Proceedings of the Combustion Institute 40 (2024) 105608

Contents lists available at ScienceDirect

Proceedings of the Combustion Institute

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

A multi-fidelity framework for developing digital twins of combustion systems from heterogeneous data: Application to ammonia combustion

Aysu Özden [a,b,c,*], Matteo Savarese [a,b,d], Lorenzo Giuntini [a,b,e], Alberto Procacci [a,b], Riccardo Malpica Galassi [f], Axel Coussement [a,b], Francesco Contino [c], Alessandro Parente [a,b,g]

[a] Aero-Thermo-Mechanic Department, Université Libre de Bruxelles, Brussels, Belgium
[b] Brussels Institute for Brussels Thermal Energy (BRITE), ULB and VUB, Belgium
[c] Institute of Mechanics, Materials, and Civil Engineering, UC Louvain, Louvain-la-Neuve, Belgium
[d] Service de Thermique et Combustion, Université de Mons, Mons, Belgium
[e] Department of Civil and Industrial Engineering, University of Pisa, L.go Lucio Lazzarino, 56122, Pisa, Italy
[f] Mechanical and Aerospace Engineering Department, Sapienza University of Rome, Roma, Italy
[g] WEL Research Institute, Avenue Pasteur, 6, 1300, Wavre, Belgium

ARTICLE INFO

Keywords:
Chemical reactor network
CoKriging
Manifold alignment
Multi-fidelity
Proper orthogonal decomposition
Reduced order model

ABSTRACT

In the present work, we introduce a multi-fidelity reduced-order model (MF-ROM) framework that effectively blends high-fidelity evaluations (accurate but expensive) with low-fidelity ones (approximate, less expensive) to predict the thermo-chemical state at unexplored operating conditions. The methodology combines Proper Orthogonal Decomposition (POD) for data compression, manifold alignment for blending information from high- and low-fidelity data, and CoKriging for regression. To assess the methodology, two-dimensional Reynolds-Averaged Navier Stokes (RANS) Computational Fluid Dynamics (CFD) simulations, along with a Chemical Reactor Network (CRN) derived from these simulations, are employed to develop the MF-ROM to predict the spatial fields of thermo-chemical variables at unexplored design conditions. Results show that the MF-ROM attains competitive predictive accuracy against the single-fidelity ROM built with only high-fidelity data for temperature and main chemical species distribution while considerably lowering the computational costs. This new framework allows to predict unexplored scenarios in a wide range of conditions, proving useful in preliminary design explorations, troubleshooting and addressing what if scenarios for a limited computational budget by incorporating simulations of different fidelities.

1. Introduction

The design of innovative combustion systems is a cornerstone of the energy transition, aiming to facilitate the efficient use of carbon-free fuels to reduce greenhouse gas emissions in those sectors that are challenging to electrify. The use of numerical models and simulations is a crucial aid in developing new technologies, especially in combustion applications, given the advances in detailed modeling techniques such as computational fluid dynamics (CFD), both in terms of the quality of predictions and the increased availability of computational resources. However, particularly with alternative fuels and their mixtures, such as hydrogen and ammonia, the operating conditions to be explored can be a prohibitively wide range. In addition, since the prediction of minor pollutants is crucial for characterizing the environmental performance of combustion systems, CFD simulations with sufficiently detailed kinetics must necessarily be employed. This ensemble of factors raises the overall computational cost of these simulations [1], making CFD an unsuitable tool for tasks such as parametric studies, optimization, or uncertainty quantification, where many model evaluations are required.

With the emergence of data-driven techniques, the development of reduced-order models (ROMs) is attracting the interest of the scientific and industrial community [2]. The goal is often to create, given a parsimonious amount of data, a digital twin of a system, namely a function that maps the input–output conditions of such a system, that allows obtaining an instantaneous response [3]. The ROM is trained on a minimal amount of detailed simulations or experimental results, whose quantity largely depends on the variability of the system's response to the input conditions. Therefore, for highly non-linear systems, such as combustion ones, and when the input parameters are many and have a wide range, several full-order model samples are required, thereby involving a considerable allocation of resources.

* Corresponding author at: Aero-Thermo-Mechanic Department, Université Libre de Bruxelles, Brussels, Belgium.
E-mail address: aysu.ozden@ulb.be (A. Özden).

https://doi.org/10.1016/j.proci.2024.105608
Received 4 December 2023; Accepted 5 July 2024
Available online 16 August 2024

To this end, the integration of computationally efficient but lower-fidelity models into ROM training is certainly an attractive perspective to economize and optimize the computational budget [4]. Various levels of fidelity might be achieved in several ways, i.e. changing computational grid size or using alternative numerical tools [5]. The development of multi-fidelity reduced-order models (MF-ROM) is based on finding a compromise between the number of high-fidelity (HiFi) samples and low-fidelity (LoFi) ones in order to ensure due model accuracy and, at the same time, to save resources. The core idea behind data fusion in MF-ROMs is that the model learns the discrepancy between LoFi and HiFi data in the locations where both are available, and it is informed with suited confidence by the LoFi data where the HiFi ones are not available.

Researchers have been using multi-fidelity methods in different areas in the past decade. A variable-fidelity surrogate model was developed to produce aerodynamic data and shape optimization on an airfoil [6]. Perron et al. [7] implemented a non-intrusive MF-ROM to assess the performance on a transonic airfoil to show that the computational training cost was reduced with comparable accuracy counter to the single-fidelity approach. For a more thorough exploration of applications employing multi-fidelity in different research fields, including uncertainty propagation, inference, and optimization, the reader is referred to [8].

In combustion, Chemical Reactor Networks (CRNs) are often utilized to perform lower-fidelity chemically reacting flow simulations by schematizing a complex flow field as a series of interconnected ideal reactor models [9]. They have been successfully applied to gas turbines and industrial furnaces to predict pollutants in a computationally efficient manner [10–13]. The benefits of CRNs extend beyond merely lowering computational expenses; they also represent physics-based models capable of accommodating detailed kinetics. However, since CRN models represent the complex flow fields by means of 0-D or 1-D models, a loss in model accuracy is unavoidable but, on the other hand, it can be quantified [14,15]. Hence, these characteristics make CRN models an ideal candidate for MF-ROM applications.

We introduce a novel methodology where the LoFi model derived from CRN simulations is used in conjunction with HiFi data from CFD simulations, particularly Reynolds-Averaged Navier Stokes (RANS) simulations, to generate multi-fidelity ROMs for a prototype combustor designed for the multi-stage combustion of ammonia. The methodology relies on Proper Orthogonal Decomposition (POD), which seeks to project both HiFi and LoFi data onto sets of directions where most of their data variance is captured. The multi-fidelity nature of the approach is introduced by blending HiFi and LoFi-generated manifolds through a manifold alignment technique. Finally, the obtained modes are mapped as a function of the input design parameters using advanced regression tools, namely Co-Kriging, from which the high-fidelity response of the system can be recovered at any location within the design space.

2. Methods

2.1. Proper orthogonal decomposition

Proper Orthogonal Decomposition, also known as Karhunen–Loéve transformation [16] or Principal Component Analysis [17], is a dimensionality reduction technique applied to the data matrix $\mathbf{X} \in \mathbb{R}^{p\times n}$ assuming that it has a large number of interrelated variables [17], where p represents the number of features, i.e., number of computational cells times the number of Quantity of Interest (QoIs, e.g., temperature and species), $p = n_{cells} \times n_{QoIs}$, and n is the number of observations corresponding to n design/operating conditions. Dimensionality reduction is obtained by retaining only the first k eigenvectors, i.e., defining the truncated matrices $\boldsymbol{\Phi}_k$, which is called POD modes, or eigenflames [18], and $\mathbf{Z}_k$ which is called the POD coefficients matrix, with dimensions $(p \times k)$ and $(k \times n)$, respectively, where $k < min(p, n)$ (for typical combustion problems $p \gg n$). The dimension k of the latent subspace is obtained by calculating the Singular Value Decomposition (SVD) of the matrix $\mathbf{X}$, which corresponds to solving the eigenproblem of the covariance matrix $\mathbf{C} = \frac{1}{n-1}\mathbf{X}^T\mathbf{X}$. Relative Information Content (RIC) criteria [19], which can be interpreted as the total variance captured by the POD, is used to determine the value of k. The matrix $\mathbf{X}$ can be approximated by using the truncated matrices $\boldsymbol{\Phi}_k$ and $\mathbf{Z}_k$ as:

$$\mathbf{X} \approx \boldsymbol{\Phi}_k \mathbf{Z}_k^T \tag{1}$$

where modes are ordered for decreasing energy content, i.e., singular value. As the eigenmodes of $\mathbf{C}$ are orthogonal and the principal components are linear combinations of the observed features, it is possible to reconstruct the high-fidelity response $\mathbf{x}_{pred}(\boldsymbol{d}^*) = \boldsymbol{\Phi}_k\ \boldsymbol{\alpha}(\boldsymbol{d}^*)$ for any unexplored design condition $\boldsymbol{d}^*$ where $\boldsymbol{\alpha}$ is the vector of the k scalar regression models built using the POD coefficients over the design space.

2.2. Manifold alignment

In a multi-fidelity framework, challenges may arise in finding a common low-dimensional representation that enables mapping, ensuring the preservation of the inherent structure within each dataset due to topological or dimensionality differences. Manifold alignment provides a solution by aligning the underlying low-dimensional representations [20], assuming a common low-dimensional subspace is shared between these datasets due to their intrinsic similarity [7]. The Procrustes manifold alignment technique is one of the methods that utilizes dimensionality reduction before performing Procrustes analysis [21] to adjust the translational, rotational, and scaling parameters to accomplish an optimal alignment between these manifolds. The HiFi and LoFi datasets are identified as $\mathbf{X}\in\ \mathbb{R}^{p\times n}$ and $\mathbf{Y}\in\ \mathbb{R}^{q\times m}$, respectively, with $n < m$. The value of n must be smaller than m since the fundamental principle behind building an MF-ROM is using scarcer HiFi data (n) and plentiful LoFi data (m) to reach acceptable prediction accuracy. The LoFi datasets are organized to include a linked subset $\mathbf{Y}_L \in \mathbb{R}^{q\times n}$ with shared input design parameters identical to those of the HiFi dataset, while the remaining data forms the unlinked subset $\mathbf{Y}_U \in \mathbb{R}^{q\times(m-n)}$, thus $\mathbf{Y} = [\mathbf{Y}_L, \mathbf{Y}_U]$. Applying dimensionality reduction on both datasets provides the eigenflames $\boldsymbol{\Phi}\ \in \mathbb{R}^{p\times k}$ and $\boldsymbol{\Psi}\ \in \mathbb{R}^{q\times k}$, and POD coefficients $\mathbf{Z} \in \mathbb{R}^{k\times n}$ and $\mathbf{W}\in\ \mathbb{R}^{k\times m}$, for the HiFi and LoFi datasets, respectively. Thereupon, the LoFi POD coefficients $\mathbf{W}$ has also linked and unlinked subsets, $\mathbf{W} = [\mathbf{W}_L, \mathbf{W}_U]$. The orthogonal Procrustes theorem finds the manifold $\mathbf{T}$ minimizing $\|\mathbf{Z} - \mathbf{T}\|$ and the solution of this problem is performed by using the procedure of Wang and Mahadevan [22]. Prior to performing manifold alignment, translation requires both $\mathbf{Z}$ and $\mathbf{W}_L$ to be centered at the origin. Rotation and scaling are achieved by calculation of the scaling factor (s) and rotation matrix ($\mathbf{Q}$). These parameters are obtained through computation of SVD, $\mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T = \mathbf{W}_L\mathbf{Z}^T$, where $\mathbf{U} \in \mathbb{R}^{k\times k}$, $\mathbf{V} \in \mathbb{R}^{k\times k}$, and $\boldsymbol{\Sigma} \in \mathbb{R}^{k\times k}$ are the left and right singular vectors, and the diagonal matrix containing the singular values, respectively. Then, the scaling factor and rotation matrix are calculated as:

$$\mathrm{s} = \frac{\mathrm{tr}(\boldsymbol{\Sigma})}{\mathrm{tr}(\mathbf{W}_\mathrm{L}\mathbf{W}_\mathrm{L}^\mathrm{T})} \tag{2}$$

$$\mathbf{Q} = \mathbf{V}\mathbf{U}^T \tag{3}$$

Finally, the transformed matrix that is a representation of the LoFi data, $\mathbf{Y}$, in the latent space of the HiFi data, $\mathbf{X}$, is obtained as:

$$\mathbf{T} = s\,\mathbf{Q}\mathbf{W} \tag{4}$$

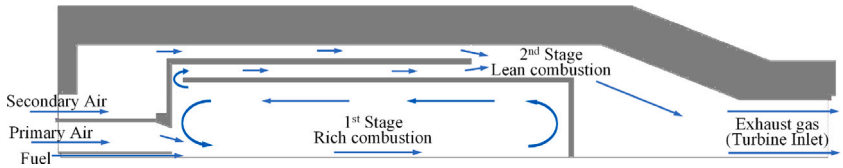

Fig. 1. Schematic view of the mGT-like combustor investigated in this work [15].

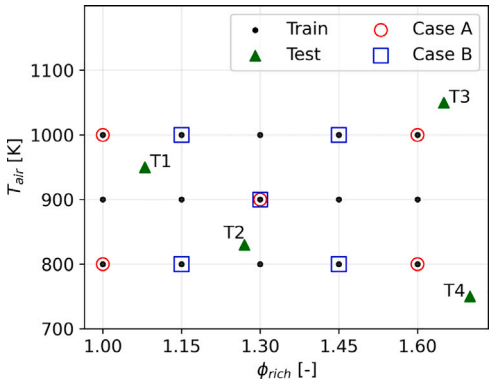

Fig. 2. Design space considered in this work. The full black dots represent the training cases, while the full triangles represent the test cases with two interpolation cases, namely T1 and T2, and two extrapolation cases, namely T3 and T4. The empty red and blue circles represent the two employed initial configurations, namely case A (red circles) and case B (blue squares). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.3. Regression model: CoKriging

The Gaussian Process Regression (GPR) [23], also known as Kriging is a very popular regression technique that models the complex relationships between variables, and when new input data are introduced, it predicts a distribution over possible values which includes a mean value and a measure of uncertainty. However, Kriging is a single-variable model, and the nature of a multi-fidelity framework requires a model that can predict one variable by combining the two correlated variables. CoKriging [24–26], which is an extension of Kriging, is suitable for these situations where one variable is measured with better accuracy and it improves the prediction quality by exploiting the spatial relations between correlated variables. The CoKriging model used for this work follows the autoregressive model of Kennedy and O'Hagan [24] with the following assumption:

$$\mathrm{cov}\{\alpha_{\mathrm{HiFi}}(\boldsymbol{d}), \alpha_{\mathrm{LoFi}}(\boldsymbol{d}')|\alpha_{\mathrm{LoFi}}(\boldsymbol{d})\} = 0 \quad \text{for all } \boldsymbol{d} \neq \boldsymbol{d}' \tag{5}$$

This Markov property indicates that the LoFi data cannot improve the prediction accuracy if HiFi data are available at the same training point [27]. The CoKriging works as if building two Kriging models in sequence. The first Kriging model is constructed by using the LoFi data only, and the second model is built on the residuals of the HiFi and LoFi data:

$$\alpha_{\mathrm{HiFi}}(\boldsymbol{d}) = \rho\, \alpha_{\mathrm{LoFi}}(\boldsymbol{d}) + \delta(\boldsymbol{d}), \tag{6}$$

where ρ is a scaling factor, and δ is the modeling of the difference between the LoFi and HiFi data.

2.4. Incremental sampling strategies

The role of parameter space sampling is key for achieving a satisfactory balance between accuracy and cost. In the case of MF-ROMs, the amount and distribution of both HiFi and LoFi samples need to be determined. A-priori sampling strategies have often proved to be inefficient due to a lack of knowledge about the system response. To this end, adaptive sampling strategies enable the progressive addition of training points only in strategic locations of the design space that improve the ROM predictivity. In this work, we always keep the full amount of LoFi data and we add the HiFi simulations progressively to the MF-ROM training set by using two different incremental strategies: the first one is heuristic and adds samples in the design space locations where the MF-ROM variance is larger (variance-based sampling), assuming that where variance is larger, additional HiFi information is needed; the second one is a trial-and-error strategy which finds the location that minimizes the prediction error (error-based sampling).

3. Test case

The system investigated in this study is a stagnation-point reverse flow combustor, sketched in Fig. 1, which consists of a first, rich, and non-premixed combustion stage that is followed by a secondary lean combustion stage. The design space is composed of the air inlet temperature (T_{air}) and the equivalence ratio in the fuel-rich stage (ϕ_{rich}). These parameters range between 800K–1000K for the temperature and 1.0–1.6 for the equivalence ratio (Fig. 2) . The exhaust temperature at the combustor outlet is set to 1200 K and used as a design constraint to infer the air mass flow rate for the CFD simulations via an energy balance accounting for the nominal input power, and the pressure is fixed at 5 bar, as typically seen in micro gas turbine applications [15].

3.1. CFD model (HiFi)

The 2D numerical simulations were performed using the commercial software ANSYS Fluent 19.1. Due to the symmetry of the setup, for the 2D simulations, the cross-section of a cylindrical representation of the combustor was chosen. A total of 19 steady-state Reynolds-Averaged Navier Stokes (RANS) simulations were carried out. The training set includes 15 of such simulations, placed within the parameter space, with a 0.15 spacing for ϕ_{rich} and a 100 K spacing for T_{air} (black markers in Fig. 2). The other 4 simulations serve as a test set, 2 inside (T1 and T2) and 2 outside (T3 and T4) of the training space (green triangle markers in Fig. 2). The closure of Reynolds Stresses was achieved using the Realizable k-ϵ model with enhanced wall treatment. Radiation was modeled using the Discrete Ordinates (DO) model. Turbulence-chemistry interactions, specifically to close the mean reactive source term, were handled with the Eddy Dissipation Concept (EDC). Otomo et al.'s [28] detailed kinetic mechanism was used for the ammonia oxidation. The 2D axisymmetric computational domain consists of 50,000 triangular/quadrilateral elements with five prism layers at the walls. A uniform boundary condition was applied at the inlets for the mass flow rate, temperature, and molar composition. Each of the numerical simulations required approximately 18 h on a 20-core CPU (Intel® Xeon® processor E5-2630 v4).

3.2. CRN model (LoFi)

The combustor was also modeled by means of a compact Chemical Reactor Network (CRN) composed of 6 ideal chemical reactor models, with the structure shown in Fig. 3. The CRN was previously calibrated to match the output average value of NO from the CFD

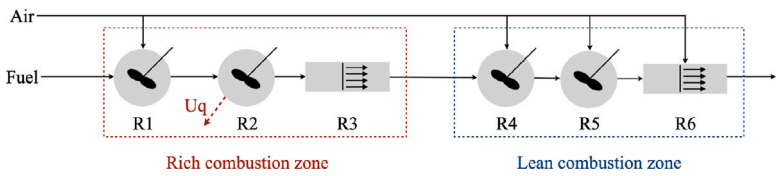

Fig. 3. Scheme of the Chemical Reactor Network of the combustor [15].

simulations [15]. In this work, the calibrated CRN model was employed as the low-fidelity counterpart. CRN simulations were carried out using the NetSMOKE++solver [29] at the same 15 locations in parameter space (black circles in Fig. 2). The values of the QoIs at each grid point of the 2D flow field are obtained through the point-wise prediction in each reactor being projected onto each CFD grid point. The grid point is chosen as the one that minimizes the discrepancy between the CFD values in that cell and all the reactors.

3.3. MF-ROM setup

The MF-ROM was built to reconstruct the temperature and species concentration fields (NO, NH_3, H_2, O_2, OH, H_2O, NO_2, NH_2, and N_2) predicted by the CFD simulations in the combustor. To build the HiFi and LoFi datasets, both models were evaluated at every training point, resulting in 15 observations for each dataset which are shown as black dots in Fig. 2. Additionally, for testing purposes, HiFi observations were also obtained for both interpolation (T1 and T2) and extrapolation test cases (T3 and T4). With reference to the definitions in Section 2.1, the dataset dimensions are $p = q = n_{cells} \times n_{QoIs} = 474800$, $m = 15$, and $n = 5, 6, .., 15$ as the number of HiFi observations increases in each new iteration of the sampling strategies. The results are shown with increasing numbers of HiFi simulations, starting with a subset (seed) of 5 CFD simulations. This is a requirement for the Co-Kriging regression model, while the LoFi simulations are always kept to 15. We emphasize that the choice of the initial subset of HiFi simulations influences the model performances. Therefore, we assigned two alternative initial simulation subsets, as depicted in Fig. 2 by red circles for case A and blue squares for case B. The first subset *(case A)* contains the four corners of the design space and the simulation at the center ($T_{air} = 900$ K, $\phi_{rich} = 1.30$). The second subset *(case B)* includes the four corners of the inner rectangular region plus the simulation at its center. Case B corresponds to the scenario where an initial Design of Experiments (DoE) is extended to include conditions that were not originally planned.

4. Results

We first show the sensitivity of the MF-ROM with respect to the variance cut-off, which determines the dimensionality k of the HiFi latent space. Next, we show how the two sampling algorithms behave in order to strategically determine the next HiFi simulation to be included in the MF-ROM training. Lastly, we examine both the interpolation (on T1 and T2) and the extrapolation capacity of the model (on T3 and T4).

4.1. Sensitivity to cut-off variance

We assessed the effect of the variance cut-off, i.e., the percentage value that determines the number of modes k in Eq. (1), on the model performance. Here, the MF-ROM is trained with an increasing number of HiFi simulations and tested only on interpolation test cases T1 and T2 for case A, so that the test cases are always within the initial design space boundaries. The results in the form of averaged normalized root-mean-squared error ($\mathcal{E}_{\text{NRMS}}$) across all QoIs for the 2 test cases are shown in Fig. 4 for cut-off variance ranging from 95% to 100%. We observe that the effect of the cut-off variance threshold on prediction accuracy is smaller than that of the addition of HiFi simulations until

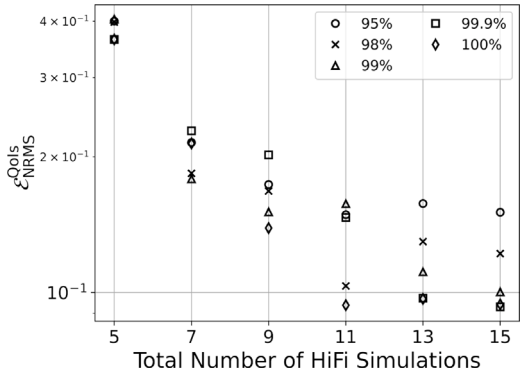

Fig. 4. The average NRMSE of QoIs ($\mathcal{E}^{QoIs}_{\text{NRMS}}$) for interpolation test data (T1 and T2) for different total variance values for case A.

a saturation happens after 11 HiFi simulations and no further accuracy improvement is observed even with 100% variance. We conjecture that when the total number of HiFi simulations is small, the loss of information for lower cut-off variances is negligible because only the mostly linear features are captured by the MF-ROM (e.g., temperature and major species), thus only the first few modes are actually needed. In some cases, such as # 7, the addition of the last modes is even detrimental. However, as the dimensionality of the training data increases with additional HiFi simulations, also the non-linear features (e.g. minor species) start to be correctly captured with the information contained in the lower-energy modes (high k). Hence the loss of information becomes significant with low cut-off variances because such modes are removed. In fact, with lower cut-offs, the model's performance plateaus despite the addition of HiFi samples. Given the minimal cost differences among various cut-off values, we consider the 99.9% cut-off variance as the most suitable option for this study.

4.2. Incremental sampling

We show the MF-ROM performance using two different incremental sampling algorithms, namely Variance Based (VB) and Error Based (EB), used for adding HiFi samples progressively. The results of the two methods are reported in Fig. 5, both for case A (Fig. 5(a)) and for case B (Fig. 5(b)). As a reference, the ($\mathcal{E}_{\text{NRMS}}$) averaged across all QoIs for 4 test cases is compared between the MF-ROM and the Single-Fidelity ROMs obtained both with only LoFi samples (LoFi SF-ROM) and only with the given number of HiFi samples (HiFi SF-ROM). As we observe for case A, the MF-ROM delivers improved accuracy with respect to the HiFi SF-ROM when the number of HiFi samples is low, namely equal to 6 and 7. This was the sought-after main result, testifying to the effectiveness of the MF-ROM approach in integrating information from the LoFi model when the number of HiFi samples is scarce. The accuracy of the MF-ROM and the HiFi SF-ROM become comparable when the number of HiFi samples is equal to or higher than 8, with performance highly similar to the best-case scenario (i.e., 15 HiFi samples). This means that, after enough information is gathered for the reconstruction of all the QoIs, no more can be learned from additional HiFi samples.

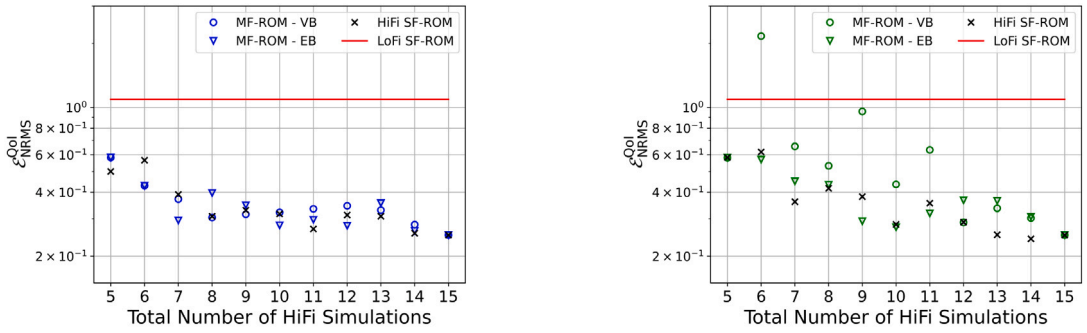

(a) Case A (b) Case B

Fig. 5. The average NRMSE of QoIs ($\mathcal{E}^{QoIs}_{\text{NRMS}}$) (*a*) for case A and (*b*) case B for variance and error-based sampling strategies.

Regarding the differences between cases A and B, it is clear that case A is a preferable alternative as both sampling approaches converge faster to an MF-ROM with comparable performance to the best-case scenario, while more scatter in the results is observable for case B (Fig. 5(b)). This is because the ROM necessitates more samples to include the variability in the data that must be captured, while it is already present in a more significant manner within the starting seed of case A. It is also observed that the first two simulations by the EB method were located at the corners of the design space, confirming the key role of corner simulations for developing a reliable MF-ROM, and the sampling algorithm aims to surround the design space as fully as possible. The effect of the sampling algorithm is evident for case B, as case A already contains an acceptable amount of variability in the POD modes. Even though both sampling strategies perform similarly for case A, the EV method is impractical to use for real-life design optimization problems since it requires the knowledge of all HiFi simulations that are not expected to be generally available. Therefore, case A and variance-based sampling were considered as baseline training settings for further analysis.

Summarizing, the important result is that, both for case A and case B, the MF-ROM can reconstruct the QoIs with better accuracy than the SF-ROM build with the same number of HiFi samples, for a relatively low amount of CFD simulations included, namely 6 for case A and 11 for case B. This implies that additional HiFi samples can actually be avoided by training an MF-ROM with similar accuracy to even the best-case scenario.

4.3. Interpolation vs extrapolation capabilities

We compare the abilities of the MF-ROMs trained with the best-selected options, namely case A and variance-based sampling, to both interpolate within the design space's boundaries (test cases T1 and T2) and to extrapolate towards unexplored conditions (T3 and T4) since extrapolation remains a frequent vulnerability of fully data-driven tools.

First, we report in Fig. 6 the average NRMSE between T1 and T2 cases obtained from the MF-ROM for different amounts of HiFi samples, evaluated for each QoI. It can be noted that the reconstruction of variables with substantially linear behavior over the operating range considered (such as temperature, which depends rather linearly on ϕ_{rich} and T_{air}, and stoichiometry-dependent species such as O_2, H_2O and N_2) is effective, while minor species remain more challenging to predict as they have a strongly non-linear behavior with both operational parameters considered. This also implies that, as the low-fidelity model is able to capture the trend of major species extremely well and that of pollutants and radicals less efficiently, the number of HiFi samples required for an acceptable prediction of the latter quantities is generally higher. Nevertheless, with 11 HiFi samples, all quantities are reconstructed with virtually identical performance to the best-case scenario.

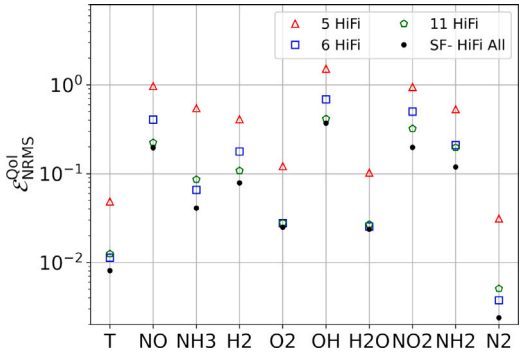

Fig. 6. The average NRMSE of QoIs ($\mathcal{E}^{QoIs}_{\text{NRMS}}$) for test data inside the design space (T1 and T2) for each QoI.

Since the prediction error on some variables, such as NO, might appear to be slightly high, we report the direct comparison between the original CFD flow field and the one predicted by the MF-ROM for test case T2, as visible in Fig. 7. As it is shown, remarkable agreement is already observed when 6 HiFi are used to train the MF-ROM, with very similar results to the 11 HiFi samples scenario, while a considerable mismatch is observed when only 5 HiFi samples are considered. For the 11 HiFi samples case, a relative error of 0.3% was observed in the prediction of the NO peak value, testifying to a high degree of accuracy in interpolating even hard-to-predict variables.

Afterward, we directly compare interpolation and extrapolation performance, and the results are shown in Fig. 8. It is observed that the average error obtained for the interpolation cases T1 and T2 (blue triangles) is always lower than the one obtained for extrapolation cases T3 and T4 (green squares). This can be attributed to the fact that it is more difficult to identify the strong nonlinear trend for characterizing cases that are quite far from the training range. Consequently, more HiFi samples are needed to include such variability in the POD modes. As a further testament to this, we show the contours of NO in Fig. 9 for the extrapolated case, T3, obtained with 11 HiFi samples, which is where MF-ROM is the closest to its best-case reference, showing a good agreement with the CFD counterpart.

The computational cost of the MF-ROM includes data generation, training, and prediction. While training and reconstructing costs occur on the order of seconds, – less than 1% of the time for a full CFD solution – the primary computational expense for MF-ROM lies in data generation. Besides achieving precise predictions, it is essential for the MF-ROM to utilize a training dataset with an optimal number of HiFi simulations to maximize computational efficiency. Fig. 8 demonstrates two notable improvements in the accuracy of MF-ROM for interpolation cases, where the $\mathcal{E}_{\text{NRMS}}$ stabilizes: one when the total number of HiFi simulations reaches 6, and another when it reaches 11. The MF-ROM

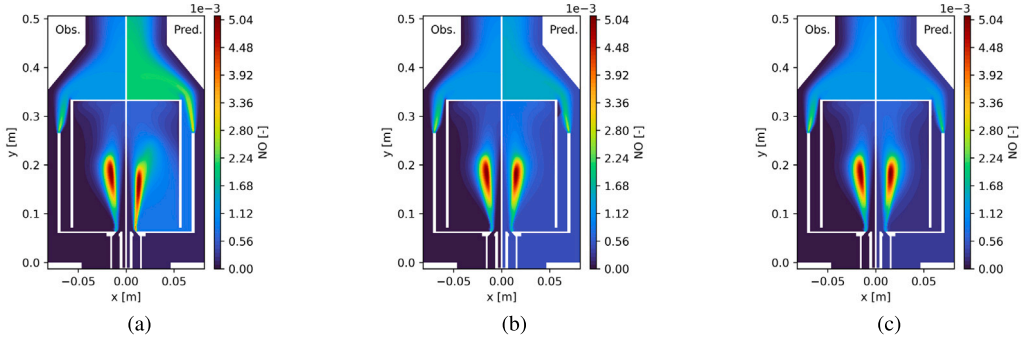

Fig. 7. NO field predictions of MF-ROM (*right*) built with (a) 5 HiFi, (b) 6 HiFi and (c) 11 HiFi with respect to the CFD simulation (*left*) for T_{air} = 840 K and ϕ_{rich} = 1.27 (T2).

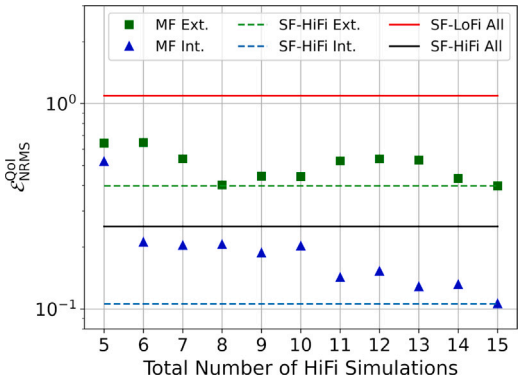

Fig. 8. The average NRMSE of QoIs ($\mathcal{E}^{QoIs}_{NRMS}$) for internal (T1 and T2) and external (T3 and T4) test data.

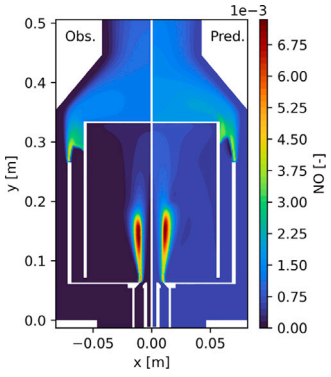

Fig. 9. NO field prediction of MF-ROM (*right*) with respect to the CFD simulation (*left*) for T_{air} = 1050 K and ϕ_{rich} = 1.65 (T3) with 11 HiFi simulations.

with 11 HiFi simulations guarantees 85% accuracy while offering a 25% reduction in overall computational cost during the data generation phase. Furthermore, the performance of the MF-ROM for the scenario involving 6 HiFi simulations is improved by up to 50% compared to the baseline case with 5 HiFi MF-ROM. This makes it possible to further reduce the training cost by up to 60%.

Interestingly, for extrapolation cases, the MF-ROM achieves performance close to the HiFi SF-ROM made with all available HiFi for a relatively low number of HiFi in the training. This is mainly because the variance related to T3 and T4 cases can only be covered within the POD modes by adding HiFi samples outside the range, and the level of information needed to predict these cases is therefore contained in a very similar manner in the POD modes of the low-fidelity model, while the addition of HiFi in interpolation continuously improves the results on T1 and T2, as more HiFi simulations enrich the available information even in lower energy modes. Therefore, an important consequence of this aspect is that the low-fidelity model can be employed to expand the boundaries of the design space in a computationally efficient manner, and subsequently, HiFi samples can be strategically performed in smaller numbers using the MF-ROM strategy, thus allowing for saving computational budget with respect to a solely HiFi ROM approach.

5. Conclusion

We introduced a novel MF-ROM design procedure that blends data from simulations at different levels of fidelity for enhanced modeling and prediction in combustion processes. This framework is assembled by integrating Proper Orthogonal Decomposition (POD) for dimensionality reduction and Procrustes Manifold Alignment for the projection of the manifolds at different levels of fidelity. The CoKriging regression model is used to predict the POD coefficients as a function of the input parameters, from which the full state of the system can be recovered via the projection matrix. A multi-stage, ammonia-fueled combustor prototype was analyzed across various sampling strategies and initialization techniques within and beyond the parameter space as a benchmark test case. The high-fidelity dataset was constructed using 15 RANS simulations within a 2-parameter design space, while 4 extra simulations were used as test cases. The low-fidelity dataset was assembled by performing cost-effective CRN simulations from a previously designed equivalent CRN system model at the same 15 operating conditions. By training several MF-ROMs with different strategies, we could draw the conclusions summarized within the following snapshots:

- The method is generally effective in predicting quantities of interest, even slightly outside of the design range. Furthermore, the multi-fidelity approach delivers a reduced model with higher accuracy than the single-fidelity model obtained with the same number of high-fidelity samples when the latter is scarce. When the high-fidelity dataset is enriched, the role of the low-fidelity information becomes secondary.
- The method is sensitive to the initial set of high-fidelity samples included in the training of the multi-fidelity model if the dataset size is small. In particular, by including the extremes of the design space at the beginning of the training procedure, the multi-fidelity approach converges with fewer high-fidelity samples to performance similar to the best high-fidelity model.
- The variance cut-off, on the other hand, has a less pronounced effect, and its importance is greater in the interpolation phase.

A high variance cut-off, in general, allows for a constant improvement in the model when adding more high-fidelity samples because lower-energy modes are retained.
- The variance-based sampling algorithm is effective in choosing the subsequent addition of high-fidelity samples to the training set than the error-based one.
- In interpolation, progressively adding high-fidelity samples consistently improves performance, whereas, in extrapolation, additional out-of-range samples are needed to cover new conditions. This method initially covers the unexplored range with the low-cost model and then strategically adds high-fidelity samples, optimizing the computational budget.

In summary, the proposed method efficiently blends information from data with structurally different levels of fidelity. The results demonstrate that the approach helps alleviate the computational budget needed to train reliable reduced-order models mimicking the behaviour of complex combustion systems. As for future work, we plan to investigate the definition of error estimates to optimally place the HiFi samples in the parameter space and improve the accuracy of the model.

Novelty and Significance Statement

The methodology presented in this work is innovative and contributes to the emerging literature on data-driven methods for the development of accurate and reliable reduced-order models of complex combustion systems. The approach demonstrates, for the first time, the possibility of effectively blending simulation frameworks that are fundamentally different, enabling accurate prediction of spatial variable distributions in unexplored design scenarios, particularly tailored for combustion systems. The physics-based multi-fidelity approach presented here fuses the zero/one-dimensional Chemical Reactors Network and two-dimensional RANS simulations data by means of Proper Orthogonal Decomposition and manifold alignment and exploits the spatial relations between these correlated variables with CoKriging. The significance arises from the imperative need, especially in the context of energy transition, to employ numerical tools, such as the multi-fidelity reduced order framework, that are both accurate and inexpensive for activities such as design exploration, optimization, and real-time access to predictions in combustion systems.

CRediT authorship contribution statement

Aysu Özden: Conceptualization, Methodology, Software, Writing – original draft, Visualization. **Matteo Savarese:** Conceptualization, Modelling, Data management, Visualization, Writing – original draft. **Lorenzo Giuntini:** Data management, Writing – review & editing. **Alberto Procacci:** Methodology, Software, Writing – review & editing. **Riccardo Malpica Galassi:** Conceptualization, Methodology, Writing – review & editing, Supervision. **Axel Coussement:** Supervision. **Francesco Contino:** Supervision. **Alessandro Parente:** Conceptualization, Methodology, Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

A. Özden acknowledges the financial support of the Fonds de la Recherche Scientifique de Belgique (FRS-FNRS) project METRIC, under grant No. T.0175.21. M.Savarese and A. Procacci acknowledge the financial support of the Fonds National de la Recherche Scientifique (FRS-FNRS) through an FRIA fellowship. R.Malpica-Galassi acknowledges the financial support from the Italian Ministry of University and Research in the framework of the project "National Center for "High-Performance Computing, Big Data, and Quantum Computing" funded by the European Union – NextGenerationEU. This work was supported by the Walloon Region through the Fonds de la Recherche Scientifique-FNRS, Belgium for the FRFS-WEL-T under Grant No. WEL-T-CR-2023 A-07.

References

[1] T. Lu, C. Law, Towards accomodating realistic fuel chemistry in large-scale computations, Prog. Energy Combust. Sci. 35 (2009) 192–215.
[2] P. Benner, S. Gugercin, K. Willcox, A survey of projection-based model reduction methods for parametric dynamical systems, SIAM Rev. 57 (4) (2015) 483–531.
[3] E. Glaessgen, D. Stargel, The digital twin paradigm for future NASA and US air force vehicles, in: 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA, 2012, p. 1818.
[4] A. Parente, N. Swaminathan, Data-driven models and digital twins for sustainable combustion technologies, iScience 27 (4) (2024) 109349, http://dx.doi.org/10.1016/j.isci.2024.109349.
[5] Y. Kuya, K. Takeda, X. Zhang, A.I. Forrester, Multifidelity surrogate modeling of experimental and computational aerodynamic data sets, AIAA J. 49 (2) (2011) 289–298.
[6] Z.-H. Han, Zimmermann, S. Görtz, Alternative cokriging model for variable-fidelity surrogate modeling, AIAA J. 50 (5) (2012) 1205–1210.
[7] C. Perron, D. Rajaram, D. Mavris, Development of a multi-fidelity reduced-order model based on manifold alignment, in: AIAA Aviation 2020 Forum, 2020, p. 3124.
[8] B. Peherstorfer, K. Willcox, M. Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization, SIAM Rev. 60 (3) (2018) 550–591.
[9] H. Khodayari, F. Ommi, Z. Saboohi, A review on the applications of the chemical reactor network approach on the prediction of pollutant emissions, Aircraft Engineering and Aerospace Technology 92 (4) (2020) 551–570.
[10] M. Savarese, A. Cuoci, W. De Paepe, A. Parente, Machine learning clustering algorithms for the automatic generation of chemical reactor networks from CFD simulations, Fuel 343 (2023) 127945.
[11] V. Fichet, M. Kanniche, P. Plion, O. Gicquel, A reactor network model for predicting NOx emissions in gas turbines, Fuel 89 (9) (2010) 2202–2210.
[12] J. Park, T. Huu Nguyen, D. Joung, K. Huh, M. Lee, Prediction of NOx and CO emissions from an industrial lean-premixed gas turbine combustor using a chemical reactor network model, Energy Fuels 27 (2013) 1643–1651.
[13] T.H. Nguyen, J. Park, C. Sin, S. Jung, S. Kim, Numerical investigation of the pressure effect on the NOx formation in a lean-premixed gas turbine combustor, Energy Fuels 35 (8) (2021) 6776–6784.
[14] Y. Sajjad, G. Bourque, R. Monaghan, Bayesian inference and uncertainty quantification for hydrogen-enriched and lean-premixed combustion systems, Int. J. Hydrog. Energy 46 (46) (2021) 23927–23942.
[15] M. Savarese, L. Giuntini, R. Malpica Galassi, S. Iavarone, C. Galletti, W. De Paepe, A. Parente, Model-to-model Bayesian calibration of a chemical reactor network for pollutant emission predictions of an ammonia-fuelled multistage combustor, International Journal of Hydrogen Energy 49 (2024) 586–601.
[16] M. Gastpar, P.L. Dragotti, M. Vetterli, The distributed Karhunen–Loeve transform, IEEE Trans. Inform. Theory 52 (12) (2006) 5177–5196.
[17] I.T. Jolliffe, Principal Component Analysis for Special Types of Data, Springer, 2002.
[18] G. Aversano, A. Bellemans, Z. Li, A. Coussement, O. Gicquel, A. Parente, Application of reduced-order models based on PCA & kriging for the development of digital twins of reacting flow applications, Comput. Chem. Eng. 121 (2019) 422–441.
[19] R. Pinnau, Model reduction via proper orthogonal decomposition, in: Model Order Reduction: Theory, Research Aspects and Applications, Springer, 2008, pp. 95–109.
[20] C. Wang, S. Mahadevan, A general framework for manifold alignment, in: 2009 AAAI Fall Symposium Series, 2009.

[21] B. Luo, E. Hancock, Feature matching with procrustes alignment and graph editing, in: Image Processing and Its Applications, 1999. Seventh International Conference on (Conf. Publ. No. 465), vol. 1, 1999, pp. 72–76.
[22] C. Wang, S. Mahadevan, Manifold alignment using procrustes analysis, in: Proceedings of the 25th International Conference on Machine Learning, 2008, pp. 1120–1127.
[23] C.K. Williams, C.E. Rasmussen, Gaussian Processes for Machine Learning, vol. 2, MIT press Cambridge, MA, 2006.
[24] M.C. Kennedy, A. O'Hagan, Predicting the output from a complex computer code when fast approximations are available, Biometrika 87 (1) (2000) 1–13.
[25] P. Perdikaris, D. Venturi, J.O. Royset, G.E. Karniadakis, Multi-fidelity modelling via recursive co-kriging and Gaussian–Markov random fields, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 471 (2179) (2015) 20150018.
[26] L. Le Gratiet, J. Garnier, Recursive co-kriging model for design of computer experiments with multiple levels of fidelity, Int. J. Uncertain. Quantif. 4 (5) (2014).
[27] A.I. Forrester, A. Sóbester, A.J. Keane, Multi-fidelity optimization via surrogate modelling, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 463 (2088) (2007) 3251–3269.
[28] J. Otomo, M. Koshi, T. Mitsumori, H. Iwasaki, K. Yamada, Chemical kinetic modeling of ammonia oxidation with improved reaction mechanism for ammonia/air and ammonia/hydrogen/air combustion, Int. J. Hydrog. Energy 43 (5) (2018) 3004–3014.
[29] S. Trespi, H. Nicolai, P. Debiagi, J. Janicka, A. Dreizler, C. Hasse, T. Faravelli, Development and application of an efficient chemical reactor network model for oxy-fuel combustion, Energy Fuels 35 (9) (2021) 7121–7132.