of centering and scaling on the principal component structure has not been addressed. The objective of the present paper is to review the PCA procedure and highlight the role of the available pre-processing techniques on the robustness of PCA and its ability to identify a low-dimensional representation

of a thermo-chemical manifold. The sensitivity of PCA to modifications of the database from which the low-dimensional basis is extracted is also considered, to investigate the universality of the PCA method.

Section 2 provides a review of PCA as well as a discussion on outlier removal (2.1), data centering and scaling (2.2), and dimension reduction (2.3). Section 3 applies PCA to several experimental datasets from the Sandia non-premixed flame datasets to illustrate the effect of pre-processing and scaling on the PCA reduction. Finally, the invariance of the chemical manifold with respect to the Reynolds number is demonstrated for a set of piloted flames at a range of Reynolds numbers.

2. Principal component analysis

Principal component analysis (PCA) [8,9] provides a rigorous mathematical formalism for the identification of the most active directions in multivariate datasets. PCA identifies correlations among the variables defining the state space. As a result, a new coordinate system is identified in the directions of maximal data variance, which allows less important dimensions to be eliminated while maintaining the primary structure of the original data. Details of the PCA reduction have been already provided [1]. Here, the PCA concept will be reviewed briefly whereas the impact of

* Corresponding author. Address: Avenue F.D. Roosevelt 50, 1050 Bruxelles, Belgium. Fax: +32 2 650 27 10.

E-mail address: Alessandro.Parente@ulb.ac.be (A. Parente).

pre-processing and post-processing on PCA results will be discussed in detail.

In PCA, n observations of Q variables are assigned to an $(n \times Q)$ matrix \mathbf{X} whose rows represent individual observations of all Q variables \mathbf{x} . For the combustion applications considered in this paper, the Q columns in \mathbf{X} are taken to be the temperature and species mass fractions.¹ PCA projects \mathbf{x} onto a rotated basis obtained from the eigenvalue decomposition of the $(Q \times Q)$ covariance matrix,

$$\mathbf{S} = \frac{1}{n-1} \mathbf{X}^T \mathbf{X} = \mathbf{A} \mathbf{L} \mathbf{A}^T, \quad (1)$$

where \mathbf{A} and \mathbf{L} are the eigenvectors and eigenvalues of \mathbf{S} . The rotated basis, defined by the eigenvectors \mathbf{A} , may be truncated to retain the most energetic directions (those columns of \mathbf{A} associated with the largest eigenvalues of \mathbf{L}), providing the non-square matrix \mathbf{A}_q on which the original data are projected to obtain the principal components (PC), \mathbf{Z}_q ,

$$\mathbf{Z}_q = \mathbf{X} \mathbf{A}_q. \quad (2)$$

Eq. (2) can be inverted to obtain an approximate reconstruction of the original $(n \times Q)$ dimensional sample:

$$\mathbf{X}_q = \mathbf{Z}_q \mathbf{A}_q^T. \quad (3)$$

Eq. (3) is a linear reconstruction. The intrinsic linearity of the PCA approach represents a major possible drawback of the technique to deal with strongly non-linear processes such as combustion. However, this limitation can be partially overcome using local PCA [1,10]. Alternatively, non-linear reconstructions can provide more accurate mappings from \mathbf{Z}_q to \mathbf{X}_q [7]. The PCA reduction process is represented schematically in Fig. 1.

Several procedures are required prior to performing the PCA reduction process (Fig. 1):

1. **Outlier removal.** Experimental datasets usually contain a few unusual observations which can strongly affect the data covariance structure and, therefore, the structure of the principal components. If we refer to a one-dimensional problem, the outliers can be classified as those observations which are either very large or very small with respect to the others. In high dimensions, there can be outliers that do not appear as outlying observations when considering each dimension separately and, therefore, they will not be detected using univariate criteria. Thus, a multivariate approach must be pursued. PCA itself represents an ideal tool for the identification and removal of outlier observations.
2. **Centering and scaling.** Data are usually centered and scaled before PCA is carried out. Centering represents all observations as fluctuations, leaving only the relevant variation for analysis. Scaling is a crucial operation when analyzing the thermochemical state of a reacting system since temperature and species concentrations have different units and vary over different scales. The choice of scaling significantly affects the subsequent PCA analysis: different scalings allow to emphasize correlations among different groups of state variables, providing an effective tool for targeting the PCA analysis on the variables which are most relevant for an investigated application.

Section 2.1 presents a technique to identify outliers, while Section 2.2 addresses centering and scaling.

¹ Formally, pressure should also be included, but for low mach number flows in open domains, it is safely neglected.




Fig. 1. PCA reduction process.

2.1. Outlier detection and removal with PCA

The usual procedure for outlier detection in multivariate data analysis is to measure the distance of each realization i of the Q observed variables, from the data center, using the so called Mahalanobis distance:

$$D_M = (\mathbf{X} - \bar{\mathbf{X}})^T \mathbf{S}^{-1} (\mathbf{X} - \bar{\mathbf{X}}), \quad (4)$$

where $\bar{\mathbf{X}}$ is a matrix containing the average values, $\bar{x}_j = \frac{1}{n} \sum_{i=1}^n x_{ij}$, of the original variables. The observations associated with large values of D_M are classified as outliers and then discarded. The Mahalanobis distance can be related to the principal components: it can be shown, in fact, that the sum of squares of the PC, standardized by the eigenvalue size, equals the Mahalanobis distance for observation i :

$$\sum_{k=1}^Q \frac{z_{ik}^2}{l_k} = \frac{z_{i1}^2}{l_1} + \frac{z_{i2}^2}{l_2} + \dots + \frac{z_{iQ}^2}{l_Q} = D_{M,i}. \quad (5)$$

This realization can be exploited for building a robust methodology based on PCA for outlier identification and removal. As mentioned previously, the first few principal components have large variances and explain most of the variation in \mathbf{X} . Therefore, those components are strongly affected by variables with relatively large variances and covariances. Consequently, the observations that are outliers with respect to the first few components usually correspond to outliers on one or more of the original variables. On the other hand, the last few principal components represent linear functions of the original variables with minimal variance. These components are sensitive to the observations that are inconsistent with the covariance structure of the data but are not outliers with respect to the original individual variables. Based on the above considerations, the following detection scheme can be proposed, as suggested by [11]:

1. **Multivariate trimming.** A fraction γ of the data points characterized by the largest value of D_M are classified as outliers and removed. $\bar{\mathbf{X}}$ and \mathbf{S} are then computed from the remaining observations. The trimming process can be iterated to ensure that $\bar{\mathbf{X}}$ and \mathbf{S} are resistant to outliers.
2. **Principal components classifier.** The classifier consists of two functions, one from the major, $\sum_{k=1}^q \frac{z_{ik}^2}{l_k}$, and one from the minor principal component, $\sum_{k=q-r+1}^Q \frac{z_{ik}^2}{l_k}$. The first function can easily detect observations with large values on some of the original variables; in addition, the second function helps detect the observations that do not conform to the covariance structure of the sample. The number of major components, q , is determined by retaining the minimum number of PC required to account for at least 50% of the original data variance, while r is chosen so that the minor components used for the definition of the classifiers are those whose variance is less than $0.2 \cdot \bar{l}$, where \bar{l} is the average value of the eigenvalues of \mathbf{S} . This ensures that the selected minor components account for a very marginal variance and they only represent linear relations among the variables. Based on the above definition, an observation \mathbf{X}_i is classified as an outlier if:




Fig. 2. Outlier identification and removal.

$$\sum_{k=1}^q \frac{z_{ik}^2}{l_k} > c_1 \quad \text{or} \quad \sum_{k=Q-r+1}^Q \frac{z_{ik}^2}{l_k} > c_2, \quad (6)$$

where c_1 and c_2 are chosen as the 99th quantile of the empirical distributions of $\sum_{k=1}^q \frac{z_{ik}^2}{l_k}$ and $\sum_{k=Q-r+1}^Q \frac{z_{ik}^2}{l_k}$.

The convergence of the algorithm is verified by looking at the third or fourth order moments of the major principal components. Since the structure of the data is frequently non-normal, the skewness and kurtosis are monitored from one iteration to the other and convergence is achieved when the rate of change of such quantities falls below an *a priori* defined tolerance, (e.g., 10^{-6}) or a maximum number of iterations is reached. A schematic representation of the outlier removal process is shown in Fig. 2.

An example of the outlier detection scheme applied to a dataset consisting of 62,766 observations of 10 state variables [12] is shown in Fig. 3a and b. Outliers were artificially introduced in the experimental data: specifically, 1000 observations have been

generated from a matrix (1000×10) of random numbers between 0 and 1 and scaled using the standard deviation, s_j , of the variables \mathbf{x}_j . The effect of the outliers on the PCs is very clear from Fig. 3a. The introduced outliers (black circles) are mostly outliers with respect to the original variables and they are visible in the plot of the first two PCs: a small cluster of points, separated from the majority of observations, appears in the plot of the first and second PC scores. They are also apparent (although less so) in the plot of the last and second-last scores as observations scattered around the main cloud of points. If the outlier detection scheme is applied (Fig. 3b), the introduced outliers are completely removed; in addition, outliers present in the original experimental dataset, affecting the first and last PC scores (univariate and multivariate outliers), are also detected with the procedure described. A closer look at Fig. 3b also indicates that the elimination of the outliers results in a slight modification of the first two PC scores, which are rotated counter-clockwise and compressed (especially in the \mathbf{z}_2 direction).


Outliers must be treated with care as they can strongly affect the covariance matrix, thus leading to the identification of false PCs.

2.2. Centering and scaling


When the variables are centered and scaled, a reduced variable can be defined as:

$$\tilde{\mathbf{x}}_j = \frac{(\mathbf{x}_j - \bar{\mathbf{x}}_j)}{d_j}, \quad (7)$$

where d_j is the scaling parameter for variable \mathbf{x}_j . PCA, as discussed above, is applied on $\tilde{\mathbf{x}}$ rather than \mathbf{x} . Centering is always applied



(a) Principal component scores showing the original data (red points) and artificial outliers (black circles).



(b) Principal component scores after outlier removal.

Fig. 3. Demonstration of removal of outliers artificially inserted into a dataset and the effect on the resulting PCA structure.

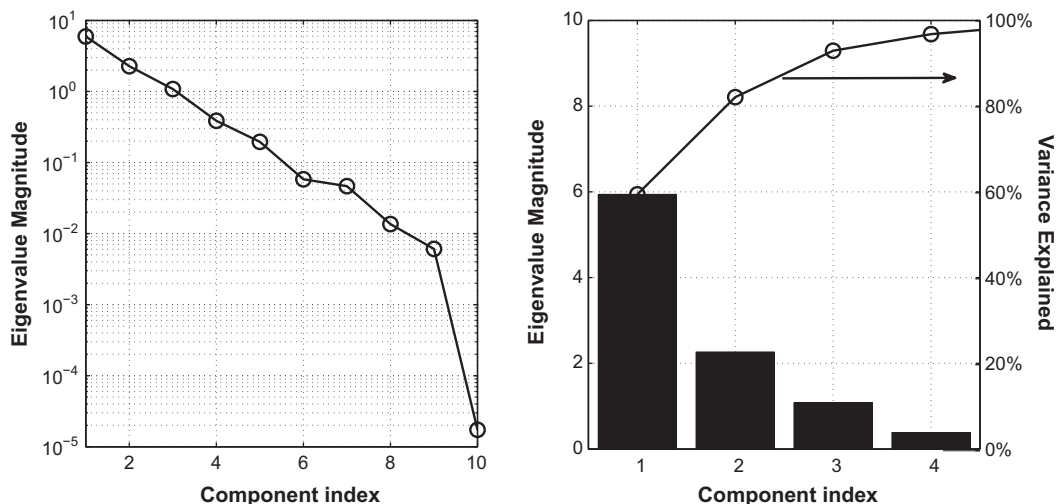


Fig. 4. Scree plot for the determination of the number of principal components. The left frame shows the eigenvalue magnitudes (log scale) while the right frame shows eigenvalue magnitudes (linear scale) along with the associated variance in the data recovered from retaining the given number of eigenvalues.




Fig. 5. Scatter plot of the first two (left frame) and last two (right frame) principal components from the JHC dataset before (black circles) and after (red circles) outlier removal. Scaling method: auto scaling. Trimming fraction, γ : 0.1%. Outliers are indicated by the blue circle.

in conjunction with scaling. We consider the following scaling methods:

1. *Auto scaling*. Also called unit variance scaling, auto scaling uses the standard deviation, s_j , as the scaling factor for each \mathbf{x}_j . After auto scaling, all the elements of \mathbf{X} have a standard deviation equal to one and therefore the data is analyzed on the basis of correlations instead of covariances.
2. *Range scaling*. Range scaling adopts the difference between the minimal and the maximal value, $(\max(\mathbf{x}_j) - \min(\mathbf{x}_j))$, as scaling factor. A disadvantage of range scaling with respect to other scaling methods is that only two values are used to estimate the range, while for the standard deviation all measurements are taken into account. This makes range scaling more sensitive to outliers (see Section 2.1). To increase the robustness of range scaling, the range could also be determined by using robust estimators for maximum and minimum sample values, or after outliers have been removed.
3. *Level scaling*. The mean values of the variables, \bar{x}_j , are used as scaling factors. As with range scaling, level scaling can be affected by outliers. Therefore, a more robust estimator of the mean (the median) could be used or the mean could be determined after outlier removal. Level scaling can be used when large relative changes are of specific interest. However, in the case of the thermochemical state of a system, this could exaggerate the role of chemical species which appear in very small concentrations (e.g., radicals).
4. *Max scaling*. The variables are normalized by their maximum values, $\max(\mathbf{x}_j)$, so that they are all bounded between zero and one. As for the range and level scaling, a robust estimator of maximum values or a procedure for outliers removal should be employed.
5. *VAST scaling* [13]. VAST is an acronym for variable stability scaling and it is an extension of auto scaling. It focuses on variables which do not show strong variation, using the product between the standard deviation and the so-called coefficient of variation, defined as s_j/\bar{x}_j . Such scaling results in a higher importance for variables with a small relative standard deviation.
6. *PARETO scaling* [14]. PARETO scales each variable by the square root of its standard deviation. As a consequence, PARETO gives the variable under evaluation a variance equal to its standard deviation instead of a unit variance, as it happens for auto scaling.

The impact of the different scaling methods will be discussed for several datasets in Section 3.

2.3. Choosing a subset of principal components

The major objective of PCA is to replace the Q elements of $\bar{\mathbf{X}}$ with $q < Q$ principal components, while minimizing information loss. The most obvious criterion for choosing q is to select a cumulative fraction of the total variance that the PCs have to account for. The required number of PCs, q , is then the smallest value of q for

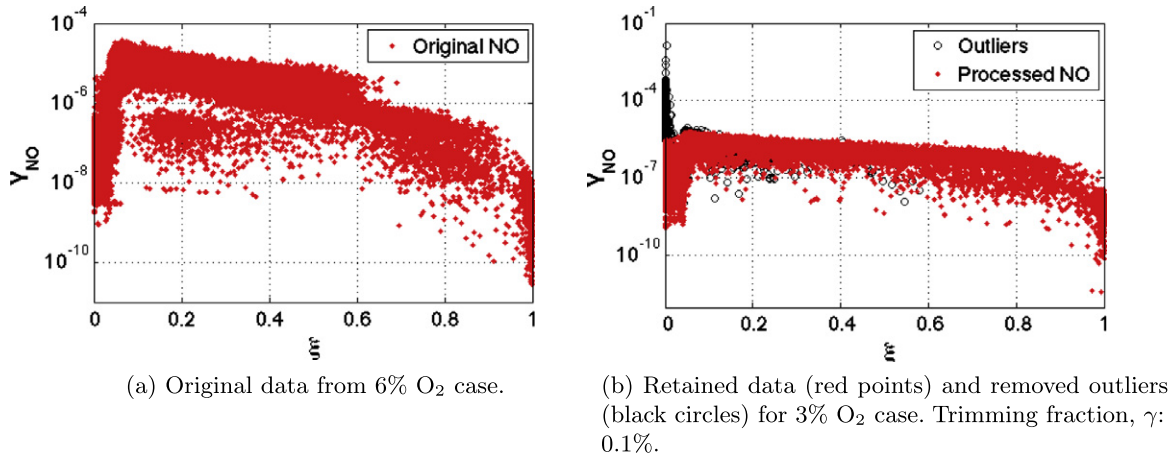


Fig. 6. Scatter plot of NO mass fraction as a function of the mixture fraction. Effect of outlier removal for the 3% O₂ dilution case. JHC dataset [21].




Fig. 7. Weights of the original variables on the first (a) and second (b) principal components before (black bars) and after (gray bars) outlier removal for the JHC data. Scaling method: auto scaling. Trimming fraction, γ : 0.1%.

which this chosen percentage is exceeded. The cumulative variance in the data can be obtained as

$$\sum_{k=1}^Q l_k = \sum_{j=1}^Q \text{var}(\mathbf{x}_j). \quad (8)$$

Then the fraction of the total variance accounted for by retaining q of the Q eigenvectors can be defined as:

$$t_q = \frac{\sum_{k=1}^q l_k}{\sum_{k=1}^Q l_k}. \quad (9)$$

Importantly, it can be shown that the definition of t_q is equivalent to the so-called R^2 value,

$$R^2 = 1 - \frac{\sum_{i=1}^n (x_{q,ij} - \hat{x}_{ij})^2}{\sum_{i=1}^n (x_{ij} - \bar{x}_j)^2}, \quad (10)$$

where $x_{q,ij}$ is the reconstructed i th observation of \mathbf{x}_j . Following the derivation of t_q , an appropriate measure of lack-of-fit of the rank q linear approximation of \mathbf{X} can be related to the size of the discarded eigenvalues, i.e.

$$\epsilon_j = \sum_{k=q+1}^Q l_k = \sum_{i=1}^n \sum_{j=1}^Q (x_{q,ij} - \hat{x}_{ij})^2. \quad (11)$$

For a given number (q) of retained components, it is also possible to determine the variance accounted for each variable by the retained eigenvectors as:

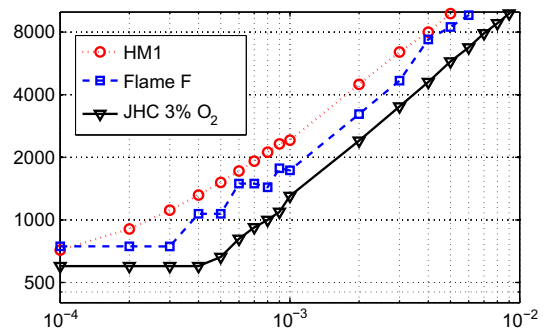


Fig. 8. Effect of γ on the number of points removed during the outlier identification process for the 3% O₂ JHC, flame F and HM1 datasets. Scaling method: auto scaling.

$$t_{q,j} = \sum_{k=1}^q \left(\frac{a_{jk} \sqrt{l_k}}{s_j} \right)^2, \quad (12)$$

where a_{jk} is the weight of the j th variable on the k th eigenvector of \mathbf{S} .

A less rigorous method to identify the number of retained PCs uses a *scree plot*, as shown in Fig. 4 for the jet in hot co-flow dataset, presented in Section 3. This is a simple plot of the eigenvalue magnitudes sorted in descending order against their indexes, and provides a graphical interpretation of the information encoded in each dimension. As previously observed [1,15], there is an exponential decay in the information encoded in each succeeding dimension.

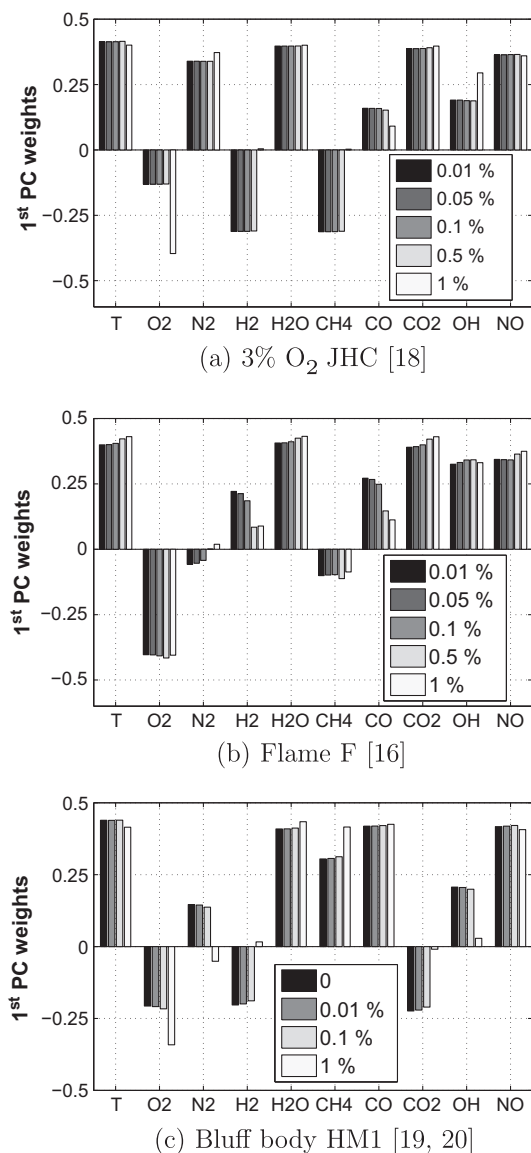


Fig. 9. Effect of the trimming fraction γ on the structure of the first PC. Scaling method: auto scaling.

3. Results

High fidelity experimental data provided under the framework of the Workshop on Measurement and Computation of Turbulent Non-premixed Flames (TNF workshop) [16] are analyzed in the present paper. In particular, the following TNF datasets are employed:

- Turbulent non-premixed CO/H₂/N₂ (0.4/0.3/0.3 by vol.) jet flame [17]. This flame represents an ideal test-case due to its simplicity in terms of turbulence/chemistry interactions.
- Flames C–F, a set of four piloted CH₄ jet flames [12], are characterized by an increasing Reynolds number and exhibit increasing non-equilibrium phenomena, including local extinction and re-ignition.
- The jet in hot co-flow (JHC) burner [18], designed to emulate MILD conditions. It consists of a central fuel jet (80% CH₄ and 20% H₂) within an annular co-flow of hot exhaust products from a secondary burner mounted upstream of the jet exit plane. O₂

in the co-flow is controlled at three different levels, 3, 6 and 9 mol%, while the temperature and exit velocity are kept constant.

- A bluff-body stabilized flame [19,20]. The experimental data used in this paper, designated as HM1, refer to an equimolar mixture of CH₄/H₂ with a fuel velocity of 118 m/s and coflow air velocity of 40 m/s.

“Instantaneous” (as opposed to ensemble-averaged) measurements were used for all analyses presented here.

It should be emphasized that experimental data are “incomplete” in the sense that we do not have simultaneous measurements of all species and temperature, as is possible from computationally obtained data (from, e.g., DNS). In previous studies that employed computational data, we employ all species and temperature in the analysis. However, for the purposes of this paper, this is not an issue, since we are focused on data preprocessing strategies.

In the following, the effect of data pre-processing on PCA results will be discussed in detail, pointing out the possible impact of outliers and scaling methods on the PC structure.

3.1. Effect of outliers on the PCA structure

Outlier detection and removal is particularly important when using PCA with experimental data. If outliers are not removed, the resulting PCA can show significant sensitivity to their existence, thereby complicating interpretation of the PCA.

Figure 5 shows the effect of the outlier removal process on the principal component structure for the JHC dataset, scaled using auto scaling and a trimming fraction γ of 0.1%. In contrast to the example discussed in Section 2.1, the dataset is not augmented with artificial outliers, but processed to identify the existence of experimental measurements inconsistent with the primary structure of the data.

It can be observed that, using the original data without any pre-processing, the scatter plot of the first two principal component scores show the existence of a few observations which strongly differ from the main multi-variate structure of the data. Those can be classified as *univariate outliers* (see Section 3.1), as they correspond primarily to the components associated with the largest eigenvalues. Mathematically, those observations are flagged as outliers because, as explained in Section 2.1, the PC classifier related to the first few PCs is larger than the 99th quantile of the actual PC distribution, indicating that the scores associated with those observations largely deviate from the main data structure. Multivariate outliers are also present in the original dataset, as indicated by the plot of the last two principal components.

To confirm the existence of univariate outliers, NO mass fraction is plotted against the mixture fraction, ξ , for the original data, at 3% and 6% O₂ mass fraction in the co-flow. Figure 6 shows, for the 3% O₂ case, high concentrations of NO on the oxidizer side ($\xi = 0$), which are not observed for the other dilution cases (e.g., 6% O₂, Fig. 6a) and that determine the extreme score values observed in Fig. 5. Figure 6b points out that some “feasible” observations are also removed during the outlier detection process (black circles behind red² dots). This does not affect the statistical value of the analysis since only approximately 1500 out of more than 60,000 observations are removed, and only a few hundred of those are in the feasible NO range.

The eigenvectors (A) of the covariance matrix provide insight into the effect of outliers on the principal components. Figure 7

² For interpretation of color in Figs. 1–3, 5, 6, 8, 10–12, and 15 the reader is referred to the web version of this article.




Fig. 10. Effect of the trimming fraction γ on the temperature distribution (plotted against the mixture fraction, ξ) for the JHC and flame F datasets.

shows a comparison between the first two PCs obtained for the JHC dataset with (gray) and without (black) outlier observations. In particular, the figures graphically indicate the weight of the original variables on the first two components. From Fig. 7, it is clear that, for the JHC dataset, outlier removal results in the chemical species NO being eliminated from the first two PCs, while the remaining weights in **A** remain largely unaffected. This confirms, as indicated by Fig. 6, that the outliers identified in Fig. 5 are related to NO measurements, leading to the overestimation of such species in the PC structure when the dataset is not pre-processed.

The results shown above indicate the relevance that outlier observations, may have on the covariance structure of the data, confirming the need for effective outlier removal tool, as the one employed here and based on PCA.

3.1.1. Effect of the threshold parameter, γ

The trimming fraction γ (see Section 2.1) plays a critical role in the outlier removal process: large values of γ may result too many samples being eliminated, resulting in an unphysical modification of the PC structure. Figure 8 shows the number of removed points as a function of γ for the HM1, 3% O₂ JHC and flame F datasets. Figure 9 shows the effect of γ on the structure of the first PC for the 3% O₂ JHC Fig. 9a, flame F Fig. 9b and HM1 Fig. 9c datasets. The PC structure is relatively constant for $\gamma < 0.05\%$, but begins to change noticeably for $\gamma > 0.1\%$. Flame F and the HM1 were included in such analysis with the JHC dataset as they show specific features which can help identifying appropriate ranges for γ , to avoid over-aggressive outlier removal during the outlier identification process. Indeed, those systems show significant extinction: flame F is close to global extinction [12] and the HM1 bluff-body stabilized flame is known to show intermittent local extinction being at 50% blow-off conditions [19,20] Fig. 9c indicates that $\gamma < 0.05\%$ leads, in all cases, to a PC structure which is unaffected by outlier removal, for all analyzed datasets.

Figure 10 illustrates the effect of over-aggressive outlier removal ($\gamma = 1\%$) on the temperature distribution for the JHC and flame F datasets. Over-aggressive outlier removal eliminates observations corresponding to extinction for the JHC dataset while removing fully-burning regions for the flame F dataset. A similar effect (not shown here) is observed for the HM1 bluff-body dataset, where large γ also results in removal of points corresponding to extinction. Figure 9a indicates that the corresponding PC structure is significantly altered in both cases when the choice of γ is too large.

Figures 9b and 11 indicate that, for flame F, the outlier removal process with an appropriate choice for γ does not significantly impact the PC structure, although some of the realizations inside the

main “data cloud” are removed because they exceed the 99th quantile of the experimental distribution of the first and last eigenvectors. As a consequence, the PC determined before and after the outlier removal procedure show very minor differences among the weights, as shown by Fig. 9b.

Based on the observations above, we recommend $0.01\% < \gamma < 0.05\%$, which effectively removes outliers but does not remove enough physically meaningful datapoints to alter the PC structure.

3.2. Effect of scaling

We now consider the effect of scaling strategies outlined in Section 2.2 on the PCA reduction, focusing on the Sandia CO/H₂ jet flame dataset. Figure 12 shows the normalized³ eigenvalue size distribution obtained by applying the different scaling options. It indicates that the VAST and PARETO scaling methods result in larger weights for the first few eigenvalues while the other scaling options are all very similar in their eigenvalue size distribution. This is a consequence of the high importance given by the VAST and PARETO scaling methods to temperature over the chemical species mass fractions. This effect is accentuated for the PARETO scaling, where the square root of the standard deviation is used to scale variables: this enhances the relevance of temperature with respect to the other variables defining the state-space. Indeed, the application of PARETO scaling results in temperature being the variable carrying most of the data variance and is, therefore, equivalent to forcing the first principal component to align with the temperature. Such behavior is a consequence of the size-dependency of PCA for non-homogeneous datasets (where the variables have very different scales) as is characteristic of combustion. Therefore, the choice of such scaling does not appear very useful for the analysis carried out in the present paper, as it is equivalent to an *a priori* choice of the PC. However, PARETO scaling can be extremely appealing for the definition of reduced-order combustion models, as the choice of temperature within the set of PC has a dramatic influence on the model's accuracy [7].

Table 1 shows t_q and $t_{q,j}$ (see Eqs. (9) and (12)) obtained by applying range, max, VAST and level scaling to the CO/H₂ dataset. Results indicate that auto scaling is the only scaling technique that provides a uniform reconstruction of the state variables (for $q = 3$), as evidenced by relatively high values of $t_{q,j}$ for all variables. Range and max scaling, whose behavior is very similar (as expected), perform slightly better than auto scaling for most of the main species and temperature. However, they cannot properly capture NO variation, even with $q = 3$. Similarly, VAST scaling concentrates on

³ The eigenvalues obtained using different scaling methods are normalized between 0 and 1, to allow comparison of different pre-processing techniques.

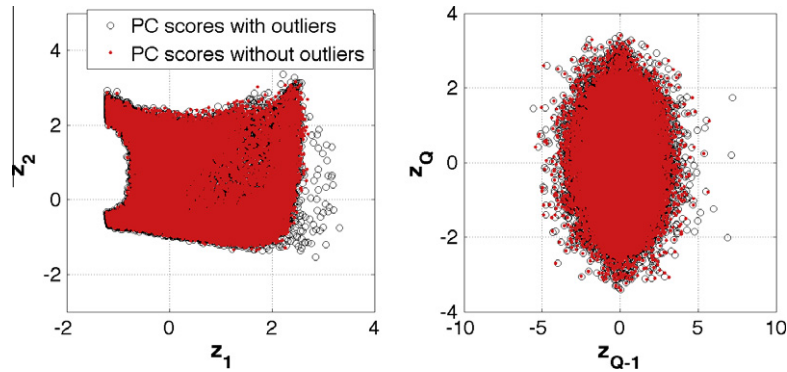


Fig. 11. Scatter plot of the first and last two principal components from flame F before (black circles) and after (red circles) outlier removal. Scaling method: auto scaling. Trimming fraction, γ : 0.5%.

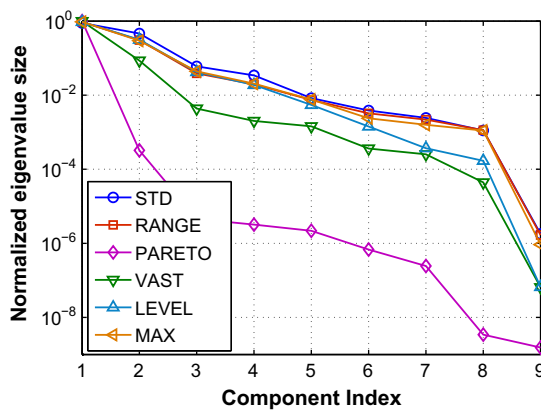


Fig. 12. Ordered normalized eigenvalue magnitudes for the CO/H₂ jet flame dataset for various scaling strategies.

extremely stable variables such as N₂, but fails to recover minor species such as OH properly. This effect is accentuated in PARETO scaling, which clearly emphasizes main species and temperature. The higher values of t_q given by range, max, VAST and PARETO scaling, compared to auto scaling, are due to the higher variance explained for the major variables. However, these scaling approaches do not preserve features related to minor species such as NO and OH. The variance accounted for OH and NO by auto scaling is up to 16% and 25% higher, respectively, than that explained by the other scaling methods. On the other hand, level scaling focuses on variables characterized by large changes (relative to their mean) and leads to an overestimation of the role of minor species in the PCA reduction. Therefore, the reconstruction of minor species such as OH and NO is very accurate, but major species such

as H₂O are poorly recovered. On the basis of the described sensitivity, it becomes clear how scaling can be constructively employed to target the desired accuracy of different subsets of state variables. In particular, auto scaling appears very well-suited when an exploratory analysis on the chemical manifold should be performed, whereas range, max, vast scaling (and PARETO) are useful for capturing the principal features of the systems and the behavior of the main chemical species. This appears very appealing for building reduced-order models of combustion systems to be used in optimization studies.

3.3. PC sensitivity to system variability

We now consider the question of how sensitive a PCA is to the characteristics of a system such as Reynolds number. To investigate this, the four piloted jet flames (Sandia flames C to F) are considered. These flames have increasing Reynolds numbers that lead to significant extinction in flames E and F, which is near blow-out. Given that significantly different regions of state space are realized in these flames (e.g., extinction), one may not think that the PCA structure should be consistent across all flames. Figure 13 shows the weights of the original variables on the first four PCs (columns of **A**) for Sandia flames C–F. The PC structure remains very similar for the first four PCs. The possible exception is weights on intermediate species such as CO, H₂ and OH, which show some variation in their contributions to the eigenvectors across the range of Reynolds numbers. This is a consequence of the increasing degree of extinction which characterizes flames C to F: the OH distribution shows a larger scatter as the Reynolds number is increased. As a result, OH contribution to the covariance matrix is decreased (OH is less correlated with the other state parameters) and the corresponding weights on the PC is reduced.

Table 1
Total, t_q , and individual variance, t_{qj} , (see Eqs. (9) and (12)) accounted for the CO/H₂ jet flame dataset, as a function of the number of retained PC, q , and the scaling option used.

	Auto (std)		Range		Max		VAST		Level		PARETO	
	$q = 2$	$q = 3$	$q = 2$	$q = 3$	$q = 2$	$q = 3$	$q = 2$	$q = 3$	$q = 2$	$q = 3$	$q = 2$	$q = 3$
T	0.971	0.973	0.983	0.991	0.979	0.990	0.992	0.992	0.896	0.943	1.000	1.000
Y_{O_2}	0.986	0.986	0.994	0.994	0.997	0.997	0.975	0.978	0.942	0.961	0.990	0.991
Y_{N_2}	0.986	0.986	0.981	0.981	0.971	0.971	1.000	1.000	0.965	0.970	0.989	0.994
Y_{H_2}	0.968	0.969	0.962	0.963	0.957	0.960	0.945	0.947	0.991	0.991	0.965	0.967
Y_{H_2O}	0.930	0.936	0.945	0.945	0.944	0.944	0.940	0.978	0.870	0.884	0.917	0.968
Y_{CO}	0.994	0.994	0.995	0.997	0.990	0.994	0.979	0.980	0.987	0.987	0.999	0.999
Y_{CO_2}	0.973	0.977	0.979	0.987	0.977	0.988	0.981	0.985	0.908	0.959	0.967	0.993
Y_{OH}	0.738	0.940	0.731	0.991	0.745	0.992	0.660	0.687	0.870	0.993	0.554	0.567
Y_{NO}	0.772	0.930	0.728	0.795	0.729	0.802	0.744	0.970	0.701	0.926	0.759	0.813
t_q	0.924	0.966	0.946	0.975	0.942	0.975	0.992	0.996	0.949	0.973	0.999	0.999

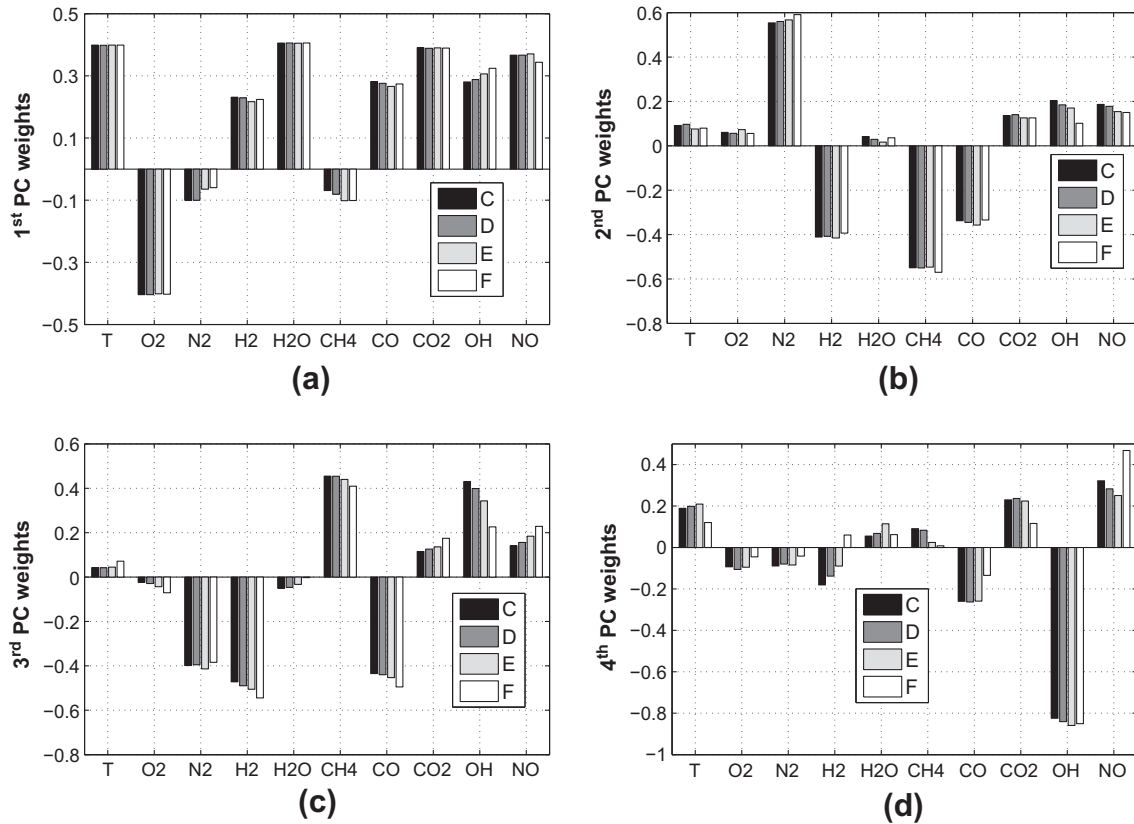


Fig. 13. Structure of the first four PC for PCA applied to flames C–F in the TNF series [16]. Scaling criterion adopted: auto scaling.

Although Fig. 13 indicates that the PCA structure is largely unchanged, the question remains of how sensitive the PCA reconstruction of state variables is to slight variation in the PCA structure. In other words, can the low-dimensional representation obtained via PCA for one of the four systems can be exploited for the others, without performing a new decomposition? This question is crucial to assess the *universality* of the PCA approach for identifying manifolds in reacting systems. To answer this question, a PCA was performed for flames C and F and employed to reconstruct the other datasets. This implies projecting the scores of each system onto a single PC basis (C or F):

$$\mathbf{X}_{q,i} = \mathbf{Z}_{q,i} \mathbf{A}_{q,k}^T \quad k = C, F. \quad (13)$$

Table 2 lists $t_{q,i}$ (the R^2 values for a linear reconstruction) for

1. flames C–F using PCA on each dataset (labeled as $t_{q,i}$)

Table 2

Individual variance, $t_{q,i}$, accounted for the Sandia flame C–F datasets by the PCA reduction, as a function of the number of retained PCs, q . Note that $t_{q,i}^C$ and $t_{q,i}^F$ refer to the accuracy by which variables are reconstructed using the PCA obtained for flame F and flame C, respectively.

	$t_{q,i}$ (%)				$t_{q,i}^F$ (%)			$t_{q,i}^C$ (%)		
	C	D	E	F	C	D	E	D	E	F
T	0.985	0.985	0.976	0.971	0.982	0.981	0.974	0.984	0.977	0.974
Y_{O_2}	0.987	0.986	0.980	0.979	0.983	0.983	0.979	0.986	0.980	0.977
Y_{N_2}	0.982	0.982	0.980	0.980	0.983	0.981	0.980	0.981	0.979	0.979
Y_{H_2}	0.975	0.969	0.964	0.970	0.973	0.966	0.966	0.966	0.964	0.965
Y_{H_2O}	0.989	0.989	0.986	0.984	0.987	0.986	0.984	0.988	0.985	0.984
Y_{CH_4}	0.987	0.987	0.984	0.984	0.986	0.985	0.984	0.986	0.985	0.985
Y_{CO}	0.972	0.968	0.962	0.969	0.970	0.963	0.962	0.964	0.962	0.970
Y_{CO_2}	0.987	0.986	0.976	0.974	0.983	0.980	0.975	0.985	0.977	0.975
Y_{OH}	0.999	0.999	0.995	0.978	0.990	0.990	0.984	0.999	0.999	0.998
Y_{NO}	0.945	0.932	0.887	0.892	0.942	0.931	0.895	0.933	0.877	0.850

2. flames C–E using the PCA obtained from flame F (labeled as $t_{q,i}^F$).
3. flames D–F using the PCA obtained from flame C (labeled as $t_{q,i}^C$).

Results indicate that the low-dimensional representation found for flame F provides a very satisfactory representation of the other systems. In all cases, the relative error with respect to an optimal reconstruction ($t_{q,i}$ versus $t_{q,i}^F$) is less than 1%. When the basis found for flame C is employed, a very interesting result is observed: the reconstruction of most state variables slightly improves and the accuracy in NO reconstruction decreases. This is probably due to the increasing degree of extinction determined by the increase of Re , which leads to a large variability of NO species, as shown in Fig. 14. As a consequence, the basis identified directly from the flames reflect the major variability of NO, leading to larger weights on the first components. This is not the case when the basis is extracted from flame C, leading to less accurate NO predictions.




Fig. 14. NO distribution with increasing Reynolds number (from left to right) for flames C–F.

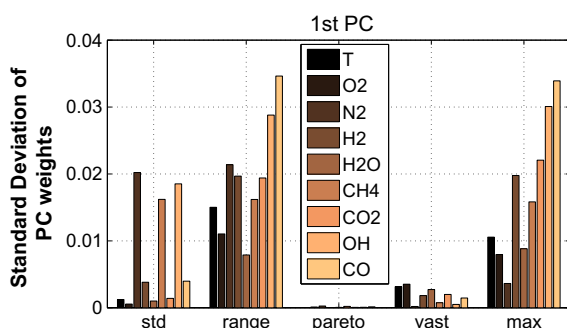


Fig. 15. Standard deviation of the weights on the first principal component across flames C–F when applying different scaling methods.

3.3.1. Effect of scaling on manifold invariance

For completeness, we also consider the effect of scaling on the manifold invariance. Figure 15 shows the standard deviation of the first PC weights on each original variable considering PCA performed on flames C–F independently with different scaling methods. Large standard deviations indicate an alteration in the PCA structure across flames C–E when using the given scaling method. In all cases, there is not a significant variation in the PC structure (with the standard deviation remaining below 0.04 in all cases), but they appear extremely stable for VAST and PARETO scaling methods, which emphasize major and stable variables of the state-space.

The results shown in this section indicate the potential of exploiting a PCA-reduced representation even when the system characteristics are modified. In particular, the relative independence of the basis on the Reynolds number indicate the invariability of the manifold in a range of operating conditions. Nevertheless, further study considering more systems over wider ranges of Re is warranted before concluding that PCA is entirely independent of Re .

4. Conclusions

PCA has recently been proposed as a technique to identify correlations among the multivariate datasets ubiquitous to turbulent combustion. These correlations imply the existence of manifolds in the chemically reactive systems, and PCA has shown promise in identifying these manifolds [1,5–7]. This paper has explored the details of data pre-processing for use in PCA. Specifically, scaling and centering the data as well as outlier removal have been discussed.

The existence of outliers in the dataset can significantly alter the determination of the PC structure and this can lead to the

overestimation of the role of specific variables, or sets of variables, for which outlier observations exist. A method based on PCA has proved very satisfactory for the elimination of the observations which differ from the main multi-variate structure, based on PC classifier built from the first and last few PC, respectively. The effectiveness of the approach was proven for different systems, including a jet in hot co-flow, a set of piloted flames and a bluff-body stabilized burner.

The choice of scaling in particular has a significant impact on the resulting PCA structure by altering the relative importance of various species and temperature. Indeed, different scaling choices may be made depending on the goal of the resulting PCA to optimize the reconstruction of specific classes of state variables. In particular, auto-scaling appears the best option where a balanced reconstruction of the state-space is required for exploratory analysis, whereas level scaling enhances the role of minor species. All the other tested scalings (range, max, VAST) appear ideal for the optimization of stable and major species.

Finally, for the TNF flame datasets, we have demonstrated that the PCA structure remains nearly invariant with Reynolds number across the range from flame C to flame F. This observation is further substantiated by the fact that reconstructing flame C data from a PCA obtained on flame F (or vice versa) is nearly as accurate as reconstructing data from a PCA obtained directly on that dataset.

Acknowledgments

The research was sponsored by the National Nuclear Security Administration under the Accelerating Development of Retrofittable CO_2 Capture Technologies through Predictivity Program through DOE Cooperative Agreement DE-NA0000740 and by National Science Foundation PetaApps award 0904631.

References

- [1] A. Parente, J.C. Sutherland, L. Tognotti, P. Smith, Proc. Combust. Inst. 32 (2009) 579–1586.
- [2] U. Maas, D. Thévenin, Proc. Combust. Inst. 27 (1998) 1183–1189.
- [3] C.E. Frouzakis, Y.G. Kevrekidis, J. Lee, K. Boulouchos, A.A. Alonso, Proc. Combust. Inst. 28 (2000) 75–81.
- [4] S.J. Danby, T. Echekki, Combust. Flame 144 (2006) 126–138.
- [5] A. Parente, J.C. Sutherland, B.B. Dally, L. Tognotti, P.J. Smith, Proc. Combust. Inst. 33 (2) (2011) 3333–3341.
- [6] J.C. Sutherland, A. Parente, Proc. Combust. Inst. 32 (2009) 1563–1570.
- [7] A. Biglari, J.C. Sutherland, Combust. Flame.
- [8] I.T. Jolliffe, Principal Component Analysis, Springer, New York, NY, 1986.
- [9] E.J. Jackson, A User's Guide to Principal Components, Wiley, New York, NY, 1991.
- [10] Axel Coussement, Olivier Gicquel, Alessandro Parente, MG-local-PCA method for reduced order combustion modeling, Proceedings of the Combustion Institute, Available online 30 June 2012, ISSN 1540-7489, <http://dx.doi.org/10.1016/j.proci.2012.05.073>.
- [11] M. Shyu, S. Chen, K. Sarinapakorn, L. Chang, A novel anomaly detection scheme based on principal component classifier, in: Proceedings of the IEEE

- Foundations and New Directions of Data Mining Workshop, in conjunction with the Third IEEE International Conference on Data Mining, 2003, pp. 172–179.
- [12] R.S. Barlow, J.H. Frank, *Proc. Combust. Inst.* 27 (1998) 1087–1095.
- [13] H.C. Keun, T.M.D. Ebbels, H. Antti, M.B. Bollard, O. Beckonert, B. Holmes, J.C. Lindon, J.K. Nicholson, *Anal. Chim. Acta* 490 (2003) 265–276.
- [14] I. Noda, *J. Mol. Struct.* 883–884 (2008) 216–227, <http://dx.doi.org/10.1016/j.molstruc.2007.12.026>. <<http://www.dx.doi.org/10.1016/j.molstruc.2007.12.026>>.
- [15] A. Parente, *Experimental and Numerical Investigation of Advanced Systems for Hydrogen-Based Fuel Combustion*, Ph.D. Thesis, Università di Pisa, 2008.
- [16] Turbulent Nonpremixed Flames Workshop <<http://www.ca.sandia.gov/TNF/abstract.html>>.
- [17] R.S. Barlow, G.J. Fiechtner, C.D. Carter, J.-Y. Chen, *Combust. Flame* 120 (2000) 549–569.
- [18] B.B. Dally, A.N. Karpetis, R.S. Barlow, *Proc. Combust. Inst.* 29 (2002) 1147–1154.
- [19] B.B. Dally, F.D. F. A.R. Masri, *Combust. Theory Modell.* 2 (1998) 193–219.
- [20] B.B. Dally, A.R. Masri, R.S. Barlow, G.J. Fiechtner, *Combust. Flame* 114 (1–2) (1998) 119–148.
- [21] B.B. Dally, E. Riesmeier, N. Peters, *Combust. Flame* 137 (2004) 418–431.