



ECOLE
POLYTECHNIQUE
DE BRUXELLES



A multi-fidelity, reduced-order modeling framework for the development of digital twins of combustion systems

Thesis presented by Aysu Özden

in fulfillment of the requirements of the PhD Degree in
Engineering Sciences and Technology (ULB & UCL - "Doctorat en
Sciences de l'ingénieur et technologie")

Academic year 2025-2026

Supervisor : Prof. Alessandro Parente

[Aero-Thermo-Mechanics Laboratory, Université Libre de Bruxelles]

Co-supervisor : Prof. Francesco Contino

[Thermodynamics and Fluid Mechanics Laboratory - Université
Catholique de Louvain]

Thesis jury :

Prof. Axel Coussement (Université Libre de Bruxelles, Chair)
Prof. Aude Simar (Université Catholique de Bruxelles, Secretary)
Prof. Laura Mainini (Imperial College London)
Prof. Temistocle Grenga (University of Southampton)
Prof. Riccardo Malpica Galassi (Sapienza Università di Roma)

To my parents,
who gave me the gift of dreams
and the ability to realize them..

*Hayal edebilme ve o hayalleri
gerçekleştirebilme gücünü bana veren
canım annem ve babama...*

To my younger self,
who decided to go through
this challanging life experience..

Acknowledgements

I wasn't sure how to begin this part, but a quote I came across captured my feelings quite well. I would like to share these words from Dr. James Beattie's PhD thesis:

"Embarking on a PhD is a treacherous task for even the most bold and brave of this world, and yet here I sit, writing my way through the final hoops of this life-changing experience."

This PhD has truly been an extraordinary and life-changing experience. Although there were days when I felt there was no light at the end of the tunnel, I am deeply grateful to have reached its end, surrounded by people who supported me throughout this journey.

I would like to begin by sincerely thanking my supervisors, Prof. Alessandro Parente and Prof. Francesco Contino, for giving me the opportunity to embark on this incredible journey. Prof. Parente, I truly appreciate that you always put me first, even when we had our differences over the past four years. You cared about me as a person, not just about the work, and I can't thank you enough for that. I've never had a professor as supportive as you. Even when I doubted myself, you believed in me and helped me find confidence again. Your knowledge, guidance, and kindness have left a lasting impact on me, and I'll always carry that with me. I would also like to thank Prof. Contino, for being there for me when I needed, and believing in me.

I would also like to sincerely thank the members of the jury: Prof. Axel Coussement, Prof. Aude Simar, Prof. Laura Mainini, Prof. Temistocle Grenga, and Prof. Riccardo Malpica Galassi. I am especially grateful to Prof. Axel

Coussement for his patience with all my clustering challenges and for his valuable guidance throughout this process. I also appreciate the jury's patience during the difficult task of scheduling my private defense and their time and effort in providing constructive feedback to improve the manuscript. Your comments and suggestions have been truly invaluable.

I would like to express my deep gratitude and respect to Riccardo in particular. You have been a true mentor, and I owe much of the researcher I am today to you. You supported me not only academically but also personally when I needed it most. I struggled a great deal in the beginning, but your patience, humility, and kindness helped me navigate the hardest part of this journey. You always had answers to my questions, and knowing you were there has been a constant source of comfort and confidence for me.

I am also incredibly grateful to have found not just colleagues, but true friends at ATM. Your unconditional support and encouragement made me stronger. I treasure all the ups and downs we went through together — I genuinely wouldn't have wanted to walk this path with anyone else.

I want to especially thank Laura. You are not just a friend — you are a sister to me. I am endlessly grateful to have found you. You are the person I can laugh without holding back, and the one in front of whom I can cry without shame. You showed up for me every single day, through every challenge, and you never wavered. You have been my rock, my safe place, and one of "my people" — for life. Of course, when I think of you, it is impossible not to think of Lorenzo. Lorenzo has always amazed me with his patience, intelligence, humility, and kindness. Every time I see the two of you together, I understand why you are soulmates — you truly are perfect for each other.

I am also deeply grateful for my Spanish girls, Patricia and Eva. Patricia, you are so much like me that sometimes I'm surprised by how easily we understand each other. You've been my travel partner, my confidant, and my go-to person for all my important decisions. Eva, your positivity and your laughter always brightened my days. I will never forget your selfness and kindness when you set everything aside to help me with the PROCI paper at a time I felt completely overwhelmed. That gesture meant so much to me and I will always carry that memory with me.

Laura, Patricia, Eva — thank you for allowing me to be my true self, always, without judgment. You have been the best friends I could have ever wished for.

A big thank you to the best officemates — Chiara, Lorenzo P., and Emanuele — for your friendship and for making the effort to speak English (most of the time!). Thanks to you, I also improved my *italiano* along the way.

I would like to express my sincere gratitude to Alberto. He was one of the people I could always rely on throughout this PhD. He patiently answered every one of my questions and consistently supported me with his brilliant ideas and valuable feedback.

I also wish to thank the other members of our group. Himanshu, a former colleague, taught me that there will always be ups and downs during a PhD, but that in the end, each challenge helps you grow. I am especially grateful to Marianna, Mahdi, Luisa, Alessandro P., Natalia, and all my past and present colleagues in the department — you all contributed to creating such an enjoyable and supportive working environment. And, of course, I sincerely thank you for welcoming Thor — though I must admit, I sometimes felt I was just his plus-one!

I can not forget two annoying brothers I gained during my time in the VKI, Ali Can and Ludovico. You are the motivators behind the reason that I started this PhD. Even though you like to make fun of me constantly, I am so luckily to find two *genco's* who will always have my back.

I would also like to thank my lifelong friends, without whom I wouldn't be the person I am today. Zeyküd - a.k.a. balıçörek - you are one of the strongest people I know, and your strength makes me stronger. Maggie, my sister who may be physically far but is always close to my heart — I constantly feel your unwavering love and support, no matter the distance.

I would like to thank to Thor, my loyal companion. Thank you for walking beside me through every step of this journey. You were there for the quiet mornings, the long nights, the stressful deadlines, and the joyful victories. Your unconditional love, your calming presence, and even your playful distractions were exactly what I needed, exactly when I needed them. You reminded me to pause, to breathe, to go outside, and to find comfort in the simple things. You have been my friend, my family, and my silent supporter through it all. You are also the reason I met my *park friends* - Noelle, Barbara, Sophia, Gabriella, Ilenia,

Alessandro & Valentina, Riccardo & Rosanna, Sarah & John, Alain & Sophia – who have listened to all my problems and have been incredibly supportive of me for the past 4 years. So, thank you, Thor, for being by my side. I couldn't have done this without you.

Finally, Mom and Dad, there are no words that can truly capture the depth of my love and gratitude for you. Every step I've taken, every challenge I've overcome, and every milestone I've reached — they are not just mine, they are ours. You always believed in me on the days I struggled to believe in myself. You always reminded me that there is nothing I cannot achieve, and you did everything in your power to support my dreams. You made countless sacrifices to give me the best life possible. *Annem*, you listened to me for hours, day after day — sometimes to laugh, sometimes to complain, sometimes to cry. You are the strongest woman I have ever known, and if I can carry even a small piece of your strength within me, I will consider myself truly fortunate. *Babam*, you showed me the greatest love. Even when you were enduring the greatest pain of your life, you put me first. I will never forget that. Just talking to you about my problems has always been enough to help me find my way. You are the strongest man I know and the mountain I can lean on. Thank you both for your wisdom, your endless support, and for always being my safe place. Your love has been my greatest strength and my guiding light — *benim kanaat liderlerim*.

All my achievements are ours, and this PhD is for you.

Abstract

Climate change is driving dramatic shifts worldwide, including wildfires, droughts, and unusually high temperatures. Without urgent efforts to reduce greenhouse gas (GHG) and pollutant emissions, exceeding the established limits for global temperature rise will become inevitable. A crucial step is to develop innovative solutions that address the link between energy demand and pollutant emissions. Since most global emissions stem from traditional combustion-based processes, finding ways to reduce emissions from these sources will have the greatest impact.

Researchers are actively working on novel combustion technologies that can achieve high energy efficiency, offer fuel flexibility, reduce greenhouse gas and pollutant emissions, and ensure robust and reliable operation. However, the design process of such systems is time-consuming and costly, as changes in design parameters directly affect system performance. Evaluating the operational behavior of these technologies heavily depends on extensive experimental campaigns and high-fidelity numerical simulations. Due to the multi-scale and multi-physics nature of combustion systems, achieving both physical insight and robust designs require tools that are accurate yet computationally efficient to explore and predict a wide range of operating conditions and device geometries.

Exclusively relying on experimental diagnostics is not practical, as accessing the combustion zone is both technically challenging and prohibitively expensive. From a computational perspective, although computational fluid dynamics tools have advanced significantly in recent years, accurately predicting complex reacting flows—especially those involving detailed chemical kinetics and advanced turbulence-chemistry interactions—still demands immense computational resources, which remains a critical barrier.

The growing interest and progress in machine learning algorithms offer a promising pathway to address this challenge. Data-driven reduced order models (ROMs) can replace computationally expensive full-order simulations, accelerating design exploration, optimization, and real-time predictions. However, these models typically require large training datasets generated from high-fidelity simulations, and the associated computational cost still represents a major limitation in their practical use.

Multi-fidelity reduced order models provide a solution by striking a balance between accurate but costly high-fidelity data and less accurate but inexpensive low-fidelity data, enabling reasonably accurate system predictions at significantly reduced computational cost. Given the typically large number of quantities of interest and the high dimensionality of numerical domains, training datasets for such models suffer from the curse of dimensionality. To address this, dimensionality reduction techniques are commonly applied to simplify the regression problem. In multi-fidelity frameworks, a shared low-dimensional representation is established across both fidelity levels using manifold alignment techniques. These shared manifolds are then employed to train regression models capable of accurately predicting the system's full state under new operating conditions.

The overarching goal of my work has been to develop a multi-fidelity reduced order modeling framework for combustion systems. Specifically, my research combined proper orthogonal decomposition for dimensionality reduction, Procrustes manifold alignment to establish shared low-dimensional manifolds, and CoKriging—an extension of Gaussian Process Regression—to predict the system's state.

In particular, to address one of the key questions in multi-fidelity modeling—determining the minimum number of high-fidelity data points required for a desired level of model accuracy—I developed incremental sampling algorithms to (i) reduce the model prediction error, and (ii) minimize the uncertainty of the framework. The accuracy of the multi-fidelity framework and the effectiveness of the incremental sampling strategies were first assessed on a furnace operating with a methane-hydrogen fuel blend under Moderate and Intense Low-oxygen Dilution combustion conditions. The results showed that sampling strategies focused on reducing model uncertainty were most effective, and that

the multi-fidelity model achieved prediction accuracy comparable to its single fidelity counterpart while reducing the number of training degrees of freedom by approximately half.

The same framework and sampling strategies were also applied to ammonia combustion in a stagnation-point reverse-flow combustor to test the framework's adaptability to different fuels and configurations. In addition, I investigated the effect of CoKriging predictions in extrapolation scenarios — specifically, test conditions outside the original design space. The results demonstrated that the multi-fidelity framework is flexible and effective across different combustion systems and fuels. While prediction accuracy and computational savings were lower in extrapolation cases compared to interpolation, the framework still successfully captured the qualitative behavior of key quantities of interest outside the training range.

The training dataset for the ammonia case was initially limited, which motivated expanding the design space through additional simulations. This expansion revealed significant changes in combustion characteristics within the combustor as the equivalence ratio and air inlet temperatures decreased, largely due to ammonia's narrow flammability limits. To further improve prediction accuracy, a hierarchical clustering technique was introduced to optimize the distribution of training samples. The results showed that this clustering approach not only provided a more structured strategy for model development but also improved prediction accuracy while reducing computational cost by tailoring the training process to localized behaviors within the design space.

Finally, a multi-level multi-fidelity reduced order model was developed by integrating experimental measurements as high-fidelity data, while 3D and 2D RANS simulations served as mid- and low-fidelity datasets, respectively. The results demonstrated that incorporating additional high-fidelity experimental data progressively stabilized the model's response and significantly improved its prediction accuracy compared to the multi-fidelity model trained exclusively on numerical simulations. It was also observed that the framework was highly sensitive to the initial selection of the linked datasets and to the constraints applied when selecting subsequent high-fidelity data points within the design space.

Overall, this work shows that multi-fidelity reduced order models provide an effective solution for balancing accuracy and computational cost. The adaptability of this approach to different fuels and combustion systems makes it particularly valuable for the combustion community, as it accelerates design and optimization processes that are otherwise computationally intensive and costly, especially in combustion applications. Future work will focus on improving model accuracy by expanding the training dataset, optimizing hyperparameters, exploring alternative regression methods, and integrating physical knowledge to develop physics-informed models.

Table of contents

List of figures	xvii
List of tables	xxiii
Nomenclature	xxv
1 Introduction	1
1.1 Global energy crisis and the path toward decarbonizing combustion	2
1.2 Advanced combustion technologies for a sustainable future	4
1.2.1 Formation mechanisms for NO_x	4
1.2.2 MILD combustion	6
1.2.3 Staged combustion	8
1.3 Combustion modeling and challenges	9
1.4 Machine learning and combustion science	11
1.5 Multi-fidelity surrogate modeling	15
1.6 Motivation and outline	18
2 Mathematical background	21
2.1 Proper orthogonal decomposition	23
2.2 Manifold alignment	26
2.3 CoKriging regression	29
2.3.1 Simple CoKriging	30
2.3.2 Multi-level CoKriging	33
2.4 Sampling strategies	34

2.4.1	Adaptive sampling based on prediction error	36
2.4.2	Adaptive sampling based on model uncertainty	37
2.5	Sparse sensing	37
3	Experimental setup	39
3.1	ULB semi-industrial flameless furnace	40
3.1.1	Temperature measurements	42
3.1.2	Energy and mass balance	44
4	Data generation	51
4.1	ULB semi-industrial flameless furnace	52
4.1.1	CFD data generation	52
4.1.2	Experimental data generation	54
4.1.2.1	Sparse sensing reconstruction	58
4.2	Stagnation point reverse-flow combustor	59
4.2.1	CFD data generation	62
4.2.2	CRN data generation	63
5	Development of the multi-fidelity reduced order model	65
5.1	Methane-hydrogen mixtures in the ULB furnace	66
5.1.1	Dataset and workflow settings	66
5.1.2	Reference high-fidelity digital model	67
5.1.3	Multi-fidelity digital model	68
5.1.3.1	Error-based incremental sampling method	69
5.1.3.2	Uncertainty-based incremental sampling method	73
5.1.3.3	Performance comparison of incremental sam- pling methods	76
5.2	Summary	79
6	Adaptability of the MF-ROM to different type of combustion sys- tems	81
6.1	Ammonia combustion in a stagnation point reverse-flow combustor	82
6.1.1	Test case	82

6.1.1.1	MF-ROM setup	82
6.2	Results	83
6.2.1	Model sensitivity to cut-off variance	83
6.2.2	Incremental sampling	84
6.2.3	Interpolation and extrapolation capabilities	86
6.3	Summary	90
7	Clustering based multi-fidelity reduced order model for ammonia combustion	93
7.1	Regime identification: Unsupervised clustering	94
7.1.1	Hierarchical sampling	94
7.1.2	Clustering evaluation metrics	95
7.2	Case Study: Two-staged ammonia combustor	99
7.2.1	MF-ROM Configuration	99
7.3	Results	100
7.3.1	Global clustering for MF-ROM improvement	100
7.3.2	Localized clustering for performance refinement	106
7.4	Summary	109
8	Multi-level multi-fidelity framework for methane-hydrogen combustion: integrating CFD and experimental data	111
8.1	Case Study: ULB semi-industrial furnace	112
8.1.1	Multi-level multi-fidelity (MLMF-ROM) configuration	113
8.2	Results	114
8.2.1	Performance of MLMF-ROM	114
8.2.2	Impact of dataset limitations on model performance	117
8.3	Summary	119
9	Conclusions and future work	121
9.1	Conclusions	122
9.2	Future work	124
	List of publications	127

References	129
Appendix A Prediction and error metrics	141

List of figures

1.1	Global surface air temperature ($^{\circ}\text{C}$) increase above the average for the 1850–1900 designated pre-industrial reference period, based on several global temperature datasets, shown as annual averages since 1967 [18].	3
1.2	Simplified scheme of formation pathways of the NO species: thermal (red arrows), prompt (blue arrows), NNH-intermediate (yellow arrows), N_2O -intermediate (green arrows), adapted from [49].	5
1.3	Different combustion modes for a $\text{CH}_4/\text{O}_2/\text{N}_2$ mixture in terms of T_{in} and ΔT , adapted from [15].	7
1.4	The number of review articles published annually between 2014 and 2024 on ScienceDirect using the search query “ <i>combustion</i> ” AND (“ <i>machine learning</i> ” OR “ <i>data mining</i> ” OR “ <i>feature extraction</i> ”).	14
2.1	Schematic drawing of POD.	25
2.2	Steps of the Procrustes manifold alignment technique: the original datasets (<i>top left</i>), translated manifolds (<i>top right</i>), scaled manifolds (<i>bottom left</i>), and aligned manifolds (<i>bottom right</i>).	29
2.3	Example of CoKriging regression using the toy datasets provided in [31].	33
3.1	Schematic of the ULB furnace (<i>left</i>), its cross section (<i>right top</i>), and the nozzle (<i>bottom right</i>). Only some probe ports are shown for the sake of clarity.	40

3.2	Schematic piping and instrumentation diagram of the ULB experimental test setup.	42
3.3	Schematic view of the suction pyrometer.	43
3.4	Simplified sketch of the ULB furnace for the global power balance	46
3.5	Separate sketches of the combustion chamber (<i>left</i>), and the heat exchanger (<i>right</i>) considered for the global power balance.	47
3.6	Idealized sketch of the window to estimate the radiative power transmitted through the window.	48
4.1	Design space composed of fuel composition, equivalence ratio, and injector diameter. Points are colored by the injector diameter value.	53
4.2	The sketch of the ULB furnace with thermocouple ports and approximate measurement locations.	54
4.3	CFD model and experimental temperature measurements for TCA at axial positions, (a) $z=0.1$ m, (b) $z=0.2$ m, (c) $z=0.3$ m, (d) $z=0.4$ m, (e) $z=0.5$ m and (f) $z=0.6$ m.	57
4.4	CFD model and experimental temperature measurements for TCB at axial positions, (a) $z=0.1$ m, (b) $z=0.2$ m, (c) $z=0.3$ m, (d) $z=0.4$ m, (e) $z=0.5$ m and (f) $z=0.6$ m.	58
4.5	NRMSE of temperature ($\mathcal{E}_{\text{NRMS}}^T$) for each experimental data for 16mm injector diameter for different total variance values.	59
4.6	NRMSE of temperature ($\mathcal{E}_{\text{NRMS}}^T$) for each experimental data for 25mm injector diameter for different total variance values.	60
4.7	Temperature profiles of CFD results and SpS reconstruction for (a) Dia = 16mm, $H_2 = 0.05$, $\phi=0.83$ (TC12), and (b) Dia = 25mm, $H_2 = 0.80$, $\phi=0.90$ (TC6).	60
4.8	Schematic view of the stagnation point reverse-flow combustor investigated in this work [90]	61

4.9	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)	62
4.10	Design space considered in this work. The full black dots represent the training cases, while the full triangles represent the test cases with eight interpolation cases (T1–T8).	63
4.11	Scheme of the Chemical Reactor Network of the combustor [90]	64
4.12	Schematic of the LoFi data generation, CFD field (<i>top</i>) and CRN extrapolation (<i>bottom</i>).	64
5.1	30-fold statistics (mean, standard deviation and kernel density estimation) of the QoIs NRMS errors $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ exhibited by the reference high-fidelity ROM built by using the full set of 41 HiFi simulations.	68
5.2	Slice of the 3D flow field compared between the true field from the CFD simulation (<i>left half</i>) and the digital twin prediction (<i>right half</i>) in a selected test point for temperature (a) and OH mass fraction (b).	68
5.3	30-fold statistics (mean, standard deviation and kernel density estimation) of the QoI-averaged NRMS error E_{NRMS} for Base MF-ROM and SF-ROM compared to the SF-ROM with 41 HiFi simulations (Full SF-ROM).	70
5.4	30-fold statistics (mean, standard deviation, and kernel density estimation) of the QoI-averaged NRMS error E_{NRMS} with increasing HiFi samples employing the error-based strategy.	71
5.5	30-fold-statistics of the QoIs NRMS errors $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ exhibited by MF-ROMs built by using 9, 18, 27 and 41 HiFi simulations employing the error-based sampling strategy.	72

5.6	Comparison between the MF-ROMs for the error-based method built with 9 (left), 18 (center), 27 (right) HiFi simulations and 41 LoFi, temperature (top row), and OH mass fraction (bottom row). Each plot compares a slice of the predicted 3D field (right half) against the same (mirrored) slice of the test simulation (left half).	73
5.7	30-fold-statistics of the relative errors ε_{obs} on wall temperature, peak OH value, peak OH position, and flame length exhibited by MF-ROMs built by using 9, 18, 27, and 41 HiFi simulations employing the error-based sampling strategy.	74
5.8	30-fold statistics (mean, standard deviation and kernel density estimation) of the NRMS error E_{NRMS} with increasing HiFi samples employing the uncertainty-based strategy.	75
5.9	30-fold-mean of the QoIs NRMS errors $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ exhibited by MF-ROMs built employing the uncertainty-based sampling strategy.	75
5.10	30-fold-statistics of the QoIs NRMS errors of the QoIs NRMS errors $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ exhibited by MF-ROMs built by using 9, 18, 27 and 41 HiFi simulations employing the uncertainty-based sampling strategy.	76
5.11	Comparison between the MF-ROMs for the uncertainty method built with 9 (left), 18 (center), 27 (right) HiFi simulations and 41 LoFi, temperature (top row), and OH mass fraction (bottom row). Each plot compares a slice of the predicted 3D field (right half) against the same (mirrored) slice of the test simulation (left half).	77
5.12	30-fold-statistics of the relative errors ε_{obs} on wall temperature, peak OH value, peak OH position, and flame length exhibited by MF-ROMs built by using 9, 18, 27, and 41 HiFi simulations employing the uncertainty-based sampling strategy.	78
5.13	Comparison between error-based and uncertainty-based incremental sampling with Guénot et al. [40]’s adaptive sampling strategies on the 30-fold-statistics of the NRMS error E_{NRMS} . . .	78
6.1	The average NRMSE of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) for interpolation test data (T1 and T2) for different total variance values for case A	84

6.2	The average NRMSE of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) (a) for case A and (b) case B for variance and error-based sampling strategies	86
6.3	The average NRMSE of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) for test data inside the design space (T1 and T2) for each QoI	87
6.4	NO field predictions of MF-ROM (<i>right</i>) built with (a) 5 HiFi, (b) 6 HiFi and (c) 11 HiFi with respect to the CFD simulation (<i>left</i>) for $T_{\text{air}} = 840 \text{ K}$ and $\phi_{\text{rich}} = 1.27$ (T2)	88
6.5	The average NRMSE of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) for internal (T1 and T2) and external (T3 and T4) test data	89
6.6	NO field prediction of MF-ROM (<i>right</i>) with respect to the CFD simulation (<i>left</i>) for $T_{\text{air}} = 1050 \text{ K}$ and $\phi_{\text{rich}} = 1.65$ (T3) with 11 HiFi simulations	90
7.1	Schematic of hierarchical sampling.	95
7.2	Davies-Boudin index to assess evaluate the clustering.	96
7.3	Silhouette index to assess evaluate the clustering.	97
7.4	Calinski-Harabsz index to assess evaluate the clustering.	98
7.5	Design space with 3 clusters	98
7.6	Temperature profiles of two training cases with different flow structures.	101
7.7	Cluster distribution after global clustering is applied.	102
7.8	Average NRME of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) of test cases $T3 - T8$ (<i>minimum error is shown with the red dotted line</i>).	103
7.9	Average NRME of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) of test cases $T1\&T2$ (<i>minimum error is shown with the red dotted line</i>).	104
7.10	Temperature profiles for the test case $\phi = 0.75, T = 625\text{K}$ for 8 HiFi simulations for (a) <i>MF-ROM with no clustering</i> and (b) <i>clustering-based MF-ROM</i>	105
7.11	NO field predictions of MF-ROM (<i>right</i>) built with 8 HiFi simulations for (a) no clustering, and (b) global clustering with respect to the CFD simulation (<i>left</i>) for T2.	106
7.12	Cluster distribution after local clustering is applied to C_1	107

7.13	Average NRMSE of temperature ($\mathcal{E}_{\text{NRMS}}^T$) of test cases for local clustering.	108
7.14	Temperature field predictions of MF-ROM (<i>right</i>) built with 8 HiFi simulations for (a) no clustering, (b) global clustering, and (c) local clustering with respect to the CFD simulation (<i>left</i>) for T5.109	
7.15	Average NRMSE of NO ($\mathcal{E}_{\text{NRMS}}^{\text{NO}}$) for global (<i>left</i>) and local (<i>right</i>) clustering. Data shown are for the test case with $\phi = 1.4, T = 730K$ for 30 HiFi simulations.	110
8.1	Design space composed of fuel composition, equivalence ratio, and injector diameter. Points are colored by the injector diameter value. The crosses represent the condition where three-fidelity data are available.	112
8.2	NRMSE of temperature ($\mathcal{E}_{\text{NRMS}}^T$) averaged over T1 and T2 to compare MLMF-ROM and MF-ROM prediction accuracy. . . .	115
8.3	Temperature profiles of CFD results and SpS reconstruction for T2 for (a) MLMF-ROM with 6 experimental data, (b) MLMF-ROM with 9 experimental data, and (c) MF-ROM with 18 3D RANS simulations.	116
8.4	NRMSE of temperature ($\mathcal{E}_{\text{NRMS}}^T$) averaged over T1 and T2 for different constraints applied to MF-ROM.	118

List of tables

- 4.1 Design parameters for the experimental data 56
- 7.1 DBI index values with increasing number of clusters for global clustering. 102
- 7.2 DBI index values with increasing number of clusters for local clustering. 107
- 8.1 Design parameters for the test cases T1 and T2 113
- 8.2 Design parameters for the initial seed of HiFi dataset for 25mm air-injector diameter 114

Nomenclature

Roman Symbols

C	Covariance matrix
d	Design condition matrix
d	Each design condition
E	Frobenius reconstruction error
k	Dimensionality of the low-dimensional manifold
l	Levels of fidelity
m	Number of designs/snapshots of low-fidelity data
n	Number of designs/snapshots of high-fidelity data
p	Dimensionality of high-fidelity data
Q	Rotation matrix
q	Dimensionality of low-fidelity data
r	Number of physical degrees of freedom
s	Scaling factor
T	Aligned low-dimensional manifold of low-fidelity data
U	Left singular vector

V	Right singular vector
$\mathbf{W}$	Low-dimensional projection of low-fidelity data matrix
$\mathbf{X}$	High-fidelity data matrix
$\mathbf{x}$	Column vector for each simulation
$\mathbf{Y}$	Low-fidelity data matrix
$\mathbf{Z}$	Low-dimensional projection of high-fidelity data matrix

Greek Symbols

α	Regression model
δ	Gaussian process for the discrepancy between the high- and low-fidelity manifolds
Σ	Diagonal matrix of singular values
$\mathcal{E}$	Normalized root mean square error
ι	Unit imaginary number $\sqrt{-1}$
Φ	High-fidelity spatial modes
Ψ	Low-fidelity spatial modes
ρ	Gaussian process scaling factor
σ	Gaussian process uncertainty

Subscripts

desVars	Design variables
in	inlet
L	Linked dataset
pred	predicted

rec	reconstructed
si	self-ignition
U	Unlinked dataset

Acronyms / Abbreviations

CFD	Computational Fluid Dynamics
CHI	Calinski-Harabasz Index
DBI	Davies-Bouldin Index
DNS	Direct Numerical Simulations
DoE	Design of Experiments
DT	Digital Twin
GHG	Greenhouse Gases
GPR	Gaussian Process Regression
HiFi	High-Fidelity
HPC	High Performance Computing
ID	Inner Diameter
LES	Large Eddy Simulations
LHS	Latin Hypercube Sampling
LHV	Lower Heating Value
LoFi	Low-Fidelity
MFC	Mass Flow Controller
MF-ROM	Multi-Fidelity Reduced Order Model
MidFi	Mid-Fidelity Data

MILD Moderate or Intense Low-oxygen Dilution

MLMF-ROM Multi-Level Multi-Fidelity Reduced Order Model

NLP Natural Language Processing

NRMS Normalized Root Mean Square

PCA Principal Component Analysis

PLM Product Lifecycle Management

POD Proper Orthogonal Decomposition

QoI Quantity of Interest

RANS-TE Reynolds-Averaged Navier Stokes Transport Equation

RANS Reynolds-Averaged Navier Stokes

RBF Radial Basis Function

RIC Relative Information Content

ROM Reduced Order Modeling

SI Silhouette Index

SpS Sparse Sensing

SVD Singular Value Decomposition

WSR Well-stirred Reactor

Chapter 1

Introduction

This chapter discusses the ongoing global energy crisis and the transition toward decarbonizing combustion through advanced combustion technologies. It also highlights the challenges in combustion modeling and the growing use of machine learning techniques, particularly multi-fidelity surrogate modeling, to address them. Finally, the chapter presents the motivation behind the thesis and outlines its structure.

1.1 Global energy crisis and the path toward decarbonizing combustion

One of the most significant challenges facing humanity in the 21st century is the global energy crisis. This crisis arises from the growing demand for energy coupled with the limited availability of fossil fuel resources, creating an urgent need to combat climate change. The impacts of climate change are increasingly evident, with more frequent and severe extreme events such as deadly heatwaves, intense rainfall, hurricanes, wildfires, and droughts. To prevent the most severe consequences of global warming and ensure a liveable future for all, the Paris Agreement established a goal to keep the global average temperature increase well below 2°C above pre-industrial levels while striving to limit the rise to 1.5°C [100].

However, 2024 marked a critical turning point as the first year when the global average temperature surpassed the 1.5°C threshold above pre-industrial levels (Figure 1.1). Alarming, 11 out of 12 months in 2024 recorded global average surface air temperatures above this threshold. The combined average temperature for 2023 and 2024 was measured at 1.54°C above pre-industrial levels [18].

The rise in surface temperatures is primarily driven by the increased release of greenhouse gases (GHG) — namely carbon dioxide, methane, and nitrous oxide — into the atmosphere. These emissions come from obvious sources like fossil-fueled power plants and factories, but also from less apparent contributors such as ships, planes, landfills, animal farming, and rice cultivation fields. According to data from Climate TRACE [97], power generation was responsible for 25% of global emissions in 2024, followed by manufacturing processes at 18%. Fossil fuel operations, including oil and gas production, refining, and coal mining, accounted for 16%, while transportation and agriculture each contributed 14% to global emissions. To effectively reduce GHG emissions from these sectors, decarbonizing the technologies they rely on is essential.

Decarbonization offers a promising solution not only for reducing greenhouse gas emissions and meeting the targets set by the Paris Agreement but also for




Fig. 1.1 Global surface air temperature ($^{\circ}\text{C}$) increase above the average for the 1850–1900 designated pre-industrial reference period, based on several global temperature datasets, shown as annual averages since 1967 [18].

promoting a sustainable future by enhancing air quality and improving public health. Decarbonizing energy-intensive sectors can significantly contribute to lowering global GHG emissions. Fossil fuel fed traditional combustion technologies, which have been the predominant energy source for decades, are major emitters of greenhouse gases and pollutants, leading to severe environmental damage [99]. Therefore, targeting combustion technologies presents a critical opportunity to reduce the carbon footprint of these sectors. Developing innovative combustion systems is essential for the energy transition, as it aims to enable the efficient utilization of carbon-free fuels and minimize greenhouse gas emissions, particularly in sectors where electrification remains challenging.

1.2 Advanced combustion technologies for a sustainable future

Advanced combustion technologies are a critical component of the global effort to decarbonize energy systems and reduce greenhouse gas emissions. These technologies aim to enhance the efficiency of combustion processes, minimize pollutant formation, and enable the use of low-carbon or carbon-free fuels. Key approaches include oxy-fuel combustion [92, 32], which uses pure oxygen instead of air to produce a concentrated CO_2 stream that is easier to capture; chemical looping combustion [2, 19], which employs solid oxides as oxygen carriers to achieve intrinsic CO_2 separation. When combined with carbon capture techniques, these technologies help to reduce atmospheric greenhouse gases and mitigate climate change by capturing CO_2 emissions from sources such as power plants and either storing it underground or utilizing it in products.

Pollutant-reducing and fuel-flexible combustion techniques aim to enhance efficiency while minimizing environmental impact. By optimizing combustion conditions and enabling the use of diverse fuels, these approaches lower emissions of harmful pollutants and support cleaner, more adaptable energy production. Flameless [107] or Moderate or Intense Low-oxygen Dilution (MILD) [15] combustion, reduces peak temperatures to significantly lower NO_x emissions. Furthermore, hydrogen combustion and ammonia-fuelled systems are gaining attention as promising carbon-free energy carriers. These advancements are not only crucial for mitigating emissions from traditional power generation but also for decarbonizing hard-to-electrify sectors such as aviation, shipping, and heavy industry.

1.2.1 Formation mechanisms for NO_x

In combustion systems, NO_x formation primarily occurs through the insertion of radicals such as O , CH_x , and H into the triple bond of molecular nitrogen in the air [73], producing reactive nitrogen intermediates like cyanides, N_2O , and NNH that are subsequently oxidized to NO_x . NO_x can also form via oxidation of organic nitrogen present in the fuel. In conventional combustion, the thermal route

and the fuel-N route are the main pathways when fuel contains nitrogen. In non-conventional regimes such as MILD combustion, additional pathways—prompt, N_2O -intermediate, and NNH-intermediate mechanisms—become significant. A graphical summary of these pathways leading to NO formation from molecular nitrogen is shown in Figure 1.2.




Fig. 1.2 Simplified scheme of formation pathways of the NO species: thermal (red arrows), prompt (blue arrows), NNH-intermediate (yellow arrows), N_2O -intermediate (green arrows), adapted from [49].

Under conventional combustion conditions with nitrogen-free fuels, the thermal route, first described by Zeldovich [114], is the dominant source of NO emissions. This mechanism becomes significant at temperatures above 1800 K in the presence of sufficient oxygen, with the NO formation rate doubling for every 35 K increase beyond 1800 K [7]. Because of this strong temperature dependence, most NO-reduction strategies focus on suppressing temperature peaks and/or limiting residence time in oxygen-rich, high-temperature zones.

The prompt route, first identified by Fenimore [28], arises when carbon-based fuels burn under fuel-rich, low-temperature conditions, as can occur in surface burners, staged combustion systems, and gas turbines. Although its contribution to total NO_x emissions from stationary combustors is currently small, its relative significance is expected to grow as advanced designs and mitigation strategies reduce other NO_x sources. This pathway involves CH_x radicals attacking molecular nitrogen, forming NCN , which subsequently reacts with O, H, and OH radicals to produce NO, as illustrated in Figure 1.2. Importantly, the prompt

route does not occur with carbon-free fuels such as hydrogen or ammonia, since no CH_x radicals are present.

The NNH-intermediate route [8] may gain importance when non-conventional fuels containing relevant amounts of hydrogen are oxidized. Konnov et al. [55] noted that the NNH route becomes the predominant source of NO at 1500 K, not only in rich conditions but also in lean and stoichiometric mixtures.

The N_2O -intermediate pathway [68] (green line in Figure 1.2) becomes significant in fuel-lean systems at low temperatures and high pressures, making it especially relevant for gas turbines and compression-ignition engines. As these devices are operated at increasingly low temperatures to suppress thermal NO_x formation, the relative contribution of the N_2O -intermediate mechanism is gaining importance.

1.2.2 MILD combustion

MILD combustion is an advanced combustion regime distinct from deflagration, detonation, or diffusion flames and can be described as a super diluted explosion or continuous auto-ignition [15]. Achieving and maintaining this unique combustion regime requires careful control of flame stabilization, which is essential for burner design. In particular, recirculation of hot exhaust gases rates play a vital role in ensuring flame stability and providing reliable operating conditions [107]. In traditional combustion systems, when temperatures are below the self-ignition temperature (T_{si}) and recirculation rates are high, the flame tends to become unstable and may blow out. However, when furnace temperatures and recirculation rates are sufficiently high, MILD combustion can be achieved, resulting in a stable combustion regime without a visible flame, which is why it is often referred to as *flameless* combustion [107].

The MILD combustion process is well explained by Cavaliere and de Joannon [15] using a well-stirred reactor (WSR) model with a stoichiometric $\text{CH}_4/\text{O}_2/\text{N}_2$ mixture at atmospheric pressure. As shown in Figure 1.3, three combustion modes illustrate various inlet temperature (T_{in}) and temperature increase (ΔT) scenarios. Feedback and High Temperature Combustion modes represent traditional combustion processes where heat release is sufficient to maintain the minimum

temperature required for sustained reactions. In contrast, MILD combustion requires preheating of reactants ($T_{in} > T_{si}$) to initiate the process. Additionally, the maximum temperature increase ($\Delta T = T - T_{in}$) in MILD combustion is lower than the mixture's self-ignition temperature due to reduced O_2 content.




Fig. 1.3 Different combustion modes for a $CH_4/O_2/N_2$ mixture in terms of T_{in} and ΔT , adapted from [15].

Preheating reactants beyond their self-ignition temperatures and mixing them with large amounts of recirculated flue gases create highly diluted local conditions with a uniform temperature distribution. The absence of high-temperature zones not only enhances thermal efficiency by minimizing radiative heat losses and ensuring more energy is transferred to the working fluid, but also substantially reduces thermal NO_x formation, as the Zel'dovich mechanism becomes active only at temperatures around 1800 K. Furthermore, the gradual oxidation process ensures a more complete combustion of the fuel, effectively minimizing the release of unburnt hydrocarbons and carbon monoxide. This improvement in combustion efficiency is particularly beneficial for addressing environmental concerns as it reduces emissions of NO_x , unburnt hydrocarbons, CO, and soot,

while also lowering greenhouse gas emissions associated with the fuel's carbon content.

MILD combustion provides a highly efficient and low-emission approach to energy conversion by operating at lower and more uniform temperatures, which not only reduces NO_x formation and incomplete combustion products but also minimizes thermal stresses on equipment, extending its lifespan and reliability. Its adaptability to a wide range of low- and zero-carbon fuels, including hydrogen, biofuels, and syngas, allows for flexible fuel switching without compromising performance, making it particularly valuable for decarbonizing energy-intensive sectors. Combined with improved thermal efficiency, enhanced safety, and scalability, MILD combustion represents a promising strategy for transitioning to cleaner, more sustainable combustion systems.

1.2.3 Staged combustion

Staged combustion, also known as rich-lean combustion, is an effective technique for reducing NO and unburned NH_3 formation (when using NH_3 as fuel) by controlling flame temperature and reaction chemistry. This approach involves dividing the combustion process into multiple stages by introducing either part of the fuel (fuel staging), the oxidizer (air staging), or both into the primary combustion zone. The remaining portion is then supplied downstream into a secondary zone, allowing for better control over temperature profiles and chemical reactions to minimize pollutant formation.

In the first stage of fuel staging, combustion occurs under slightly fuel-lean conditions, resulting in high NO_x production and elevated temperatures. In the second stage, additional fuel is introduced, creating a fuel-rich environment within the reburn zone [88]. Under these fuel-rich conditions, a reducing atmosphere is established, promoting the reaction of the excess fuel with the NO_x formed in the primary zone. This interaction leads to the effective breakdown and reduction of NO_x emissions.

Air-staging begins with a fuel-rich combustion zone in the primary stage, where the limited availability of oxidizer suppresses thermal NO_x formation by maintaining lower flame temperatures compared to conventional combustion

[108]. In the subsequent fuel-lean secondary stage, additional air is introduced to fully oxidize the unburned intermediates. The temperature in this stage remains lower than in the primary stage, effectively minimizing thermal NO_x formation.

Staged combustion is particularly effective in large-scale systems such as industrial furnaces, boilers, gas turbines, and internal combustion engines. Its effectiveness lies in the ability to control the stoichiometry at each stage, allowing for temperature and reaction rate management. The reduction in peak temperatures directly leads to lower thermal NO_x formation, which is a significant contributor to overall NO_x emissions in conventional combustion systems.

1.3 Combustion modeling and challenges

Developing advanced combustion technologies necessitates a thorough understanding of fundamental physical phenomena, including fluid dynamics, chemical kinetics, transport phenomena, and heat transfer. Computational simulations have become a critical tool in modern scientific research, allowing researchers to investigate complex processes by accurately replicating system behaviors under various conditions [110]. These simulations provide essential insights into phenomena that are difficult or impossible to observe experimentally, thereby reducing the time and cost related concerns associated with traditional experimental methods.

In the field of combustion, numerical modeling aims to accurately represent the complex chemical and physical processes governing combustion, allowing for the prediction of critical aspects such as flame structure, temperature distribution, pollutant formation, and overall system efficiency. Several modeling approaches have been developed over the years, each offering unique advantages and facing specific limitations.

Direct Numerical Simulation (DNS) is one of the most precise techniques for combustion modeling, as it solves the complete set of governing equations, species transport, and energy conservation, without relying on turbulence models. Although DNS provides highly detailed insights into combustion physics, its extreme computational cost – linked to time-space resolution requirements dictated by the smallest fluid/chemical time scale – limits its application to small-scale

systems. On the other hand, Large Eddy Simulation (LES) offers a more practical approach by explicitly resolving larger turbulent structures while modeling smaller scales, striking a balance between accuracy and computational efficiency. However, despite being more feasible than DNS, LES still presents significant computational challenges when applied to industrial-scale systems.

Reynolds-Averaged Navier-Stokes (RANS) modeling remains the most widely used approach for engineering applications due to its relatively low computational cost. By employing turbulence models like $k - \epsilon$ and $k - \omega$, RANS models all turbulent structures through averaged equations, making it efficient for both steady and unsteady scenarios with simpler flow configurations. However, its accuracy declines when applied to complex, unsteady, or highly turbulent flows where more advanced techniques like LES or DNS are necessary. To address these limitations, hybrid approaches such as RANS/LES and LES/DNS have been developed, offering enhanced accuracy for multi-scale problems by effectively capturing varying turbulence intensities across the computational domain.

Despite significant advancements in combustion modeling, substantial challenges persist. Combustion is inherently a nonlinear, multi-scale, and multi-physics phenomenon. The nonlinear nature of chemical reactions makes it difficult to accurately predict ignition and extinction, especially in regimes dominated by auto-ignition and homogeneous reactions—such as MILD combustion—where chemical and flow time scales are comparable ($Da \approx 1$). This poses challenges for turbulence–chemistry decoupling methods like Flamelet Progress Variable model. In addition, simulating fluid–flow phenomena is further complicated by the wide spectrum of length and time scales that arise even in seemingly simple problems [14]. Another major challenge lies in accurately modeling chemical kinetics. Beyond solving the continuity, momentum, and energy equations, combustion simulations often require solving thousands of chemical reactions involving hundreds of species, leading to a high-dimensional problem commonly referred to as the curse of dimensionality.

Recent advancements in high-performance computing (HPC) have significantly enhanced computational fluid dynamics (CFD) tools, providing solutions to many challenges associated with combustion modeling. While these advancements have made CFD tools more accessible and effective, they also contribute to

substantial energy consumption and, consequently, a significant carbon footprint. The establishment of various online repositories hosting extensive datasets of high-fidelity CFD simulations further highlights this issue [16]. For instance, a case study on turbulent Reynolds number of $Re_\tau = 5200$ channel flow DNS reported an estimated carbon footprint of approximately 10^6 kg CO₂, equivalent to the emissions from a round-trip Boeing 777 flight between New York and Beijing [110]. Although CFD tools initially contributed to reducing carbon emissions through improved design and optimization processes, the increasing complexity of problems and the continuous push for a deeper understanding of physical phenomena through enhanced HPC capabilities have started to counteract this benefit. Therefore, a balanced approach is necessary, as CFD tools, despite their carbon emissions, remain essential for designing and optimizing advanced combustion technologies aimed at reducing overall emissions.

1.4 Machine learning and combustion science

Machine learning's ability to process vast amounts of data, learn from historical patterns, and make predictions makes it a powerful tool across various domains [89]. Some of the most prominent application areas of machine learning include:

- Predictive analytics and intelligent decision making
- Cybersecurity and threat detection
- Traffic prediction and transportation optimization
- Healthcare diagnostics
- Product recommendations
- Natural language processing (NLP) and sentiment analysis
- Image, speech and pattern recognition
- User behavior analytics

One possible objective of machine learning models is to develop a digital twin of a system using a minimal amount of data, effectively creating a function that maps input–output conditions to provide instantaneous responses [34]. The concept of the "*digital twin*" (DT) was first introduced by Grieves in 2002 [39] within the framework of Product Lifecycle Management (PLM). The original idea envisioned every system as comprising two parts: the physical system, which already exists, and a virtual system containing all information about the physical counterpart—essentially creating a mirrored relationship between the real and virtual domains. The concept was later adopted by NASA [94] in its aeronautics and aerospace technology roadmaps, where DTs were defined as integrated, multiphysics and multiscale simulations of a vehicle or system, continuously updated using physical models, sensor data, fleet history, and other sources to replicate the lifecycle of the real asset. These ultra-realistic digital counterparts may include interdependent subsystems such as propulsion and energy storage, avionics, life support, vehicle structure, and thermal management. Since the early 2000s, digital twins have been increasingly applied across diverse fields, including aviation, healthcare, manufacturing, and precision medicine [6]. Despite variations in early interpretations, there is now broad consensus in the literature that a digital twin constitutes a living, intelligent, and evolving model capable of continuously predicting future states—including defects, damage, or failures—and enabling simulation of novel configurations through real-time interaction, communication, and synchronization with its physical twin [6].

Although there is broad consensus on the definition of a DT, the literature still shows confusion between *digital twins*, *digital shadows*, and *digital models*. A defining feature of a *digital twin* is the bidirectional flow of information between the physical and digital counterparts [43]. In this setup, not only can the physical system update the digital version, but the digital version can also influence and control the physical system, enabling real-time synchronization and feedback [46, 56].

By contrast, when the information flow is unidirectional, i.e., only from the physical system to the digital counterpart, the system is referred to as a *digital shadow* [46, 56]. Here, the digital replica continuously mirrors the physical system but lacks any feedback mechanism to affect it.

A third category often confused with DTs is *digital models* or *surrogate models*. Unlike DTs or digital shadows, digital models are static replicas of the physical system that cannot autonomously integrate real-time data. Data transfer, if any, is manual, with no real-time link in either direction [46, 56]. Such models are typically used to approximate the behavior of computationally expensive physics-based simulations using data-driven methods trained on a predefined dataset [101].

This distinction also raises the question of how DTs differ from surrogate or simulation models, since all replicate physical system behavior in a virtual environment. The key difference lies in purpose and interaction: simulation models are digital prototypes that predict how a system might behave under assumed conditions, relying on initial parameters and assumptions rather than tracking and synchronizing with the real-time states of a specific physical system [111].

While many cases described as digital twins in the literature are actually digital shadows or surrogate models, the concept of creating virtual replicas of physical systems has undeniably gained significant traction across a wide spectrum of research areas. Brunton et al. [10] reviewed the application of machine learning models in fluid mechanics, highlighting how machine learning offers a modular and flexible framework suitable for addressing various challenges such as reduced-order modeling, experimental data processing, shape optimization, turbulence modeling, and control.

Advancements in combustion science and engineering have produced extensive datasets from large-scale simulations, high-resolution experiments, and sensors. Coupled with machine learning's proven effectiveness in data analytics, a new era of data-driven research has emerged: combustion machine learning [50]. Combustion science is inherently interdisciplinary, combining turbulent fluid mechanics and chemical kinetics. The presence of complex chemical reactions makes modeling turbulent combustion even more challenging. Therefore, the application of machine learning in fluid mechanics serves as a valuable foundation for advancing research at the intersection of combustion and machine learning [116].




Fig. 1.4 The number of review articles published annually between 2014 and 2024 on ScienceDirect using the search query “*combustion*” AND (“*machine learning*” OR “*data mining*” OR “*feature extraction*”).

The recent trend in machine learning applications within the combustion field is illustrated in Figure 1.4, revealing a substantial exponential increase in the number of review articles published over the past decade. This surge demonstrates the growing interest and success of machine learning techniques in addressing complex combustion problems. These techniques have been applied to various aspects of combustion datasets with different objectives. For instance, they have been employed to explore low-dimensional representations of the thermo-chemical manifold, which simplifies the description of chemical kinetics in turbulent reacting flows [79, 112, 113]. Huang et al. [48] used a POD-Galerkin technique to successfully predict the thermoacoustic limit cycle frequency and amplitude to understand the nonlinear heat release in a combustor. Additionally, machine learning has been utilized for analysing the dynamics of reactive systems [26], which is critical for understanding the transient behavior of combustion processes.

Moreover, significant progress has been made in using machine learning to accelerate numerical simulations of combustion systems by developing efficient surrogate models or reducing computational costs through adaptive algorithms [27, 66, 1]. Another important application of machine learning in this field is the identification of local combustion regimes, which helps in better characterizing

combustion processes and improving control strategies for various practical applications [67]. The increasing use of machine learning in combustion research highlights its potential to revolutionize the field by providing powerful tools for data-driven analysis, optimization, and predictive modeling.

1.5 Multi-fidelity surrogate modeling

Data-driven models provide an efficient means of capturing the fundamental characteristics of complex systems while significantly reducing computational costs, enabling simplified yet insightful analyses [57]. These models achieve this by constructing a function, often termed a digital model, which maps input-output relationships based on training datasets derived from numerical simulations and/or experimental data. In many cases, multiple physical models exist for estimating the same output quantity, each differing in approximation accuracy and computational cost [80]. A high-fidelity (HiFi) physical model provides highly accurate predictions necessary for precise analysis but comes with significant computational expense. In contrast, a low-fidelity (LoFi) physical model delivers faster but less precise solutions, making it computationally affordable but potentially limited in accuracy. Different levels of fidelity may be obtained for CFD simulations by varying the physical model, the size of the computational grid, the computational time step, the convergence level of the simulations, and/or combining numerical data with experimental results [58].

Nature of reduced order models require extensive datasets for proper training. However, in highly nonlinear systems like combustion—in which multiple input parameters interact over a wide range — generating a sufficient number of HiFi models can become impractical. Relying solely on LoFi models, while cost-effective, may lead to inaccurate or incomplete representations of the system [115].

To overcome the challenge of balancing computational efficiency and accuracy, multi-fidelity modeling techniques have been developed to integrate high-fidelity and low-fidelity models. These methods capitalize on the strengths of both approaches, using LoFi models for broad exploration and HiFi models

for refining critical regions, ultimately enhancing predictive reliability while reducing computational costs. Peherstorfer et al. [80] conducted a comprehensive literature review, highlighting that multi-fidelity models have been employed for over two decades to accelerate optimization and uncertainty quantification. Since then, their application has expanded across numerous scientific and engineering disciplines, demonstrating their growing impact.

In the context of data-driven modeling, data assimilation methods are essential to ensure both accuracy and reliability. Their role becomes particularly critical in multi-fidelity frameworks, where the objective is to exploit complementary representations of a physical system that differ in accuracy and computational cost. Peherstorfer et al. [80] classify data assimilation strategies into three principal categories: *adaptation*, *fusion*, and *filtering*. Adaptation seeks to calibrate low-fidelity data such that it more closely approximates high-fidelity information. Fusion leverages statistical methodologies to integrate data across fidelity levels, thereby exploiting the unique strengths of each. Filtering emphasizes the extraction of the most informative content while systematically excluding less relevant or redundant information.

Since this thesis focuses on the development of surrogate models, the discussion is focused to the data assimilation strategies tailored to such models. In the context of multi-fidelity surrogate modeling, data assimilation approaches can broadly be categorized into two classes: *deterministic* and *non-deterministic (stochastic)* strategies. Deterministic approaches are typically employed when the underlying process is well understood and not strongly affected by random variability, as they establish a fixed functional relationship between input and output variables [29]. While these models aim to minimize discrepancies between observed data and model predictions, they do not explicitly account for uncertainty [102, 35]. Representative examples include polynomial regression [24], Gaussian process models with a zero nugget effect [87], and various neural network architectures [36].

In contrast, non-deterministic approaches explicitly incorporate variability and uncertainty in the modeling process [29]. Rather than yielding a single deterministic prediction, outputs are expressed as probability distributions conditioned on the input variables. Examples include Gaussian process models with

a non-zero nugget effect [3, 106] and Monte Carlo simulation-based methods [72]. These stochastic models generally achieve higher predictive accuracy than their deterministic counterparts, albeit at the cost of increased computational complexity [29].

Recent studies showcase the versatility of multi-fidelity frameworks across diverse fields. Yang et al. [109] developed a multi-fidelity reduced-order model for the rapid simulation and optimization of microfluidic concentration gradient generators. Perron et al. [84] applied a multi-fidelity approach to analyse high-dimensional displacement and stress fields in structural mechanics, addressing variations in discretization and topology. In naval architecture, Liu et al. [63] utilized a multi-fidelity Co-Kriging surrogate model for hull form optimization to improve resistance and wake performance in a Japan Bulk Carrier (JBC).

In atmospheric sciences, Nony et al. [75] integrated LES with Reynolds-Averaged Navier-Stokes transport equation (RANS-TE) models to enhance the accuracy of atmospheric flow dispersion predictions. In machine learning applications, Demo et al. [21] successfully combined DeepONet architecture with Proper Orthogonal Decomposition (POD) and gappy POD to improve sensor data predictions for parametric and nonlinear Navier-Stokes problems. Guo et al. [42] developed a multi-fidelity Gaussian process model inspired by hierarchical Kriging to predict flame frequency response using both harmonic and broadband forcing data. Lastly, Tong et al. [96] applied a POD-Hierarchical Kriging multi-fidelity framework to optimize the combustion performance of an aero-engine combustion chamber.

These studies highlight the growing impact of multi-fidelity modeling in addressing complex problems across various scientific and engineering domains, demonstrating its ability to balance computational efficiency with high-accuracy predictions. Despite these advancements, several critical challenges remain. One issue is that the common assumption that high-fidelity models represent the absolute truth, whereas, in reality, all models—regardless of fidelity—are approximations of physical systems. Furthermore, many existing multi-fidelity approaches assume a simple hierarchical relationship between models, failing to capture more intricate interdependencies. Another key challenge is broadening the scope of multi-fidelity frameworks beyond computational models to

incorporate additional data sources, including experimental measurements, expert knowledge, and empirical lookup tables. This necessitates not only determining which information sources to utilize but also optimizing where and under what conditions to sample them. Statistical techniques, such as experimental design and value of information analysis, offer valuable insights into these challenges, reinforcing the need for more sophisticated and adaptable multi-fidelity model management strategies [80].

1.6 Motivation and outline

The design and optimization of advanced combustion technologies rely extensively on both experimental studies and numerical modeling at various fidelity and cost. The data generated from these approaches play a critical role in shaping the development of new technologies that drive the transition toward decarbonization. However, the inherently nonlinear, multi-scale, and multi-physics nature of combustion imposes significant limitations on experimental and numerical data collection. Additionally, accurately capturing these complex phenomena in numerical simulations is highly challenging, as it requires detailed chemical mechanisms, which in turn demand substantial computational resources.

Over the past two decades, multi-fidelity reduced-order modeling has gained considerable attention due to its ability to balance computational efficiency and predictive accuracy. By combining data from models of varying fidelity, this approach reduces the cost of generating training datasets while minimizing the loss of accuracy. Its application is becoming increasingly relevant in fields like combustion, where both experiments and numerical simulations are costly and problem specific. The high expense associated with these methods often result in a limited amount of training data, which poses a challenge for data-driven models. Multi-fidelity techniques help mitigate this issue by leveraging lower-cost models while incorporating high-fidelity data to enhance overall predictive capability.

Although reduced-order modeling has become increasingly common in the combustion community, the application of multi-fidelity surrogate modeling remains limited due to the intrinsic complexities of combustion. While there has

been growing interest in multi-fidelity approaches, most efforts have focused on uncertainty quantification and optimization rather than on achieving accurate predictive capabilities for the underlying physical problem. Furthermore, many of the methodologies reported in the literature have not demonstrated broad applicability across different types of combustion systems. It is well recognized that surrogate models can provide highly accurate predictions for a specific problem, yet they often fail when input conditions or the system's intrinsic dynamics change. Another important gap is the integration of physical knowledge into data-driven surrogate models, which has not been clearly addressed in prior studies. Finally, the scarcity of experimental and numerical data in combustion further constrains progress in this area. Consequently, research has largely concentrated on bi-fidelity reduced-order models, with limited attention given to the development of comprehensive multi-fidelity frameworks.

This thesis aims to address these gaps within the combustion community, building on the proven effectiveness of multi-fidelity surrogate modeling in other research domains. The primary objective of this thesis is to develop a multi-fidelity reduced order model (MF-ROM) framework for combustion applications that delivers accurate predictions while maintaining a limited computational budget. In general, the fusion of low-fidelity and high-fidelity data in the construction of an MF-ROM responds to both an accuracy requirement and a computational budget availability. In other words, some HiFi simulations are replaced with LoFi ones to reduce the training cost, with the simultaneous goal of incurring in a limited accuracy loss of the resulting multi-fidelity reduced order model. Therefore, a fundamental question is: given a design of experiment (DoE), how many and which HiFi simulations can be replaced with LoFi ones to train an MF-ROM whose prediction error remains satisfactory? Another challenge associated with MF-ROMs lies in the regression methods used for model training. While these methods generally perform well within the boundaries of the DoE, their accuracy deteriorates when applied to extrapolation. Therefore, the interpolation and extrapolation capabilities of the proposed MF-ROM are examined. Additionally, to enhance prediction accuracy while reducing the number of required simulations, the potential of unsupervised learning techniques are explored. In particular, clustering is investigated as a method to identify distinct combustion

regimes by grouping similar data points, thereby improving model accuracy. Finally, the study assesses the impact of incorporating experimental data as a third level of fidelity, evaluating its effectiveness in enhancing prediction accuracy and further reducing computational costs.

The thesis is structured as follows:

- Chapter 2 provides the mathematical background of the machine learning techniques employed throughout the thesis
- Chapter 3 details the experimental campaign of the ULB furnace for the $\text{CH}_4\text{-H}_2$ blend.
- Chapter 4 explains the how the training and test data obtained for the MF-ROM both for the ULB furnace and the stagnation point reverse-flow combustor.
- Chapter 5 introduces the novel multi-fidelity reduced order model and its application with incremental sampling strategies on the ULB furnace for $\text{CH}_4\text{-H}_2$ fuel blend.
- Chapter 6 show the adaptability of the MF-ROM framework to different fuels and combustion systems through the stagnation point reverse-flow combustor for ammonia combustion.
- Chapter 7 investigates the impact of combustion regime identification and incorporating intrinsic physics into the MF-ROM framework using clustering techniques.
- Chapter 8 presents the integration of experimental data from Chapter 3 as a third level of fidelity and discusses the corresponding results.

Chapter 9 summarizes the key findings, presents the conclusion, and outlines the future research directions.

Additional information is provided to support the thesis. Appendix A provides the calculations related to prediction and error metrics.

Chapter 2

Mathematical background

This chapter summarizes the methodologies employed throughout the thesis, with a particular focus on data fusion within the multi-fidelity framework. At the core of this framework lies the integration of insights obtained from both HiFi and LoFi datasets. These datasets, while representing the same physical phenomena, often differ in terms of accuracy, physical modeling assumptions, geometry, dimensionality, and resolution. Effective data fusion seeks to identify and exploit shared intrinsic structures—often of reduced dimensionality—across these heterogeneous sources. The chapter covers Proper Orthogonal Decomposition, which is used to uncover latent spaces and correlations within the datasets; a manifold alignment technique that leverages these correlations to align HiFi and LoFi data; and CoKriging regression, which combines the aligned information from these latent spaces into a continuous, predictive surrogate model. Additionally, the chapter includes incremental sampling strategies, which address one of the most fundamental challenges in multi-fidelity modeling: for a given set of conditions being covered, determining how many HiFi simulations can be effectively replaced with LoFi simulations while maintaining acceptable predictive accuracy in the final MF-ROM. This chapter also presents the sparse sensing methodology employed to integrate sparse experimental data into the MF-ROM framework.

The primary objective of this thesis is to develop a multi-fidelity surrogate model that can be applied effectively to advanced combustion systems. The methodology adopted in this work is structured around three main components: (i) dimensionality reduction, (ii) manifold alignment, and (iii) regression. Dimensionality reduction is carried out using Proper Orthogonal Decomposition (POD), one of the most widely applied and effective techniques in the literature. Procrustes manifold alignment is then employed to ensure that the low-dimensional manifolds obtained from different fidelity levels are represented in a consistent and reliable manner. Finally, the input-output relationship is modeled using CoKriging as the regression technique. Sampling strategies are designed to reduce either the prediction error or the model uncertainty, while sparse sensing is introduced to incorporate high-fidelity information from experimental sensor data into the multi-fidelity surrogate model.

A critical aspect of multi-fidelity frameworks is understanding the relationships among different fidelity levels. In the literature, these differences are typically categorized into three groups: *model* differences, which arise from changes in the governing equations representing physical reality; *accuracy* differences, stemming from variations in discretization, dimensionality, or convergence; and *source* differences, which distinguish between experimental and numerical data [29]. In this context, the use of a simple discrepancy function is advantageous, particularly for regression models such as CoKriging, where the discrepancy is explicitly incorporated into the formulation.

Under normal circumstances, datasets at different fidelities are expected to represent the same physical system, with differences only attributable to model, accuracy, or source variations. However, in the event of an error where training datasets represent entirely different problems (e.g., combustion versus fluid dynamics), several diagnostic “check-points” can signal incompatibilities within the multi-fidelity reduced-order modeling framework combining POD, manifold alignment, and CoKriging. The first indicator arises during POD, where the eigenvalue spectra would reveal distinct modal energy distributions. The second could be observed during manifold alignment, where large residual errors would indicate a poor correspondence between manifolds. Finally, in CoKrig-

ing, near-zero correlations across fidelities would result in very high predictive uncertainty.

Taking into account the methodology and sub-models described in this chapter, sampling strategies provide guidance on both the placement and the number of high-fidelity data points required within the design space to achieve optimal performance. Overall, the multi-fidelity reduced-order model developed here represents one of the few such frameworks applied to combustion systems that can accurately predict system responses under varying conditions. Moreover, by incorporating sparse sensing, this framework moves beyond traditional bi-fidelity approaches, establishing itself as a "*true*" multi-fidelity reduced-order modeling framework.

2.1 Proper orthogonal decomposition

Combustion processes represent a highly nonlinear, multi-physics problem characterized by complex interactions among thermodynamic, chemical, and fluid dynamic variables — such as temperature, pressure, species mass fractions, and velocity fields — spanning a broad range of spatial and temporal scales. Detailed chemical kinetics further complicate modeling efforts, often requiring mechanisms with hundreds of reactive species and thousands of elementary reactions. For example, even a reduced mechanism for methane combustion may involve 23 species and 51 reactions [117]. As the dimensionality of the problem grows, especially in data-driven modeling frameworks, the modeling process becomes susceptible to the "curse of dimensionality". The curse of dimensionality refers to the rapid growth in computational cost and data requirements as the number of dimensions (variables or parameters) increases. In high-dimensional spaces, sampling, storing, or interpolating data becomes exponentially harder, making many algorithms inefficient or impractical. This issue significantly affects key components of machine learning applications in combustion considering that chemical kinetics involve hundreds of species and reactions, including the design of experiments, feature selection and extraction, model training, generalization, and prediction accuracy. Various techniques have been proposed in the literature

to mitigate these challenges. In the context of this thesis, Proper Orthogonal Decomposition (POD) is employed to mitigate these effects, identify the main patterns, and reduce the model efficiently without substantial loss of accuracy.

In the literature, Proper Orthogonal Decomposition (POD) is also referred to as the Karhunen–Loève decomposition [33] or Principal Component Analysis (PCA) [79], and it is commonly used for reducing the dimensionality of high-dimensional datasets; while PCA and POD are mathematically equivalent under a Euclidean inner product, POD constitutes the more general formulation by enabling physics-based inner products that better capture energy-containing structures in fluid and combustion systems. Consider a training dataset represented by a data matrix composed of n designs/snapshots, each having r physical degrees of freedom (i.e. independent variables) and corresponding to a design condition $\mathbf{d} = \{d_i\}_{i=1, \dots, N_{\text{desVars}}}$ within the DoE. Each of these n designs may represent a simulation, and each simulation is arranged as a column vector $\mathbf{x} \in \mathbb{R}^p$, where the length p equals the number of computational cells multiplied by the number of Quantities of Interest (QoIs) (i.e., temperature and chemical species mass fractions), so $p = n_{\text{cells}} \times n_{\text{QoI}}$. The data matrix $\mathbf{X}$ is formed by collecting these n simulations, i.e., $\mathbf{X} = \{\mathbf{x}_i\}_{i=1, \dots, n}$. In combustion problems, it is important to note that $p \gg n$.

The assumption of a large number of interrelated variables allows the decomposition of the data matrix ($\mathbf{X} \in \mathbb{R}^{p \times n}$) into a matrix that captures physics invariant information ($\Phi \in \mathbb{R}^{p \times n}$), and a matrix that contains parameter-dependent information ($\mathbf{Z} \in \mathbb{R}^{n \times n}$), as follows [52]:

$$\mathbf{X}(\mathbf{r}, \mathbf{d}) = \Phi(\mathbf{r})\mathbf{Z}^T(\mathbf{d}). \quad (2.1)$$

POD enables the extraction of the most energetic modes (modes containing the most important information) from the initial training dataset by projecting it onto a lower-dimensional subspace, preserving the dominant patterns while filtering out less significant variations. POD projects the columns vector $\mathbf{x}$ onto a rotated basis obtained from the singular value decomposition (SVD) of the matrix $\mathbf{X}$, which corresponds to solving the eigenproblem of the covariance matrix:

$$\mathbf{C} = \frac{1}{n-1} \mathbf{X}^T \mathbf{X}. \quad (2.2)$$

Dimensionality reduction is achieved by retaining only the first k eigenvectors, resulting in the truncated matrices Φ_k and Z_k with dimensions $(p \times k)$ and $(k \times n)$, where $k < \min(n, p)$. The last $p-k$ eigenvectors can be discarded while incurring in minimal information loss because the POD algorithm ensures the minimization of the Frobenius norm of the reconstruction error:

$$\mathbf{E} = \mathbf{X} - \Phi_k \mathbf{Z}_k^T. \quad (2.3)$$

Several approaches exist in the literature for determining the value of k . In this framework, the reduced dimension k is selected based on the Relative Information Content (RIC) criterion [85], which quantifies the total variance captured by the POD modes. Typically, a threshold of 99% or higher is adopted. With this selection, the data matrix $\mathbf{X}$ can be approximated by the truncated matrices Φ_k and Z_k (see Figure 2.1, with dimensions $(p \times k)$ and $(k \times n)$, respectively, as:

$$\mathbf{X} \approx \Phi_k \mathbf{Z}_k^T. \quad (2.4)$$




Fig. 2.1 Schematic drawing of POD.

The columns of Φ_k are referred to as the POD modes, or eigenflames [4], and the matrix Z_k is known as the POD coefficients matrix. Since the eigenvectors of $\mathbf{C}$ are orthogonal and the principal components are linear combinations of the original variables, the approximated/reconstructed system response $\mathbf{x}_{rec}$ corresponding to a given set of POD coefficients $\mathbf{z}$ can be calculated as:

$$\mathbf{x}_{rec} = \Phi_k \mathbf{z}. \quad (2.5)$$

Thus, the ROM is built by regressing the POD coefficients $\mathbf{z}$ over the design space, i.e., by building a model $\alpha^k = \alpha^k(\mathbf{d}; \mathbf{z})$ for each coefficient k , enabling one to estimate/predict the system response at virtually no cost for any unseen design condition $\mathbf{d}^*$:

$$\mathbf{x}_{pred}(\mathbf{d}^*) = \Phi_k \alpha(\mathbf{d}^*). \quad (2.6)$$

2.2 Manifold alignment

Multi-fidelity frameworks effectively integrate heterogeneous datasets that may differ in resolution, dimensionality, and underlying physical models, and manifold alignment is an essential part of this process. Without alignment, direct fusion of these datasets can lead to degraded model accuracy and inconsistencies in the learned representations. Manifold alignment seeks to uncover a shared latent space where the intrinsic relationships between different fidelities are preserved, facilitating more accurate data fusion and improved generalization in reduced-order modeling [105]. This approach is particularly important in complex systems like combustion, where high-dimensional, multi-physics phenomena exacerbate the challenges of directly comparing or combining datasets across fidelities [79].

Manifold alignment allows to identify a correlation structure between datasets and therefore transfers information from one manifold to the other by aligning their underlying low-dimensional representation [104]. The main assumption behind this technique is that these datasets do share a common low-dimensional representation which allows the comparison of the embedded data due to their intrinsic similarity [83]. Manifold alignment is applied as a useful technique in many different disciplines such as protein alignment [103], image matching and classification [44, 41] or automatic machine translation [22]. In the literature, there are two types of manifold alignment approaches. The first method searches for a shared latent space with a semi-supervised approach by constructing a joint graph Laplacian and the resultant latent space should cover the individual and shared characteristics of both manifolds [83]. The other method uses dimensionality

reduction techniques to discover the intrinsic latent space for each dataset and then apply Procrustes analysis [64] to align both spaces.

Considering the high-dimensions of each dataset used within this framework, dimensionality reduction is unavoidable. Hence, the Procrustes manifold alignment technique is chosen as the logical approach. As mentioned before, the first step of this method is applying POD to identify the intrinsic latent space for each dataset. Then, the Procrustes analysis is used to adjust the translational, rotational, and scaling components from one manifold to achieve an optimal alignment with the other.

Consider, HiFi and LoFi datasets used in the framework are $\mathbf{X} \in \mathbb{R}^{p \times n}$ and $\mathbf{Y} \in \mathbb{R}^{q \times m}$, respectively. Here n and m indicate the number of observations/snapshots with $n < m$, while p and q represent the value obtained by multiplication of grid points and features, $p = n_{cells,hf} \times n_{features}$ and $q = n_{cells,lf} \times n_{features}$. 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 dataset includes a linked subset $\mathbf{Y}_L \in \mathbb{R}^{q \times n}$ which shares the same input design parameters as the HiFi dataset, and the rest of the data belongs to the unlinked subset, $\mathbf{Y}_U \in \mathbb{R}^{q \times (m-n)}$, hence $\mathbf{Y} = [\mathbf{Y}_L, \mathbf{Y}_U]$. The first step of this method requires obtaining the reduced manifold for HiFi and LoFi datasets. Following the application of POD, for a latent subspace dimension of k , the eigenflames $\Phi \in \mathbb{R}^{p \times k}$ and $\Psi \in \mathbb{R}^{q \times k}$, with POD coefficients $\mathbf{Z} \in \mathbb{R}^{k \times n}$ and $\mathbf{W} \in \mathbb{R}^{k \times m}$, are obtained for HiFi and LoFi inputs, respectively. Consequently, the LoFi manifold $\mathbf{W}$ has also linked and unlinked coefficients subsets, $\mathbf{W} = [\mathbf{W}_L, \mathbf{W}_U]$. This orthogonal Procrustes problem is solved using the procedure of Wang and Mahadevan [28] which seeks the manifold $\mathbf{T}$ minimising $\|\mathbf{Z} - \mathbf{T}\|$.

The first step is to translate the linked manifolds $\mathbf{Z}$ and $\mathbf{W}_L$ to make sure that both were centered at the origin. To center the $\mathbf{W}_L$, the following shift is required:

$$\overline{\mathbf{W}}_L = \frac{1}{p} \sum_{i=1}^p \mathbf{W}_{L,i} = 0. \quad (2.7)$$

After translation, scaling and rotation are performed by calculating and applying a scaling factor (s) and a rotation matrix ($\mathbf{Q}$). These parameters are obtained through the computation of Singular Value Decomposition of $\mathbf{W}_L \mathbf{Z}^T$:

$$\mathbf{U} \mathbf{\Sigma} \mathbf{V}^T = \mathbf{W}_L \mathbf{Z}^T. \quad (2.8)$$

Here $\mathbf{U} \in \mathbb{R}^{k \times k}$, $\mathbf{V} \in \mathbb{R}^{k \times k}$, and $\mathbf{\Sigma} \in \mathbb{R}^{k \times k}$ are the left and right singular vectors, and the diagonal matrix containing the singular values, respectively. The scaling factor and the rotation matrix are calculated as follows:

$$s = \frac{\text{tr}(\mathbf{\Sigma})}{\text{tr}(\mathbf{W}_L \mathbf{W}_L^T)}. \quad (2.9)$$

$$\mathbf{Q} = \mathbf{V} \mathbf{U}^T, \quad (2.10)$$

Finally, the transformed manifold is obtained as:

$$\mathbf{T} = s \mathbf{Q} \mathbf{W}. \quad (2.11)$$

Each step of the Procrustes manifold alignment process is illustrated in Figure 2.2. The original datasets are presented in the top left panel. Following translation, the HiFi and translated linked LoFi manifolds are shown in the top right panel. After scaling is applied, the HiFi, scaled linked, and scaled unlinked LoFi manifolds are displayed in the bottom left panel. Finally, the bottom right panel shows the HiFi manifold alongside the aligned LoFi manifold, including both linked and unlinked components.

Manifold alignment introduces a transformed matrix which is a representation of the LoFi data, $\mathbf{Y}$, in the latent space of the HiFi data, $\mathbf{X}$. As a result of this alignment, both manifolds are embedded in the same space, which allows an easy comparison and the ability to be used in the multi-fidelity framework. By leveraging manifold alignment, it becomes possible to construct multi-fidelity surrogate models that exploit complementary information across fidelities, enhancing predictive performance while managing computational costs.




Fig. 2.2 Steps of the Procrustes manifold alignment technique: the original datasets (*top left*), translated manifolds (*top right*), scaled manifolds (*bottom left*), and aligned manifolds (*bottom right*).

2.3 CoKriging regression

Once the shared latent representation is obtained, with POD coefficients matrices $\mathbf{Z}$ and $\mathbf{T}$ for the HiFi and LoFi data respectively, a regression model is required to relate the POD coefficients with the design space, i.e., a model of the form $\alpha^k = \alpha^k(\mathbf{d}; \mathbf{z}, \mathbf{t}) - \mathbf{t}$, where $\mathbf{d}$ is the design parameters, $\mathbf{z}$ is the high-fidelity manifold, and $\mathbf{t}$ is the low-fidelity manifold, is added due to the additional fidelity data – for each POD coefficient k .

The Gaussian Process Regression (GPR) [106], also known as Kriging when a spatial covariance function or kernel is used, is a very popular machine learning technique used for predicting data points based on probabilistic models that define distributions over functions and estimate uncertainties. Kriging is a two-step process: training and prediction. The model finds the underlying patterns between input-output from the available data during training. After model training, for a given new input, the model provides a predictive distribution over possible function values at that input, including a mean value and a measure of uncertainty.

Kriging is used to predict a single variable based on the values of that same variable at observed locations. However, in the multi-fidelity approach, a regression model is required to predict one variable, combining the values of two or more correlated variables at observed locations. For this reason, CoKriging [53, 82, 37], also known as collocated Kriging, is preferred for these applications. CoKriging is an extension of Kriging and it improves the prediction accuracy by taking advantage of the spatial relations between correlated variables. This method is particularly useful in situations where one variable is measured more accurately than the others, making it a good choice for multi-fidelity frameworks. In cases where two correlated random fields are considered, simple CoKriging is employed. However, when extending to scenarios involving more than two correlated fields—corresponding to multiple fidelity levels—a hierarchical formulation is required, necessitating the use of multi-level co-Kriging to properly account for the nested correlation structures across fidelities.

2.3.1 Simple CoKriging

Consider the two datasets introduced in the previous section: the HiFi dataset $\mathbf{X} \in \mathbb{R}^{p \times n}$ and the LoFi dataset $\mathbf{Y} \in \mathbb{R}^{q \times m}$ with corresponding output vectors $\mathbf{x}$ and $\mathbf{y}$, respectively. After applying the POD and manifold alignment, the datasets are projected into a shared latent space, yielding $\mathbf{Z} \in \mathbb{R}^{k \times n}$ for HiFi data and $\mathbf{T} \in \mathbb{R}^{k \times m}$ for LoFi data, with associated outputs $\mathbf{z}$ and $\mathbf{t}$. As described in Section 2.2, the design space of the HiFi dataset, denoted by $\mathbf{P}_L$, is a subset of the design space for the LoFi dataset, $\mathbf{P}$, resulting in a nested design structure. These transformed datasets are then concatenated for subsequent multi-fidelity modeling as follows:

$$\mathbf{D} = \begin{pmatrix} \mathbf{P} \\ \mathbf{P}_L \end{pmatrix} = (\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(m)}, \mathbf{p}_L^{(1)}, \dots, \mathbf{p}_L^{(n)}). \quad (2.12)$$

The predicted value at a given point in the design space $\mathbf{D}$ is modeled as a realization of a Gaussian random variable, thereby defining the associated Gaussian random field as follows:

$$\mathbf{R} = \begin{pmatrix} \mathbf{T}(\mathbf{P}) \\ \mathbf{Z}(\mathbf{P}_L) \end{pmatrix} = (\mathbf{t}(\mathbf{p}^{(1)}), \dots, \mathbf{t}(\mathbf{p}^{(m)}), \mathbf{z}(\mathbf{p}_L^{(1)}), \dots, \mathbf{z}(\mathbf{p}_L^{(n)})). \quad (2.13)$$

The CoKriging regression model adopted in this study follows the autoregressive model of Kennedy and O'Hagan [53] which is based on the following assumption:

$$\text{cov}\{\alpha_{\text{HiFi}}(\mathbf{d}), \alpha_{\text{LoFi}}(\mathbf{d}') | \alpha_{\text{LoFi}}(\mathbf{d})\} = 0 \text{ for all } \mathbf{d} \neq \mathbf{d}'. \quad (2.14)$$

According to this Markov property, once a HiFi observation is available at a given design point, no additional information can be gained from the corresponding LoFi data at that point [31]. In this framework, the Gaussian processes $\alpha_{\text{LoFi}}(\mathbf{d})$ and $\alpha_{\text{HiFi}}(\mathbf{d})$ characterize the local features of the cheap (LoFi) and expensive (HiFi) responses, respectively. The autoregressive model assumes that the expensive response can be approximated as a scaled version of the cheap response via a factor ρ , augmented by a Gaussian process $\delta(\mathbf{d})$ that captures the discrepancy between $\rho\alpha_{\text{LoFi}}(\mathbf{d})$ and $\alpha_{\text{HiFi}}(\mathbf{d})$:

$$\alpha_{\text{HiFi}}(\mathbf{d}) = \rho \alpha_{\text{LoFi}}(\mathbf{d}) + \delta(\mathbf{d}). \quad (2.15)$$

In standard Kriging, the covariance matrix is constructed solely based on a single-fidelity input-output dataset, making its computation relatively straightforward. In contrast, CoKriging involves a more intricate covariance structure due to the inclusion of cross-correlations between more than one fidelity levels. This results in a higher number of hyperparameters to estimate in addition to the scaling parameter ρ . The resulting CoKriging covariance matrices are defined as follows:

$$\text{cov}\{\mathbf{T}(\mathbf{D}), \mathbf{T}(\mathbf{D})\} = \text{cov}\{\alpha_{\text{LoFi}}(\mathbf{P}), \alpha_{\text{LoFi}}(\mathbf{P})\}, \quad (2.16)$$

$$\text{cov}\{\mathbf{Z}(\mathbf{D}_L), \mathbf{T}(\mathbf{P})\} = \text{cov}\{\rho\alpha_{\text{LoFi}}(\mathbf{P}) + \delta(\mathbf{P}), \alpha_{\text{LoFi}}(\mathbf{P}_L)\}, \quad (2.17)$$

$$\text{cov}\{\mathbf{Z}(\mathbf{P}_L), \mathbf{Z}(\mathbf{P}_L)\} = \text{cov}\{\rho\alpha_{LoFi}(\mathbf{P}_L) + \delta(\mathbf{P}_L), \rho\alpha_{LoFi}(\mathbf{P}_L) + \delta(\mathbf{P}_L)\}. \quad (2.18)$$

Both α_{HiFi} and α_{LoFi} require the specification of a kernel function (or covariance), which is typically a radial basis function (RBF), along with the calibration of their hyper-parameters, in addition to the calibration of the scaling factor. Since the hyper-parameters values have a considerable effect on the fitting properties of the model — their choice is indeed a model calibration problem — a common approach is to optimize them. To this end, the Bayesian approach proposed by Le Gratiet [60] was utilized and then implemented in openMDAO [38], which provides a closed-form expression for the scaling factor and a numerically tractable marginal likelihood of the hyper-parameters, whose maximization returned the best estimate for the parameters given the training data.

Once fitted, the CoKriging model returns a mean prediction function that estimates $\alpha_{HiFi}(\mathbf{d})$ and its uncertainty $\sigma_{\alpha_{HiFi}}(\mathbf{d})$, with the assumption that the uncertainty associated with the LoFi data α_{LoFi} and the residual δ are uncorrelated and can be defined as:

$$\sigma_{\alpha_{HiFi}}(\mathbf{d}) = \sqrt{\sigma_{\alpha_{LoFi}}^2(\mathbf{d}) + \sigma_{\delta}^2(\mathbf{d})}. \quad (2.19)$$

This uncertainty directly reflects the model's reliability: it is typically low in regions near high-fidelity training data, indicating confident predictions, and higher in areas relying more on low-fidelity information or extrapolation, signaling reduced reliability. The estimated uncertainty also captures the strength of correlation between fidelity levels, with weak correlations increasing predictive variance and lowering confidence. In theory, prediction accuracy could be improved by incorporating more LoFi points; however, this also depends on the discrepancy between the HiFi and LoFi datasets, as CoKriging relies on a discrepancy function to relate the HiFi and LoFi manifolds.

An example of CoKriging regression is shown in Figure 2.3 using the toy datasets provided in [31]. The autoregressive feature is noticeable from the fact

that the mean prediction goes through the HiFi samples and the uncertainty in those locations is zero.




Fig. 2.3 Example of CoKriging regression using the toy datasets provided in [31].

2.3.2 Multi-level CoKriging

In the cases that there are more than two levels of data, a multi-level approach is necessary to incorporate all the information coming from multiple sources. Assuming that there are l levels of fidelity, for all $t = 1, \dots, l$, the t^{th} output is modeled by the Gaussian process α_t . Similar to simple CoKriging, both the input and the output data are stacked from the less accurate to the most accurate. For each level of data, the conditional distribution of the Gaussian process $\alpha_t(\mathbf{d})$ is determined from $\alpha_{t-1}(\mathbf{d})$.

Even when dealing with multiple levels of fidelity, it is still possible to assume the Markov property introduced by Kennedy and O'Hagan [53] as follows:

$$\text{cov}\{\alpha_t(\mathbf{d}), \alpha_{t-1}(\mathbf{d}') | \alpha_{t-1}(\mathbf{d})\} = 0 \text{ for all } \mathbf{d} \neq \mathbf{d}'. \quad (2.20)$$

This property still implies if the value of the higher fidelity data is available, nothing can be learnt about $\alpha_t(\mathbf{d})$ from the lower fidelity data. This assumption allows to use the auto-regressive model once again as following [60]:

$$\alpha_t(\mathbf{d}) = \rho_{t-1} \alpha_{t-1}(\mathbf{d}) + \delta_t(\mathbf{d}), t = 2, \dots, l. \quad (2.21)$$

Here, $\delta_t(\mathbf{d})$ denotes a Gaussian process that is independent of $\alpha_{t-1}, \dots, \alpha_1$ and ρ_{t-1} represents the scale factor linking two successive Gaussian processes, α_t and α_{t-1} . The hyper-parameters in the multi-level CoKriging framework are estimated using the same Bayesian approach as in simple CoKriging. Notably, the generalized Bayesian estimation yields a numerical complexity that is equivalent to performing l independent Kriging estimations [60].

At each fidelity level t , the design set D_t is selected such that $D_t \subseteq D_{t-1}$. While this nested structure is not strictly required, it enables closed-form solutions for the parameter estimation equations.

2.4 Sampling strategies

A critical aspect of Reduced Order Model training lies in the selection of design points within the design space that yield the most accurate predictive performance. The quantity and distribution of training samples required to achieve a target level of accuracy are highly problem-dependent, influenced by factors such as the underlying model's nonlinearity, dimensionality, noise levels, and smoothness of the input/output mapping, as well as the chosen training strategy [62]. Whether derived from physical experiments or numerical simulations, training data inherently contain uncertainties and deviate from the true solution due to measurement or modeling errors [31]. Therefore, the central objective is to sample the design space efficiently, ensuring that the resulting surrogate model provides a reliable global approximation of the underlying system behavior.

The most basic sampling strategy is *simple random sampling*, where each point in the design space has an equal probability of selection. This typically involves enumerating the design space and selecting a subset using a random number generator. *Stratified sampling* improves upon this by dividing the design space into subregions, or *strata*, based on specific characteristics, and drawing random samples from each to ensure balanced representation. Alternatively, *uniform sampling methods* aim to maximize the minimum pairwise distance between

points, ensuring even coverage of the space [74]. A widely used technique, *Latin Hypercube Sampling (LHS)*, enhances stratified sampling for multidimensional spaces by dividing each input variable's range into equally probable intervals and selecting one sample from each. This ensures that each sample is the only one in its interval along each dimension, providing efficient space-filling properties with fewer samples than purely random methods [71].

However, these strategies follow an a-priori approach, distributing training points without insight into the underlying input/output mapping, which may result in inefficient sampling. Therefore, the ROM training strategies have moved from static to dynamically updating models, due to the ability to improve their fitting capability by adaptive sampling or training [93]. The importance of adaptive sampling strategies lies in adding training points to the locations in the design space where they are the most impactful, keeping the number of model evaluations needed to build the surrogate as low as possible. Clearly, the impact of a new sample must be estimated from the knowledge gained during the process up to that point. Therefore, adaptive sampling strategies require a method to choose the coordinates of a new training point, which in turn requires a decision-making metric. This metric should target and quantify a potential/estimated improvement of the surrogate model prediction and/or a reduction of the associated uncertainty.

To this end, adaptive sampling strategies have been combined with multi-fidelity surrogate models, with the goal of deciding on a cost/effect balance basis whether an expensive HiFi sample is needed in place of a cheaper one. An augmented expected improvement for multi-fidelity Kriging was proposed by Huang et al. [47], by using multi-fidelity data to minimize the total evaluation cost by finding the overall optimum of the real system. Pellegrini et al. [81] used a sampling strategy that sequentially adds LoFi or both HiFi and LoFi samples depending on the prediction uncertainty and on a parameter based on their computational cost ratio. Cai et al. [12] used an algorithm to sample both HiFi and LoFi at every iteration based on the cross-validation error and a Voronoi partition. Liu et al. [61] applied a sampling method updating the dataset of the corresponding fidelity level that provides the maximum reduction of the MF uncertainty. Kimpton et al. [54] combined expected improvement with a repulsion function that ensures a filling design space, in a multi-level

fashion. Guénot et al. [40] used a basis improvement based adaptive sampling strategy that enhances POD-based models by selecting new sample points that improve the modal basis while exploring the parameter space. The influence of each snapshot on the modes is quantified using singular-value-weighted angles, and candidate points are evaluated based on proximity to influential snapshots. Efficient computation is achieved through thin SVDs of small matrices, avoiding costly full SVDs of the full snapshot set.

2.4.1 Adaptive sampling based on prediction error

The first incremental sampling method proposed in this work is rather a benchmark strategy that requires the availability of the full HiFi and LoFi datasets $\mathbf{X}$ and $\mathbf{Y}$. Starting from a baseline/minimal set of HiFi data and the full set of LoFi data, the choice of a new HiFi training point is driven by the calculated prediction accuracy improvement. Therefore, given the initial set of HiFi samples $D^0 \in \mathbf{X}$, the HiFi dataset is progressively augmented by adding the simulation data $\mathbf{x}'(\mathbf{d}')$ corresponding to the DoE coordinates $\mathbf{d}'$ such that, at the $i + 1$ -th increment, $\mathbf{x}'$ is:

$$\underset{\mathbf{x}' \in (\mathbf{X} - D^i)}{\operatorname{argmin}} \mathcal{E}(\mathbf{x}') := \{\mathbf{x}' | \forall \mathbf{x} : \mathcal{E}(\mathcal{M}[D^i + \mathbf{x}, \mathbf{Y}]) > \mathcal{E}(\mathcal{M}[D^i + \mathbf{x}', \mathbf{Y}])\}, \quad (2.22)$$

where $\mathcal{E}(\mathcal{M}[D^i + \mathbf{x}, \mathbf{Y}])$ is the Normalized Root Mean Square (NRMS) error defined in Appendix A computed between the test vector and the data vector predicted by the model $\mathcal{M}$ obtained using the HiFi dataset D^i augmented with a new sample $\mathbf{x}$. Therefore, the chosen new HiFi sample is the one that maximizes the prediction accuracy improvement among the pool of available samples.

Clearly, at a given stage i , this strategy requires the generation of all possible models $\mathcal{M}$ built with the HiFi training dataset D^i plus each one of the remaining training data $\mathbf{x}$. Therefore, besides requiring the full HiFi dataset up front, it encompasses the generation and evaluation of many models at each sampling increment, making it unfeasible in a practical context.

2.4.2 Adaptive sampling based on model uncertainty

In a practical context, a different approach is needed since the full HiFi dataset is generally not available upfront because of its computational cost. At the same time, the full set of LoFi data is considered to be available, being it much cheaper, although this might not be the case in general. In this case, the HiFi sample coordinates of successive dataset increments should be identified using heuristic methods.

The uncertainty of the CoKriging model is a good indicator of missing information, i.e., of the need to add a sample in the proximity of the region where it is larger. In fact, the addition of a HiFi sample will improve the CoKriging discrepancy model where it has the largest uncertainty.

Therefore, a new sample $\mathbf{x}'(\mathbf{d}')$ is added corresponding to the DoE coordinates $\mathbf{d}'$ such that, at the $i + 1$ -th increment, it is:

$$\operatorname{argmax}_{\mathbf{d}_j} \left[\sum_k^{N_{modes}} \Sigma_k \sigma_{\alpha^k}(\mathbf{d}_j) \right] \quad j = 1, \dots, p - i, \quad (2.23)$$

where the overall prediction uncertainty in the DoE coordinates $\mathbf{d}_j$ is computed by summing the uncertainty associated with each POD coefficient k , weighted with the corresponding eigenvalue Σ_k , i.e., the POD coefficients' uncertainty are weighted with the importance of their mode.

2.5 Sparse sensing

Sparse sensing (SpS) leverages the inherent sparsity present in many physical systems to enable efficient data acquisition and reconstruction with significantly fewer measurements than traditional approaches. The core principle of SpS is that complex, high-dimensional systems often exhibit low-dimensional structure, allowing the full system state to be represented by a small number of basis vectors [9]. For machine learning algorithms that seek to extract such low-dimensional patterns, SpS is particularly advantageous, as it enables accurate reconstruction of high-dimensional states from limited measurements by exploiting coherent underlying structures [69]. In this context, the full state can be approximated

as an unknown linear combination of basis vectors, and the objective of sparse sensing is to solve the resulting linear system:

$$\mathbf{y} = \mathbf{C}\mathbf{x}, \quad (2.24)$$

where $\mathbf{y} \in \mathbb{R}^s$ denotes the vector of observations/measurements, $\mathbf{C} \in \mathbb{R}^{s \times p}$ is the measurement matrix, and $\mathbf{x} \in \mathbb{R}^p$ represents the full state of the system.

Since the goal is to reconstruct the full-state by using only a few sparse measurements, $s \ll p$, which means that the matrix $\mathbf{C}$ is non-invertible. Furthermore, the system is typically ill-conditioned, rendering standard linear techniques such as the pseudo-inverse ineffective. However, by assuming that the system admits a sparse representation in an alternative basis, the full-state vector $\mathbf{x}$ can be expressed as:

$$\mathbf{x} = \Phi_k \Psi_k, \quad (2.25)$$

where Φ_k denotes the truncated transformation basis of order k , and Ψ_k represents the projection of $\mathbf{x}$ onto this truncated basis.

The transformation basis is obtained by applying POD, as described in Section 2.1, to the dataset. This yields the reformulated expression of Equation 2.24 as:

$$\mathbf{y} = \Theta \mathbf{Z}_{k,SpS}, \quad (2.26)$$

where $\Theta = \mathbf{C}\Phi_k$, and Φ_k is derived via POD. The SpS algorithm employed follows the approach of [86], which incorporates a weighted least squares regression subject to physical bounds on temperature, constrained between 300 K and 3000 K. Measurement uncertainty is accounted for through the weight matrix $\mathbf{W}$, a diagonal matrix with entries $\mathbf{W}_i, i = 1/\sigma_i$, where σ_i represents the uncertainty associated with the i -th measurement in $\mathbf{y}$. For temperature measurements, the uncertainty is assumed to be 1% of the recorded value [59].

To improve the accuracy of the SpS reconstruction and maintain numerical stability, air and fuel inlet temperatures are included into the measurement vector for each case, with the same uncertainty assumption applied to these values.

Chapter 3

Experimental setup

This chapter presents the semi-industrial furnace used in this work since the experimental measurements of this system is used in this study. The main measurement techniques temperature and the energy balance of the furnace are provided.

3.1 ULB semi-industrial flameless furnace

The experimental furnace which can operate both in conventional and flameless modes utilized in this work is shown in Figure 3.1. It has a cubic combustion chamber (1100 x 1100 x 1100 mm) which is made of stainless steel. To limit the heat loss through the calculated walls is internally insulated with a 200 mm thick high temperature ($\sim 1300^{\circ}\text{C}$) ceramic foam layer, resulting in a combustion chamber with inner dimensions of 700 x 700 x 700 mm.




Fig. 3.1 Schematic of the ULB furnace (*left*), its cross section (*right top*), and the nozzle (*bottom right*). Only some probe ports are shown for the sake of clarity.

A commercial WS REKUMAT M150 recuperative Flame FLOWX burner with a nominal power of 20 kW is installed at the bottom of the combustion chamber (Figure 3.1 (*top right*)) to preheat the combustion air by extracting

the energy from the flue gases by means of its integrated metallic finned heat exchanger. In this setup, it is not possible to actively control the preheating level, being just the result of an energy balance. Moreover, due to hardware limitations, air temperature at the combustion chamber entrance can only be delivered analytically (see Section 3.1.2).

The experimental setup is designed to accommodate variations in geometry, injection system, air excess, and load, enabling the simulation of different industrial applications. As shown in Figure 3.1 (*bottom right*), the fuel is injected through a centrally located nozzle with an inner diameter of 8 mm, surrounded by the coaxial air jet whose dimensions can be adjusted (ID 16–20–25 mm). However, it is not possible to achieve flameless combustion with this configuration when the furnace is at ambient temperature. To address this issue, industrial flameless recuperative burners have a "flame mode", in which the fuel is injected close to the air jet to ensure the flame is attached to the burner during the heating of the furnace. Once the lowest furnace temperature (typically in the lower corner) exceeds the fuel-air mixture's self-ignition temperature, the burner transitions to "flameless mode," where all fuel is delivered through the coaxial injector, enabling flameless combustion.

The test bench features an air-cooling system with four cooling tubes (Figure 3.1 (*top right*)) with an outer diameter of 80 mm, extending 630 mm into the combustion chamber from the top wall. The system enables the combustion chamber to operate under different stable conditions by adjusting the airflow, allowing simulation of the impact of variable load.

There are openings for measurements in each vertical wall of the combustion chamber. One side features 110 x 450 mm quartz windows, enabling optical diagnostics. The quartz glass (GE124) has a transmissivity of about 90% in the 200–2000 nm wavelength range, allowing OH* and CH* chemiluminescence imaging. The remaining three walls are sealed with the same insulation material and contain ports for temperature measurements.

The flow control scheme and feeding systems are presented in Figure 3.2. The system consists of gas bottles, a set of Brooks Mass Flow Controllers (MFC) SLA-585XX for various flow ranges and fuels (CH₄, NH₃, H₂, CO₂, and N₂), and a static mixer to create a homogeneous fuel blend. A natural gas (NG) line

supplies the burner through a three-way valve. Combustion air and cooling air are supplied by two blowers, with their flow rates also regulated by MFCs.




Fig. 3.2 Schematic piping and instrumentation diagram of the ULB experimental test setup.

3.1.1 Temperature measurements

There are eleven thermocouple ports installed at equal intervals of 50 mm in one of the three walls of the combustion chamber. An air-cooled suction pyrometer, equipped with a 0.5 mm diameter B-type thermocouple (Platinum-Rhodium 30% / Platinum-Rhodium 6%), is used for in-flame temperature measurements. Type B thermocouples are ideal for high-temperature applications, offering high accuracy ($\pm 0.5\%$ standard accuracy) and stability at temperatures up to 1700°C . The suction probe features a Venturi tube connected to a compressed air circuit, with a maximum pressure of 6 bar(g). The thermocouple is protected by two concentric sintered alumina shields, guarding it against chemical attack and radiation heat exchange. The outer ceramic shield has an outer diameter of 10 mm, an inner diameter of 7 mm, and a length of 400 mm, while the inner ceramic

shield is a filled alumina cylinder with a diameter of 5 mm. The inox portion of the probe has an external diameter of 12 mm (Figure 3.3).




Fig. 3.3 Schematic view of the suction pyrometer.

The compressed air drives the injector, which ensures the suction of gas at high speed (approximately 100 m/s). The hot gas then heats the thermocouples through convective transfer. The high-speed suction, coupled with the two ceramic shields, ensures that the equilibrium temperature measured at the tip of the inner ceramic layer closely matches that of the hot gases, eliminating the need for corrections for radiation losses. The response time of the suction pyrometer is approximately 10 seconds. The differential voltage signal is continuously sampled at 10 Hz and automatically converted to temperature units using a National Instruments cold-junction compensated acquisition card. According to the manufacturer's specifications, the measurement uncertainty is 0.5% of the reading. However, the final temperature measurement error is estimated to be 1%, taking into account both the repeatability of the measurements and the instrument errors.


Six K-type thermocouples (Nickel-Chromium/Nickel-Alumel) are positioned flush against the insulation on another wall of the chamber, spaced 100 mm apart, to monitor the temperature of the inner wall. Additional K-type thermocouples are placed at various fixed points to measure temperatures, including at the outer window, fuel, cooling air at the inlet and outlet, combustion air, and exhaust gases after the heat exchanger.

Temperature at the top and bottom walls is measured by two shielded N-type (Nicrosil/Nisil) thermocouples which are positioned inside the combustion chamber, which are located 250 mm off-axis. The thermocouple on the bottom wall is positioned next to the combustion chamber outlet, just before the heat exchanger, and is used to measure the flue gas temperature at the outlet. The

measurement uncertainty for both K-type and N-type thermocouples is estimated at 1.5%.

3.1.2 Energy and mass balance

The mass and energy global balances must be estimated by using the available experimental measurements in order to calculate the boundary conditions for the CFD simulations. The fuels CH_4 and H_2 are used in the scope of this thesis in the ULB semi-industrial furnace. Therefore it is possible to write the global combustion reaction in a general form using the equivalent molecule $C_xH_yO_z$ as:



where x, y, z are the coefficients for the equivalent molecule, E is the air excess, f_{N_2/O_2} is the ratio of molar fraction of N_2 and O_2 in air. The number of oxygen molecules required for the complete combustion of one mole of corresponding fuel, m , can be written as:

$$m = x + \frac{y}{4} - \frac{z}{2}. \quad (3.2)$$

The coefficients x, y, z are determined as a weighted average of the number of atoms considered in each species i , using the corresponding molar fraction as the weighting factor. For example, x can be expressed as:

$$x = \sum_{i=1}^{N_s} N_{C,i} x_i, \quad (3.3)$$

where $N_{C,i}$ represents the number of carbon atoms in species i . The injected fuel mass flow rate into the chamber is determined as:

$$\dot{m}_f = \frac{P_{th}}{LHV}, \quad (3.4)$$

where P_{th} represents the desired thermal power, LHV is the lower heating value, determined using Hess' law applied to the global chemical reaction (Equa-

tion 3.1), considering water in the gaseous state (i.e. using its standard enthalpy of formation). The combustion air mass flow rate is then given by:

$$\dot{m}_{air} = \dot{m}_f(m(1 + f_{N_2/O_2})(1 + E)), \quad (3.5)$$

where $m(1 + f_{N_2/O_2})$ is the stoichiometric coefficient for air. Then, the equivalence ratio ϕ can be defined as:

$$\phi = \frac{\dot{m}_f/\dot{m}_{air}}{(\dot{m}_f/\dot{m}_{air})_{st}}. \quad (3.6)$$

The global energy balance of the furnace can be modeled as shown in Figure 3.4. Then, it would be possible to write the energy balance as follows.

$$P_{th} + P_f + P_{air} = P_{cool} + P_{walls} + P_{rad} + P_{exh} + \Delta P, \quad (3.7)$$

where P_{th} is the thermal power generated by the combustion, while P_f and P_{air} represent the sensible heat of the fuel and combustion air, respectively (with the latter typically being negligible). P_{cool} donates the power transferred to the cooling bars, P_{walls} accounts for the heat losses to the walls, P_{rad} represents the radiative power transferred through the window, and P_{exh} corresponds to the power associated with the enthalpy of the exhaust gases. In order to close the energy balance, all remaining unaccounted losses are included as ΔP . The sensible power terms (i.e. P_f , P_{air} , P_{cool} , P_{exh}) are determined using the following equation:

$$P_i = \dot{m}_i \int_{T_1}^{T_2} c_{p,i} dT, \quad (3.8)$$

where T_1 is the inlet temperature or the reference temperature ($T_{ref} = 298.15K$) of the fluid i inside the furnace, T_2 is the temperature of the fluid i at the outlet of the furnace, and $c_{p,i}$ is the specific heat capacity at constant pressure of the fluid which is approximated with a 4th degree polynomial according to the NASA's JANAF tables [70].




Fig. 3.4 Simplified sketch of the ULB furnace for the global power balance

The process efficiency (η_p) of the furnace can be written by considering the useful heat extracted by the cooling system and the radiative losses through the window as follows:

$$\eta_p = \frac{P_{cool} + P_{rad}}{P_{th}}. \quad (3.9)$$

An important consideration is that when the window is closed with insulation material, the operator must increase the cooling heat flux to maintain constant operating parameters.

In the global power balance sketch shown in Figure 3.4, the incoming air passes through the heat exchanger, where it is preheated by the exhaust gases before entering the combustion chamber. Detailed sketches of the combustion chamber and heat exchanger are provided in Figure 3.5. However, due to hardware limitations, the air temperature at the heat exchanger exit, which subsequently enters the combustion chamber, cannot be directly measured and

must instead be determined from a secondary energy balance across the heat exchanger.

$$\varepsilon_{hex} \dot{m}_{exh} \int_{T_{exh,out}}^{T_{exh,in}} c_{p,exh} dT = \dot{m}_{air} \int_{T_{air,out}}^{T_{air,in}} c_{p,air} dT, \quad (3.10)$$

where ε_{hex} is the heat exchanger efficiency, $\dot{m}_{exh}$ and $\dot{m}_{air}$ are the exhaust and cooling airflow rates, respectively. The manufacturer's data indicates that a value of 0.9 can be used for ε_{hex} .




Fig. 3.5 Separate sketches of the combustion chamber (*left*), and the heat exchanger (*right*) considered for the global power balance.

To determine the radiative power transmitted through the window (P_{rad}), the window can be modeled as illustrated in Figure 3.6. Assuming the furnace operates under steady-state conditions, the furnace interior and exterior surroundings are large, isothermal surfaces, and the insulation layers are adiabatic, and diffuse-grey surfaces with uniform radiosity. The internal and external surfaces (A_1 and A_2), along with both sides of the insulation layer (A_R), form a diffuse-grey enclosure. The surfaces A_1 and A_2 behave as black surfaces at the interior and exterior temperatures, respectively, while A_R acts as a re-radiating surface, making its emissivity irrelevant. The net radiation passing through the window, represented by surface A_2 , can then be expressed as:

$$P_{rad} = \frac{E_{b1} - E_{b2}}{\frac{1-\varepsilon_1}{\varepsilon_1 A_1} + \frac{1}{A_1 F_{12} + [(1/A_1 F_{1R}) + (1/A_2 F_{2R})]^{-1}} + \frac{1-\varepsilon_2}{\varepsilon_2 A_2}}, \quad (3.11)$$

where ε_i is the emissivity of the surface i , and F_{ij} is the view factor between surfaces i and j . Since the outer surface represent the surroundings, the emissivity ε_2 can be taken as 1. E_{bi} is the black body emissive power of surface i , and can be expressed as:

$$E_{bi} = \sigma T_i^4. \quad (3.12)$$

where $\sigma = 5.67 \times 10^{-8} \text{ W/m}^2 \text{ K}^4$ is the Stefan-Boltzmann constant.




Fig. 3.6 Idealized sketch of the window to estimate the radiative power transmitted through the window.

The view factor F_{12} can be determined by using the relation between aligned parallel rectangles [51], as explained in [11]. F_{1R} can be estimated by using the summation rule on the surface A_1 by, $F_{1R} = 1 - F_{12}$, and due to the symmetry of the closure, $F_{2R} = F_{12}$.

Finally, the convective and radiative heat loss through the furnace walls, P_{walls} can be found as follows:

$$P_{walls} = A(h_{conv} + h_{rad})(T_{walls} - T_{surr}), \quad (3.13)$$

where A is the global external furnace wall surface, T_{walls} is the external wall temperature, T_{surr} is the air temperature of in the surroundings, and h_{conv} and h_{rad} are the convective and radiative heat transfer coefficients, respectively. Additional information about the estimation of these coefficients can be found in [11].

Chapter 4

Data generation

This chapter introduces the numerical setup employed for both the semi-industrial furnace and the stagnation point reverse-flow combustor considered in this thesis. In addition, it provides a detailed description of the data generation process for each benchmark case, outlining the modeling assumptions, boundary conditions, and operating parameters that form the basis for the subsequent analyses.

4.1 ULB semi-industrial flameless furnace

4.1.1 CFD data generation

The semi-industrial furnace described in Section 3.1 was previously simulated in 3D by Ferrarotti et al. [30]. To minimize discrepancies between the 2D and 3D simulations arising from differences in the numerical setup, the 2D simulations were performed using the same configuration as the 3D cases. Reynolds-Averaged Navier–Stokes (RANS) simulations for both dimensionalities were conducted with the commercial solver Ansys Fluent 19.1. The same design parameters were applied to both 2D and 3D simulations: H_2 mole fraction ranging from 0% to 100%, equivalence ratio between 0.7 and 1, and air injector diameters of 16 mm, 20 mm, and 25 mm. The fuel blends were assumed to enter at 343 K, while the air inlet temperature and the heat losses associated with the cooling tubes and furnace walls were determined for each design condition using the furnace energy balance described in Section 3.1.2, under the assumption of an outlet temperature fixed at 1273 K.

For the numerical setup, Reynolds stresses were resolved using the standard k – ε turbulence model. Radiation was treated with the Discrete Ordinates (DO) model, while radiative properties of the reacting mixture were accounted for using the weighted-sum-of-gray-gases (WSGG) approach with the coefficients proposed by Smith et al. [95]. Turbulence–chemistry interactions were modeled through the Partially Stirred Reactor (PaSR) approach [17], with $C_{mix} = 0.5$, as previous studies demonstrated that this value provides good agreement between simulations and experiments for methane–hydrogen MILD combustion conditions [30]. As the introduction of hydrogen leads to notable changes in simulation outcomes, accounting for diffusion effects, a non-unity Lewis number assumption was applied to both HiFi and LoFi datasets. This was addressed by accounting for binary diffusion coefficients calculated following the kinetic theory and a modification of the Chapman–Enskog formula. An effective diffusion coefficient of the species in the mixture was obtained by applying a mixing rule. A sensitivity analysis performed by Aversano et al. [5] identified the Kee chemical mechanism,

consisting of 17 species and 58 reactions, as suitable for this furnace under the investigated operating conditions.

For the 3D simulations, the numerical domain was defined as a 45° angular sector of the full furnace geometry, exploiting its geometric symmetry. The validity of this assumption was confirmed by Ferrarotti et al. [30]. Similarly, for the 2D simulations, a rectangular cross-section of the cylindrical furnace representation was employed. Following the mesh sensitivity analysis reported by Ferrarotti et al. [30], numerical domains of approximately 216k cells for the 3D case and 29k cells for the 2D case were adopted.

The training dataset based on 3D furnace simulations was generated using the Latin Hypercube Sampling (LHS) technique by Aversano et al. [5], comprising 45 design points, as illustrated in Figure 4.1. An equal number of simulations (45) were performed for both the 3D and 2D cases. The datasets include temperature fields as well as major species (CH_4 , H_2 , O_2 , H_2O) and minor species (CO and OH).




Fig. 4.1 Design space composed of fuel composition, equivalence ratio, and injector diameter. Points are colored by the injector diameter value.

4.1.2 Experimental data generation

Experimental measurements were obtained from six equally spaced thermocouple ports, located 100 mm downstream, as described in Chapter 3. A schematic of the furnace, indicating the positions of these thermocouple ports and their approximate locations on a 2D projection, is shown in Figure 4.2. The only available data from both experimental campaigns are the local temperature measurements collected from these ports. For temperature measurements, the uncertainty is assumed to be 1% of the recorded value [59].




Fig. 4.2 The sketch of the ULB furnace with thermocouple ports and approximate measurement locations.

A detailed numerical and experimental investigation of the semi-industrial furnace with a 16 mm air injector diameter is provided in Ferrarotti et al. [30], where in-flame temperature profiles, pollutant emissions, and OH^* chemiluminescence images were used to validate the numerical results. The temperature measurements employed in this study were performed using an air-cooled suction pyrometer equipped with a 1.5 mm B-type thermocouple. To protect the

thermocouple from chemical degradation and radiative heat exchange with the furnace walls, it was housed within two concentric sintered aluminium shields. The numerical simulations were conducted using a RANS-based solver, with turbulence–chemistry interactions modeled using the PaSR approach. A parametric study was performed to identify the optimal constant C_{mix} , assess turbulence model sensitivity, and determine the appropriate complexity of the chemical mechanism. Aversano et al. [5] further validated the optimized numerical setup for three test cases, showing strong agreement (average error being below 6%) with measured temperature profiles at several axial locations. For more detail, the full experimental and numerical configurations can be found in Ferrarotti et al. [30], and the CFD validation against experiments is discussed extensively in Aversano et al. [5].

Additional experiments were conducted with 25 mm air-injector diameter as part of this thesis under the design conditions listed in Table 4.1. To ensure consistency with the previous 16 mm air-injector measurements, the same experimental setup was used for in-situ flame temperature measurement. Likewise, the same numerical modeling approach was employed for simulating the 25 mm cases. Before incorporating the experimental data into the training dataset for the multi-level multi-fidelity reduced order model, it is essential to verify that the numerical simulations adequately represent the experimental measurements. A comparison of measured and computed temperature profiles at multiple axial positions ($z = 100, 200, 300, 400, 500$, and 600 mm) for the test cases TC2 and TC6 are presented in Figure 4.3 and Figure 4.4. The use of a suction pyrometer and its associated spatial averaging effect—typical for this type of measurement—has been rigorously examined in Lavadera et al. [59] and accounted for in Aversano et al. [5]. To maintain consistency and enable a fair comparison, the CFD results were volume-averaged over a spherical region with a 10 mm radius, corresponding to the probe’s diameter.

TC2 and TC6 are selected as they represent the bounds of the design space in terms of H_2 mole fraction — TC2 with the lowest and TC6 with the highest. These cases highlight how the model performs under varying design parameters when compared to experimental data. As shown in Figure 4.3, for TC2, a slight overestimation is observed in the far downstream region, but the numerical

	Dia [mm]	H ₂ [mole fraction]	ϕ [-]
TC1	25	0.35	0.73
TC2	25	0.20	0.76
TC3	25	0.45	0.85
TC4	25	0.25	0.85
TC5	25	0.54	0.74
TC6	25	0.80	0.90
TC7	25	0.65	0.91
TC8	25	0.70	0.94
TC9	25	0.40	0.95
TC10	16	0.60	0.93
TC11	16	0.30	0.78
TC12	16	0.05	0.83

Table 4.1 Design parameters for the experimental data

simulation captures the temperature measurements near the injector with an average error of 11%. In TC2, the combination of low H₂ mole fraction and lean mixture results in a MILD combustion regime, characterized by a more uniform temperature distribution and smoother gradients. Considering the spatial averaging effect of the suction pyrometer — especially near the injector — the numerical and experimental data show good agreement.

In contrast, the model behaves significantly differently for TC6. As shown in Figure 4.4, the numerical simulation fails to capture the temperature profile both near the injector and further downstream. The overestimation in temperature exceeds the average measurement uncertainty of 10K. The combustion characteristics in TC6 are influenced by the higher H₂ mole fraction, which leads to sharper temperature gradients. These steep gradients complicate the comparison between the spatially averaged experimental measurements and the localized numerical predictions. The resulting average error between experimental and numerical data for TC6 is approximately 16%.

The numerical simulations for the 16 mm air-injector diameter cases demonstrated excellent agreement with experimental measurements, with an average error below 6%, as reported by Ferrarotti et al. [30]. In contrast, for the 25mm




Fig. 4.3 CFD model and experimental temperature measurements for TCA at axial positions, (a) $z=0.1$ m, (b) $z=0.2$ m, (c) $z=0.3$ m, (d) $z=0.4$ m, (e) $z=0.5$ m and (f) $z=0.6$ m.

cases, the error increased to between 11% and 16% — approximately two to three times higher. This discrepancy can be attributed to the influence of injector diameter on turbulence intensity and mixing characteristics. A smaller injector diameter leads to higher jet velocities, more intense turbulence, and finer mixing scales, all of which affect the combustion behavior and necessitate recalibration of the mixing constant C_{mix} in the PaSR combustion model.

Ferrarotti et al. [30] used a value of $C_{mix} = 0.5$, which was found to be optimal for methane-air MILD combustion in the 16 mm configuration. To maintain consistency across all simulations, this same value was retained for the 25 mm cases. For 25 mm conditions that remain within the MILD regime, the sensitivity to C_{mix} is reduced, and the predictions remain acceptable. However, in cases where elevated H_2 mole fraction increases flame speed — and where the larger injector diameter further alters mixing behavior, this fixed C_{mix} value becomes less accurate. In such scenarios, an adjusted C_{mix} is likely required to better match the experimental results.




Fig. 4.4 CFD model and experimental temperature measurements for TCB at axial positions, (a) $z=0.1$ m, (b) $z=0.2$ m, (c) $z=0.3$ m, (d) $z=0.4$ m, (e) $z=0.5$ m and (f) $z=0.6$ m.

4.1.2.1 Sparse sensing reconstruction

The sparse sensing method reconstructs the flow field using a truncated transformation matrix, making the choice of retained variance—i.e., the percentage of variance that defines the number of modes k Equation 2.25. In this thesis, LoFi data is employed for SpS, reflecting its practical relevance in real-world applications due to its low computational cost and utility in capturing the fundamental physical behavior. To determine an appropriate retained variance threshold, sensitivity analysis is performed across all training cases—three for the 16 mm and nine for the 25 mm air injector diameters—by evaluating the average normalized root mean square error of temperature predictions ($\mathcal{E}_{\text{NRMS}}^T$) against experimental measurements.

The $\mathcal{E}_{\text{NRMS}}^T$ values for the 16 mm cases are shown in Figure 4.5. As illustrated, varying the retained variance has a negligible impact on the SpS reconstruction accuracy of the temperature field compared to experimental measurements.

Therefore, the reconstruction is robust to the choice of retained variance in these cases.




Fig. 4.5 NRMSE of temperature (ε_{NRMS}^T) for each experimental data for 16mm injector diameter for different total variance values.

In contrast, the 25 mm injector cases exhibit sensitivity to the retained variance in SpS reconstruction. As shown in Figure 4.6, the reconstruction error relative to experimental measurements decreases as more variance is retained. The minimum average error of 5% is achieved when 99% of the total variance is preserved. Given that the 16 mm cases are unaffected by the retained variance, a threshold of 99% is adopted for SpS reconstruction in the remainder of this study, based on the performance observed for the 25 mm cases.

The SpS reconstruction of the temperature field alongside the CFD results is presented in Figure 4.7 for one representative case from each injector diameter. When 99% of the variance is retained, the SpS reconstruction not only minimizes the discrepancy with the experimental measurements but also demonstrates good agreement with the CFD predictions for both injector sizes.

4.2 Stagnation point reverse-flow combustor

The stagnation point reverse-flow combustor, sketched in Figure 4.8, which consists of a first, rich, and non-premixed combustion stage that is followed by a




Fig. 4.6 NRMSE of temperature (ϵ_{NRMS}^T) for exact experimental data for 25mm injector diameter for different total variance values.




Fig. 4.7 Temperature profiles of CFD results and SpS reconstruction for (a) $\text{Dia} = 16\text{mm}$, $H_2 = 0.05$, $\phi = 0.83$ (TC12), and (b) $\text{Dia} = 25\text{mm}$, $H_2 = 0.80$, $\phi = 0.90$ (TC6).

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 ($\phi_{\text{rich}}/\phi_{\text{primary}}$). 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 [90]. Numerical domain is consisted of 48k elements [78]. The MF-ROM is developed to reconstruct the

temperature and species concentrations fields (NO , NH_3 , H_2 , O_2 , OH , H_2O , NO_2 , NH_2 , and N_2).




Fig. 4.8 Schematic view of the stagnation point reverse-flow combustor investigated in this work [90]

The first training dataset is generated for 19 design conditions shown in Figure 4.9. The design parameters range between 800K–1000K for the air inlet temperature and 1.0–1.6 for the equivalence ratio of the first stage of the combustor. This training set includes 15 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 Figure 4.9). 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 Figure 4.9).

The second set of training simulations for the stagnation point reverse-flow combustor contains 84 steady-state simulations in total as shown in Figure 4.10, distributed across the design space as 76 for training and 8 for testing. The sampling in the design space was selected to follow the structure of the first training dataset shown in Figure 4.9), with 100K intervals for temperature between 600K and 1100K and 0.15 increments for the equivalence ratio between 0.7 and 1.6. However, at low temperatures and low equivalence ratios—specifically around 600–700K and an equivalence ratio of 0.7—deviations in combustion characteristics, such as the combustion location and maximum temperature profile, were observed. To address this, a denser sampling was implemented between 600K–800K for temperature and 0.7–0.85 for the equivalence ratio (*red rectangle*). Additionally, to explore the combustion limits under these design conditions, the equivalence ratio was extended down to 0.6. The upsampled regions in the




Fig. 4.9 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)

design space correspond to conditions where the numerical model predicted no combustion (*red crosses*), and thus, they were excluded from the dataset.

4.2.1 CFD data generation

The 2D RANS 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. The closure of Reynolds Stresses was achieved using the Realizable $k-\epsilon$ model with enhanced wall treatment. Radiation was modeled using the Discrete Ordinates (DO) model and the WSGG model is employed with the coefficients proposed by Smith et al. [95]. Turbulence-chemistry interactions, specifically to close the mean reactive source term, were handled with the Eddy Dissipation Concept (EDC). Otomo *et al.*'s [76] detailed kinetic mechanism was used for the ammonia oxidation. The 2D axisymmetric computational domain consists




Fig. 4.10 Design space considered in this work. The full black dots represent the training cases, while the full triangles represent the test cases with eight interpolation cases (T1-T8).

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 hours on a 20-core CPU (Intel® Xeon® processor E5-2630 v4).

4.2.2 CRN data generation

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 Figure 4.11. The CRN was previously calibrated to match the output average value of NO from the CFD simulations [90]. In this work, the calibrated CRN model was employed as the low-fidelity counterpart. CRN simulations were carried out using the NetSMOKE++ solver [98] at the same locations in parameter space as the CFD simulations. As for the computational cost, a single

CRN solution can be obtained within seconds on 1 CPU with 2.3GHz quad-core Intel® processor.




Fig. 4.11 Scheme of the Chemical Reactor Network of the combustor [90]

The Quantity of Interest (QoI) values at each grid point of the 2D flow field are determined through point-wise predictions, where each reactor's output is projected onto the corresponding CFD grid point. The grid point is selected by minimizing the discrepancy between the CFD values in that cell and the predicted values from all reactors. The CFD grid is divided into 6 - the number of reactors - and the value of each reactor is assigned to the elements within each zone. The extrapolation of CRN data on the CFD grid for each reactor is shown in Figure 4.12 through zones A-F.




Fig. 4.12 Schematic of the LoFi data generation, CFD field (*top*) and CRN extrapolation (*bottom*).

Chapter 5

Development of the multi-fidelity reduced order model

This chapter addresses two fundamental questions regarding the multi-fidelity framework: the impact of additional high-fidelity samples and their distribution within the design space on prediction accuracy. The incremental sampling methodologies introduced in Section 2.4 are applied to combustion problem using $\text{CH}_4\text{-H}_2$ as the fuel blend.

This chapter is based on the following publication:

Aysu Özden, Alberto Procacci, Riccardo Malpica Galassi, Francesco Contino, Alessandro Parente, "Incremental sampling methods for multi-fidelity surrogate modeling: Application on a furnace operating in MILD combustion conditions", Applied Thermal Engineering (2024), 122902, doi:10.1016/j.applthermaleng.2024.122902.

5.1 Methane-hydrogen mixtures in the ULB furnace

This work presents a framework for the generation of a multi-fidelity reduced order model (MF-ROM) of a combustion furnace operating in MILD conditions. Two- and three-dimensional simulations of the MILD combustion furnace with design parameters of air injector diameter, fuel composition, and equivalence ratio are fused to build the model (See Chapter 4). Moreover, a typical question related to MF-ROMs is tackled, that is, how many high-fidelity simulations should be employed to get a favourable trade-off between accuracy and training cost.

5.1.1 Dataset and workflow settings

This work presents a framework for the generation of an MF-ROM of the semi-industrial MILD combustion furnace with a nominal power of 20 kW, whose schematic is shown in Figure 3.1. Two- and three-dimensional simulations of the MILD combustion furnace with design parameters of air injector diameter, fuel composition, and equivalence ratio are fused to build the model (See Chapter 4). Moreover, a typical question related to MF-ROMs is tackled, that is, how many high-fidelity simulations should be employed to get a favourable trade-off between accuracy and training cost.

The training dataset, namely the 45 companion HiFi/LoFi simulations, needed to be split into train/test sets. Instead of making a-priori choices of the test simulations, a k -fold cross-validation is performed, given the limited amount of data samples available, which would otherwise lead to biases in the model performance evaluation. More specifically, a 30-fold cross-validation with random test sets containing 4 simulations was performed.

The total variance retained in the POD modes after dimensionality reduction was set to 99.9%, while the CoKriging model was built with radial basis functions as covariance, and a constant regression model for the scaling factor, the latter being estimated with a Bayesian optimization together with the covariance hyper-parameters.

5.1.2 Reference high-fidelity digital model

A high-fidelity digital model of the semi-industrial MILD combustion chamber was already developed by Aversano et al. [5] using 3D simulations only, and then employed by Donato et al. [23] to assimilate experimental data. The POD method was used to decompose the physics-invariant information (POD modes) from the system-dependent information (POD coefficients), while the behavior of the POD coefficients over the design space was modeled with a Kriging regression. This strategy can be seen as the single-fidelity special case of the methodology presented in this thesis, where the HiFi information contained in the POD coefficients $\mathbf{Z}$ always overrides the low-fidelity information contained in the dataset $\mathbf{Y}$, which therefore is not used.

The average NRMS error for the features of interest, computed as per Equation A.3 on field quantities in the prediction of the test simulations are provided in Figure 5.1. More specifically, the error statistics are reported (mean, standard deviation, and kernel density estimation in the form of a violin plot) over the 30 random folds. The error for temperature and the main chemical species mass fractions and pollutants, except for CO, OH, CH₄, and H₂, are predicted around and below 10%. The higher errors for CO, OH, CH₄, and H₂ are expected since these species are more challenging to predict due to their much more localized distribution [5]. The 30-fold-mean NRMS error averaged over the QoIs is $\mu_{30}[E_{\text{NRMS}}] = 0.8$, where the contribution of the less accurate species is clearly dominant.

Figure 5.2 shows the temperature and OH field comparison between the CFD and the HiFi ROM for a selected test point, which is an unseen design condition for the ROM. The field plots show that the HiFi ROM is able to predict accurately the true distribution of the features of interest without any noteworthy difference, except for a slight misplacement of the flame edge in the injector post-tip.

The full-HiFi error figures are shown in Figure 5.1 and Figure 5.2 will represent the lower error bounds, i.e., the best achievable performance of any MF-ROMs, and will be reported in black in all the following figures.




Fig. 5.1 30-fold statistics (mean, standard deviation and kernel density estimation) of the QoIs NRMS errors ϵ_{NRMS}^{QoIs} exhibited by the reference high-fidelity ROM built by using the full set of 41 HiFi simulations.




Fig. 5.2 Slice of the 3D flow field compared between the true field from the CFD simulation (*left half*) and the digital twin prediction (*right half*) in a selected test point for temperature (a) and OH mass fraction (b).

5.1.3 Multi-fidelity digital model

Two MF-ROMs were built by employing the incremental sampling strategies outlined in Section 2.4.

Both strategies required an initial HiFi training set that acts as a seed for the incremental sampling. To this end, 6 design conditions placed close to the corners of the parameter space were picked, leaving 39 simulations for the train/test sets, performing again 30-fold cross-validations with random test sets containing 4 simulations each. The initial set of six simulations was chosen because CoKriging requires a minimum number of samples, determined by both the number of fidelities and the dimensionality of the design space, to ensure a valid Cholesky decomposition of the covariance matrix. Since this decomposition is only feasible when the matrix is symmetric and positive definite, having too few simulations may lead to an ill-conditioned matrix, preventing stable decomposition and undermining model reliability. The initial set of six simulations were chosen closer to the design space boundaries and this choice was justified by the fact that Kriging-type methods are known to be inaccurate when used for extrapolation [45].

The performance of the MF-ROM utilizing the initial 6 HiFi and 41 LoFi simulations (referred to as Base MF-ROM), is compared with a SF-ROM constructed using the same 6 HiFi simulations and no LoFi simulations (referred to as Base SF-ROM), in order to show the effect of the LoFi data on predictions. This comparison is depicted in Figure 5.3 for each feature of interest within the context of a 30-fold cross-validation scenario. The full HiFi ROM performance is reported as well for reference. As illustrated in the figure, the MF-ROM generally provides more accurate predictions with reduced bias across most features compared to the SF-ROM. Notably, for species characterized by localized and challenging-to-predict behaviors such as H_2 , OH, and CO, the average performance of MF-ROM improves that of the SF-ROM. This outcome underlines the utility of employing MF-ROMs, as they demonstrate the capability to achieve better prediction accuracy compared to their SF-ROM counterparts constructed using the same amount of HiFi data.

5.1.3.1 Error-based incremental sampling method

Here, MF-ROMs were incrementally built employing the criterion in Equation 2.22, which progressively adds the HiFi sample that provides the largest accuracy




Fig. 5.3 30-fold statistics (mean, standard deviation and kernel density estimation) of the QoI-averaged NRMS error E_{NRMS} for Base MF-ROM and SF-ROM compared to the SF-ROM with 41 HiFi simulations (Full SF-ROM).

improvement among the pool of remaining HiFi simulations. The performance of the obtained models was evaluated by comparing their predictions against the test data in the 4 randomized test points, using 30-fold cross-validation. Figure 5.4 shows the probability distribution of the NRMS error E_{NRMS} defined in Equation A.4 over the 30 folds at each increment. A decreasing trend is observed for the 30-fold-mean error and 30-fold-standard-deviation with increasing HiFi samples, which reveals that (i) the addition of HiFi simulations improves the prediction accuracy of the MF-ROM and (ii) the bias induced by the test set choice (represented by the standard deviation, or PDF spread) is less and less important because the left-out simulations are less impactful on prediction accuracy. Moreover, the PDFs are skewed towards lower errors, meaning that only a few random test case combinations have less accurate predictions, typically when test cases are clusterized towards the domain boundary. Further, it is observed that the MF-ROM accuracy largest improvement happens in the first ~ 20 increments, while it plateaus after ~ 30 HiFi samples are reached. When the full set of 41 HiFi samples is added, the accuracy of the MF-ROM recovers the performance of the high-fidelity ROM shown in Section 5.1.2, Figure 5.1, as expected.




Fig. 5.4 30-fold statistics (mean, standard deviation, and kernel density estimation) of the QoI-averaged NRMS error E_{NRMS} with increasing HiFi samples employing the error-based strategy.

After a general understanding of the behavior of this sampling strategy obtained by looking into the QoIs-averaged error, it is beneficial to examine the error behavior feature by feature. Figure 5.5 shows the QoIs errors (Equation(A.3)) at 4 select stages of the incremental sampling, namely 9, 18, 27, and 41 (full) HiFi samples. These error figures show that for every feature, the addition of HiFi simulations increases the performance of the MF-ROM, as expected. Very disparate prediction performances across the QoIs are observed, coherently with the full-high-fidelity model (black symbols, also shown in Figure 5.1), with excellent accuracy for temperature and major products, and decreasing quality for reactants and radical species. It is noteworthy to remark that the MF-ROM with 9 HiFi samples already exhibits acceptable error figures compared to the full-high-fidelity ROM.

The temperature and OH predictions for the same test case are compared in Figure 5.6 as the number of HiFi simulations increases, particularly in the ranges of 9–18–27 HiFi simulations, for the error-based method. The temperature exhibits an overestimation near the symmetry axis for small and medium number of HiFi samples. A similar overestimation trend is observed in OH predictions,




Fig. 5.5 30-fold-statistics of the QoIs NRMS errors $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ exhibited by MF-ROMs built by using 9, 18, 27 and 41 HiFi simulations employing the error-based sampling strategy.

accompanied by a less accurate depiction of the flame shape. Both predictions improve with the increasing number of HiFi samples.

Lastly, topological observables were analyzed, namely the wall temperature (i.e., the average temperature predicted on the side wall $x = 0.35$ m), the peak OH value/position, and the flame length (measured as the maximum height reached by the 20% OH iso-line). 30-fold-statistics on the relative errors ε_{obs} on such quantities (Equation A.5) are shown in Figure 5.7 at the same 4 selected stages of the incremental sampling (9, 18, 27, full). The convergent trend towards the reference full-HiFi performance is here found again, except for the 30-fold-mean peak OH position whose error does not change appreciably. The error values confirm that the OH prediction in general is not very accurate, regardless of the number of low-fidelity data injected in the model, while the wall temperature (which is more relevant in a practical context) is very well predicted, even with 9 HiFi samples only.




Fig. 5.6 Comparison between the MF-ROMs for the error-based method built with 9 (left), 18 (center), 27 (right) HiFi simulations and 41 LoFi, temperature (top row), and OH mass fraction (bottom row). Each plot compares a slice of the predicted 3D field (right half) against the same (mirrored) slice of the test simulation (left half).

5.1.3.2 Uncertainty-based incremental sampling method

In this section, MF-ROMs were incrementally built employing the criterion in Equation 2.23, which adds the HiFi sample corresponding to the DoE location where the MF-ROM uncertainty is larger. Again, the performance of the obtained models were evaluated by comparing their predictions' average NRMSE distribution against the test data in the 4 randomized test points, using 30-fold cross-validation. Figure 5.8 shows the 30-fold-statistics of the NRMS error E_{NRMS} defined in Equation A.4 at each increment. The trend of the 30-




Fig. 5.7 30-fold-statistics of the relative errors ε_{obs} on wall temperature, peak OH value, peak OH position, and flame length exhibited by MF-ROMs built by using 9, 18, 27, and 41 HiFi simulations employing the error-based sampling strategy.

fold-mean of the error with an increasing number of HiFi simulations is similar to the error-based method. However, the number of increments needed to reach the plateau is lower. A sharper decrease in the error indicates a faster and steadier convergence to the full-high-fidelity NRMS error.

The effect of additional HiFi simulations on the prediction accuracy of the QoIs is shown in Figure 5.9, where only the mean over the 30-folds is shown for the sake of clarity. All the QoIs errors exhibit a smooth decreasing trend with increasing HiFi samples, maintaining the disparities across the QoIs already observed in the previous sections.

Again, a more compact visualization of predictive performance is shown in Figure 5.10, which depicts the QoIs errors (Equation A.3) at 4 selected stages of the incremental sampling, namely 9, 18, 27, and 41 (full) HiFi samples. These error figures show that for every feature, the addition of HiFi simulations increases the performance of the MF-ROM. Moreover, the accuracy of the MF-ROM with 27 HiFi simulations and the high-fidelity ROM trained with the full set of HiFi simulations show very close accuracy metrics. Figure 5.11 shows slices of the predicted flow fields of temperature and OH in the 9/18/27 HiFi stages of the




Fig. 5.8 30-fold statistics (mean, standard deviation and kernel density estimation) of the NRMS error E_{NRMS} with increasing HiFi samples employing the uncertainty-based strategy.




Fig. 5.9 30-fold-mean of the QoIs NRMS errors $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ exhibited by MF-ROMs built employing the uncertainty-based sampling strategy.

incremental sampling, compared against the test simulation in a selected test point. The temperature field progressively improves with additional HiFi simulations, both in the flame location and in the peak value. The OH field topology improves as well, reaching a satisfying flame location with 27 HiFi simulations, with the

peak OH value being over-predicted at 9 HiFi stage and under-predicted at 18 HiFi, finally reaching a more accurate value at 27.




Fig. 5.10 30-fold-statistics of the QoIs NRMS errors of the QoIs NRMS errors $\xi_{\text{NRMS}}^{\text{QoIs}}$ exhibited by MF-ROMs built by using 9, 18, 27 and 41 HiFi simulations employing the uncertainty-based sampling strategy.

This behavior is confirmed by Figure 5.12, which shows the 30-fold statistics of the topological relative errors at the same increment stages.

5.1.3.3 Performance comparison of incremental sampling methods

The performance of error-based and uncertainty-based sampling methods are compared, alongside an adaptive sampling methodology introduced by Guénot et al. [40]. The comparison is based on the NRMS errors exhibited by the MF-ROMs resulting from these three methods, as illustrated in Figure 5.13. Guénot et al.'s adaptive sampling method was utilized with non-intrusive POD-based surrogate models. The findings reveal that all three methods exhibit similar trends, with Guénot et al.'s method closely resembling the error-based approach presented in this study. Additionally, it is observed that the uncertainty-based strategy demonstrates faster convergence in prediction accuracy with increasing HiFi samples compared to the other two methods.




Fig. 5.11 Comparison between the MF-ROMs for the uncertainty method built with 9 (left), 18 (center), 27 (right) HiFi simulations and 41 LoFi, temperature (top row), and OH mass fraction (bottom row). Each plot compares a slice of the predicted 3D field (right half) against the same (mirrored) slice of the test simulation (left half).

This result may appear counter-intuitive at first glance because one would expect the error-based strategy to be the best possible strategy if the goal is to minimize the prediction error of the MF-ROM. However, it is important to consider that this sampling strategy can be classified as a *greedy* algorithm, meaning that at each increment it selects the *locally* optimal solution. In other words, this methodology does not guarantee to uncover the globally optimal set of training samples at each iteration. This would require a combinatorial search over the set of possible training samples, which becomes computationally intractable when the number of possible training samples grows larger than a few samples.




Fig. 5.12 30-fold-statistics of the relative errors ε_{obs} on wall temperature, peak OH value, peak OH position, and flame length exhibited by MF-ROMs built by using 9, 18, 27, and 41 HiFi simulations employing the uncertainty-based sampling strategy.




Fig. 5.13 Comparison between error-based and uncertainty-based incremental sampling with Guénot et al. [40]'s adaptive sampling strategies on the 30-fold-statistics of the NRMS error E_{NRMS} .

The uncertainty-based strategy can also be classified as a greedy algorithm, and thus it incurs the same local optimization problem. However, given that the uncertainty exhibited by the MF-ROM in a specific location is informed by the entire set of LoFi training samples as well, it is believed that the uncertainty-based method approximates a better global optimization method. By choosing the uncertainty-minimizing samples, this method selects HiFi samples that efficiently explore the design space based on the combined HiFi and LoFi information, and in the long run this results in more accurate models than the ones produced by the error-based strategy.

The adaptive sampling strategy of Guénot et al. [40], which is based on basis improvement, generally performs best for smooth and energetic fields where dominant modes are well represented. However, for QoIs that are highly nonlinear, strongly localized, and less energetic in the POD decomposition – such as OH and CO – uncertainty-based sampling provides better performance, particularly when initial snapshots are insufficient or when large discrepancies exist between low- and high-fidelity data, as it directly targets poorly captured or uncertain regions.

Such a result is remarkable since it demonstrates that the uncertainty-based strategy is a satisfying incremental sampling method for the problem under investigation, as well as being the only feasible one in a practical context. Moreover, it is demonstrated that an MF-ROM built with such a strategy employing ~ 20 HiFi simulations instead of 41, plus 41 LoFi simulations, delivers a model whose accuracy is comparable to the full-high-fidelity model while reducing by approximately one-half the number of degrees of freedom needed for the training.

5.2 Summary

An MF-ROM for a MILD combustion furnace was developed and the effects of sampling strategies on the prediction accuracy of this model were investigated. This framework was based on the combination of the POD, the Procrustes Manifold Alignment, and the CoKriging regression model. The ROM was used

to predict the three-dimensional flowfields of temperature and chemical species for varying design parameters, namely air injector diameter, H_2 mole fraction, and the equivalence ratio. The training datasets included 45 two-dimensional (LoFi) and 45 three-dimensional (HiFi) CFD simulations of the same furnace for the same design conditions. The framework fused the LoFi and HiFi data with the goal of reducing the number of HiFi simulations required for a desired model accuracy. Preliminary studies indicated that the number and the design conditions of the chosen HiFi data have a crucial impact on the prediction accuracy. As a result, two incremental sampling strategies were implemented and their effectiveness was examined. The error-based strategy added the HiFi sample that provides the largest accuracy improvement among the pool of remaining HiFi simulations, while the uncertainty-based one added the HiFi sample corresponding to the DoE location where the MF-ROM uncertainty was larger. The latter is more suitable as well in a practical context since it does not require the availability of the full HiFi dataset upfront.

The uncertainty-based method reached the full-high-fidelity model accuracy on the quantities of interest with ~ 20 HiFi and 41 LoFi simulations, instead of 41 HiFi simulations, therefore exhibiting a saving of approximately half the number of degrees of freedom needed for the training. If accuracy requirements on minor species were relaxed, the MF-ROM configuration with ~ 9 HiFi and 41 LoFi simulations demonstrated as well good performance, in which case the saving would be 65%.

Future developments will seek further performance improvements in the MF-ROM accuracy, especially in minor species, by an ad-hoc tuning of the CoK-riging hyper-parameters which have a strong impact on the model predictions. Moreover, a scenario in which the DoE exploration is performed on the fly during the incremental sampling will be addressed, instead of being constrained to a fixed LHS. Lastly, additional fidelity levels will be explored using less expensive chemical reactor networks and more detailed chemical kinetic mechanisms, to improve predictions on minor species.

Chapter 6

Adaptability of the MF-ROM to different type of combustion systems

This chapter focuses on extending the MF-ROM framework to different combustion systems and fuel types, thereby demonstrating its versatility and robustness. In addition, the predictive performance of the MF-ROM is systematically assessed not only for interpolation within the sampled parameter space but also for extrapolation beyond the training data, highlighting its potential for broader applicability in practical combustion problems.

This chapter is based on the following publication:

Aysu Özden, Matteo Savarese, Lorenzo Giuntini, Alberto Procacci, Riccardo Malpica Galassi, Axel Coussement, Francesco Contino, Alessandro Parente, "A multi-fidelity framework for developing digital twins of combustion systems from heterogeneous data: Application to ammonia combustion", Proceedings of the Combustion Institute (2024), 105608, doi:10.1016/j.proci.2024.105608.

6.1 Ammonia combustion in a stagnation point reverse-flow combustor

To benchmark the MF-ROM framework on stagnation point reverse-flow combustor fueled with ammonia, 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, in most cases, the new 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 framework provides a significant advantage when discovering unexplored scenarios in the design stage for a limited computational budget by incorporating simulations that are structurally different in their intrinsic nature.

6.1.1 Test case

The system investigated in this study is a stagnation-point reverse flow combustor, sketched in Figure 4.8, which consists of a first, rich, and non-premixed combustion stage that is followed by a secondary lean combustion stage. The configuration of the 2D numerical simulations and the CRN solutions can be found in Chapter 4. In this chapter 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 Figure 4.9). 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 Figure 4.9).

6.1.1.1 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 Figure 4.9. 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 $n = 5, 6, \dots, 15$ 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 (as explained in Section 5.1.3), 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 Figure 4.9 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 ($T_{air}=900\text{ K}$, $\phi_{rich}=1.30$). 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.

6.2 Results

6.2.1 Model 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 Equation 2.4, 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 ($\mathcal{E}_{NRMS}$) across all QoIs for the 2 test cases are shown in Figure 6.1 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 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 HiFi simulations, 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.




Fig. 6.1 The average NRMSE of QoIs (ϵ_{NRMS}^{QoIs}) for interpolation test data (T1 and T2) for different total variance values for case A

6.2.2 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 Figure

6.2, both for case A (Figure 6.2a) and for case B (Figure 6.2b). 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.

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 (Figure 6.2b). 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.




Fig. 6.2 The average NRMSE of QoIs ($\varepsilon_{\text{NRMS}}^{\text{QoIs}}$) (a) for case A and (b) case B for variance and error-based sampling strategies

6.2.3 Interpolation and 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 Figure 6.3 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. 6.3 The average NRMSE of QoIs ($\epsilon_{\text{NRMS}}^{\text{QoIs}}$) for test data inside the design space (T1 and T2) for each QoI

Figure 6.3 shows that the MF-ROM shows better accuracy in predicting temperature, O_2 , H_2O , and N_2 fields. These QoIs exhibit smooth, non-localized distributions, are less sensitive to turbulence–chemistry interactions, and represent energetically dominant variables, making their behavior easier to capture and reducing surrogate model prediction error.

Because the prediction error for certain variables—such as the minor species NO—can be comparatively higher, we present a direct comparison between the original CFD flow field and the MF-ROM prediction for test case T2, as shown in Figure 6.4. 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 Figure 6.5. 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 Figure 6.6 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.




Fig. 6.4 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 $T_{\text{air}} = 840$ K and $\phi_{\text{rich}} = 1.27$ (T2)

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. Figure 6.5 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 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%.




Fig. 6.5 The average NRMSE of QoIs (ϵ_{NRMS}^{QoIs}) for internal (T1 and T2) and external (T3 and T4) test data

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.




Fig. 6.6 NO field prediction of MF-ROM (*right*) with respect to the CFD simulation (*left*) for $T_{\text{air}} = 1050$ K and $\phi_{\text{rich}} = 1.65$ (T3) with 11 HiFi simulations

6.3 Summary

A multi-stage, ammonia-fuelled 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 the design range. Furthermore, the multi-fidelity approach offers 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 more effective in choosing the subsequent addition of high-fidelity samples to the training set than the error-based one. Although the error-based algorithm performs better when the boundaries of the design space are not included in the initial seed, it has been determined that it is not suitable for real-life scenarios.
- 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 method efficiently blends information from data with profoundly different levels of fidelity. The results demonstrate that this method can potentially help alleviate the computational budget needed to train reliable reduced models and digital twins, representing a key help for the design of more efficient combustion technologies. As for future work, how to build error estimates to optimally place the HiFi samples in the parameter space should be investigated to improve the accuracy of the model.

Chapter 7

Clustering based multi-fidelity reduced order model for ammonia combustion

This chapter introduces a clustering-based multi-fidelity reduced-order modeling framework aimed at improving computational efficiency while preserving high predictive accuracy. The methodology leverages unsupervised clustering to identify distinct combustion regimes within the training dataset, enabling a more structured and targeted model development. A hierarchical clustering approach is employed, beginning with global clustering to partition the design space into different combustion regimes. Subsequently, a localized clustering strategy is applied within these subsets to further refine predictions and enhance model performance.

This chapter is based on the following publication:

Aysu Özden, Riccardo Malpica Galassi, Francesco Contino, Alessandro Parente, "Clustering-based data-driven multi-fidelity reduced order modeling of ammonia combustion", Proceedings of Combustion Institute (2025), *under review*.

7.1 Regime identification: Unsupervised clustering

The novelty presented in this chapter lies in integrating the MF-ROM architecture with clustering strategies.

Clustering is a widely used technique in data analysis for grouping data based on their features. The term *unsupervised* indicates that the algorithm identifies similarities and hidden patterns without prior knowledge of predefined labels or categories. Among the various clustering methods reviewed in [91], the k -means algorithm [65], known for its simplicity and effectiveness, is utilized in this work. The data is partitioned into a user-defined number of clusters, k , each represented by a centroid. The goal is to assign each data point to the cluster that minimizes the Euclidean distance between the point and its centroid. To determine the optimal value of k , considering both cluster homogeneity and inter-cluster distances, the Davies-Bouldin index [20] is employed as a statistical evaluation metric. Since a smaller value for DBI indicates a more compact solution, the value of k is chosen based on the minimum value of the DBI. Additional information about other indices can be found in Section 7.1.2.

7.1.1 Hierarchical sampling

This chapter introduces a hierarchical clustering framework that employs a cascade of unsupervised clustering techniques to progressively refine data selection and improve model accuracy. As illustrated in Figure 7.1, the process begins with global clustering - black dotted cluster- where the entire dataset is partitioned into broad groups based on overall similarities. This initial step provides a coarse segmentation of the design space, grouping data points with comparable characteristics (green dotted clusters). Next, local clustering is applied within each global cluster, revealing finer structures and uncovering more intricate patterns and relationships (blue dotted clusters). By systematically breaking the dataset into smaller, more homogeneous subsets, the approach enables more targeted and precise model training. This cascading strategy not only enhances computational efficiency but also improves the model's ability to capture complex patterns, ultimately leading to more accurate predictions.




Fig. 7.1 Schematic of hierarchical sampling.

In this study, first clustering is applied across the entire design space to identify distinct combustion regimes and assess how this improves MF-ROM accuracy compared to the no clustering approach. Clustering is performed using temperature and NO profiles, as they are key indicators of combustion characteristics. Then, the potential for further refinement is explored by performing an additional clustering within a specific subset of the domain, to enhance predictive accuracy locally.

7.1.2 Clustering evaluation metrics

To determine the optimal value of k , considering both cluster homogeneity and inter-cluster distances, there are different metrics used commonly in the literature. Three commonly used indices in the literature for evaluating clustering solutions are the Davies-Bouldin index (DBI) [20], the Silhouette index (SI) [25], and the Calinski-Harabasz index (CHI) [13]. The DBI measures the similarity between clusters, assuming that the data density decreases as the distance from a cluster's characteristic vector increases. It can be calculated as follows:

$$DBI = \frac{1}{k} \sum_{i=1}^k \max_{j \neq i} D_{i,j}, \quad (7.1)$$

where k is the number of clusters and $D_{i,j}$ is defined as:

$$D_{i,j} = \frac{\bar{d}_i + \bar{d}_j}{d_{i,j}}. \quad (7.2)$$

Here, $\bar{d}_i$ and $\bar{d}_j$ represent the average Euclidean distances from the points to the centroids in clusters i and j , respectively, while $d_{i,j}$ denotes the distance between the centroids of clusters i and j . In this context, the DBI measures the ratio between the average within-cluster distance and the between-cluster distance. A smaller DBI value indicates a more compact clustering solution. The DBI values for global clustering with an increasing number of clusters (from 2 to 6) are shown in Figure 7.2, where 2 clusters are identified as the optimal solution according to this index.




Fig. 7.2 Davies-Bouldin index to assess evaluate the clustering.

The SI provides a measure of affinity for each cluster and can be calculated as:

$$SI = \frac{b_j - a_j}{\max(a_j, b_j)}, \quad (7.3)$$

where b_j represents the minimum average distance between the j -th observation and the points in other clusters, and a_j denotes the average distance between the j -th observation and the points within the cluster to which it has been assigned. The SI ranges between -1, indicating that the point would be better assigned to another cluster, to +1, which indicates maximum affinity within the assigned cluster. Therefore, for different values of k , the solution with the SI

closest to +1 should be preferred. As shown in Figure 7.3, 2 clusters is the optimal choice based on the SI.




Fig. 7.3 Silhouette index to assess evaluate the clustering.

Another option is the CHI, which evaluates the overall between-cluster variance V_B and the overall within-cluster variance V_W , and is defined as:

$$CHI = \frac{V_B}{V_W} \times \frac{N_p - k}{k - 1}, \quad (7.4)$$

where N_p is the number of observations. The values of V_B and V_W can be calculated as:

$$V_B = \sum_{i=1}^k N_p^{(i)} \|c_i - \mu\|^2, \quad (7.5)$$

$$V_W = \sum_{i=1}^k \sum_{x_j \in C_i} \|x_j - c_i\|^2, \quad (7.6)$$

where $N_p^{(i)}$ is the number of observations in the i -th cluster C_i , with centroid c_i . The mean of the datasets is denoted as μ . The CHI measures both the separation of clusters (V_B) and the dispersion within a cluster (V_W) with the goal of maximizing V_B and minimizing V_W . Therefore, the optimal solution according to the CHI is the one that yields the highest CHI value for a given k .




Fig. 7.4 Calinski-Harabasz index to assess evaluate the clustering.

Figure 7.4 displays the CHI values for cluster numbers ranging from 2 to 6. According to this plot, the optimal number of clusters for the dataset is determined to be 3. The distribution of simulations for the 3-cluster solution is presented in Figure 7.5.




Fig. 7.5 Design space with 3 clusters

The cluster distribution shown in Figure 7.5 is similar to the result obtained by applying DBI in a cascaded manner (see Figure 7.7 and Figure 7.12). This raises the question: rather than using DBI hierarchically, wouldn't it be more efficient to apply CHI in a single step? While the end results may be similar,

our approach focuses on identifying the inherent similarities within the dataset. By using a cascaded approach, we also address the dispersion of the clusters, ultimately achieving the optimal solution of CHI through DBI's hierarchical use. Moreover, the observed similarity arises from the naturally clustered structure of the training set. In regions with a higher density of simulations, CHI partitions the data in a manner similar to DBI; however, under conditions of uniformly spaced simulations in the design space, CHI would underperform since it does not account for the intrinsic similarities among data points. Compared to SI, DBI is easier to compute and offers an objective framework, which is why we chose DBI as the cluster evaluation metric for this work.

7.2 Case Study: Two-staged ammonia combustor

Ammonia was chosen in this study for its specific characteristics that make it a very well-suited candidate for their proposed clustering-based multi-fidelity framework. Ammonia's narrow flammability range results in distinct combustion characteristics results in distinct combustion behaviors in the investigated design space, making it suitable to local ROMs.

The stagnation point reverse-flow combustor, illustrated in Figure 4.8 is investigated in this chapter. It features a primary non-premixed combustion stage, followed by a secondary lean one. The design space is defined by the air inlet temperature (T_{air}) and equivalence ratio in the primary stage ($\phi_{primary}$) with values ranging from 600K to 1100K for temperature and 0.6 to 1.6 for the equivalence ratio as shown in Figure 4.10.

7.2.1 MF-ROM Configuration

The MF-ROM is developed to reconstruct the temperature and species concentrations fields (NO , NH_3 , H_2 , O_2 , OH , H_2O , NO_2 , NH_2 , and N_2) at the accuracy of the CFD simulations. The dataset dimensions for the first MF-ROM are $p = q = n_{cells} \times n_{QoIs} = 474800$, with $m = 76$, and $n = 5, 6, \dots, 76$ as the number of HiFi observations increases in successive iterations of the incremental sampling strategy explained in Section 2.4.

For the MF-ROMs constructed after clustering, the dimensions of p and q remain the same. However, since the maximum number of simulations is determined by the number of simulations within each cluster, the value of m and the maximum value for n vary across different MF-ROM configurations. The results are presented with an increasing number of HiFi simulations for each MF-ROM setup, starting from an initial subset (seed) of five CFD simulations, which is a prerequisite for the Co-Kriging regression model, while the number of LoFi simulations remains fixed at the value of m . It is shown that the choice of initial subset of HiFi simulations significantly influences the model performance, and as demonstrated by Özden et al. [78], the initial seed is always selected to enclose the group of LoFi simulations within each MF-ROM configuration.

The training cost and prediction efficiency of the MF-ROM depends on the amount of HiFi data included in the training set; however, the total training and prediction time remains on the order of minutes on 1 CPU with 2.3GHz quad-core Intel® processor. The computational cost of HiFi and LoFi simulations are provided in Chapter 4.2. Hence, the dominant computational cost arises from producing the HiFi training dataset, whereas the MF-ROM training and prediction steps are comparatively inexpensive once the dataset has been generated.

7.3 Results

The following sections present hierarchical clustering analyses, demonstrating the benefits of clustering at different scales.

7.3.1 Global clustering for MF-ROM improvement

Distinct flow characteristics of temperature profiles and NO emissions were observed during the database preparation phase Section 7.2. Therefore, the sampling frequency was increased in the region with an equivalence ratio of 0.6 to 0.85 and temperatures between 600K and 800K. The temperature distribution in these regions is shown in Figure 7.6a. Combustion occurs within the primary chamber, and the temperature distribution follows the expected pattern. In contrast, Figure

7.6b illustrates a case where the air inlet temperature and equivalence ratio are too low to sustain ignition within the primary chamber. However, due to mixing and turbulence effects, ignition is observed along the walls, though the combustion exhibits weaker reactivity—an expected outcome under lean conditions with lower inlet temperatures.




Fig. 7.6 Temperature profiles of two training cases with different flow structures.

Unsupervised clustering was carried out on physical parameters – the mean and maximum temperatures in the primary combustion stage, as well as the NO concentration averaged at the outlet – that could highlight these behaviors. The clustering was applied to the LoFi data, assuming its availability at every training point. This assumption aligns with a more realistic scenario where HiFi data is limited due to computational constraints. As outlined in Section 7.1.2, the Davies-Bouldin Index (DBI) was used to determine the optimal number of clusters via k-means clustering. The number of clusters was then varied parametrically from 2 to 6, with the corresponding DBI values shown in Table 7.1. The results suggest that the optimal number of clusters is two.

After determining the optimal number of clusters, the training dataset was partitioned into two groups, namely C_1 and C_2 , using k-means clustering, as illustrated in Figure 7.7. This clustering approach effectively categorized the data based on distinct flow characteristics, ensuring that each subset captured a specific

Table 7.1 DBI index values with increasing number of clusters for global clustering.

	2	3	4	5	6
DBI	0.28	0.61	0.75	0.86	0.78

combustion regime. Through unsupervised clustering, C_1 includes 59 training




Fig. 7.7 Cluster distribution after global clustering is applied.

simulations and 6 test cases, while C_2 consists of 17 training simulations and 2 test cases.

The primary motivation for clustering is to enhance MF-ROM accuracy when distinct physical phenomena are present within the training dataset. To assess this improvement, the averaged normalized root mean square error of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) across test cases is analyzed. Both Figure 7.8 and Figure 7.9 present the $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ trends as the number of HiFi simulations added to the linked dataset increases for C_1 and C_2 , respectively. The results are also compared against the accuracy of the baseline single-fidelity reduced-order model (SF-ROM) trained solely on LoFi data, since in theory, additional HiFi samples improve the accuracy of the SF-ROM.

As demonstrated in Figure 7.8 and Figure 7.9, unsupervised clustering significantly improves MF-ROM performance by increasing prediction accuracy and achieving the unclustered models's accuracy with fewer HiFi simulations, thereby reducing computational costs. However, since C_1 represents 75% of the baseline training data, the overall improvement in MF-ROM accuracy is less pronounced in C_1 compared to the more substantial enhancement observed in C_2 .

The comparison of MF-ROM and baseline SF-ROM ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) shows that dataset diversity plays a decisive role. In Figure 7.9, the unclustered MF-ROM produces errors five times larger than the baseline SF-ROM, whereas clustering consistently reduces errors below the baseline limit—highlighting the need for dataset partitioning when training data are diverse. In Figure 7.8, both clustered and unclustered MF-ROMs have higher prediction errors compared to the baseline, indicating that further clustering within C_1 is required. The larger error gap in Figure 7.9 arises from the much smaller training dataset for C_2 , which minimally affects C_1 predictions but yields substantial gains for C_2 cases when clustering is applied.




Fig. 7.8 Average NRME of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) of test cases $T3 - T8$ (minimum error is shown with the red dotted line).

The MF-ROM with no clustering achieves a minimum (shown with red dotted line) $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ of 0.24 with 42 HiFi simulations, whereas the MF-ROM built on C_1 reaches a lower $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ of 0.1 for 52 HiFi simulations. While this represents a 60% improvement in accuracy, it comes at the cost of a 24% increase in computational expense. For the clustering-based MF-ROM, the $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ for 42 HiFi simulations is 0.19, which remains 20% more accurate than the no-clustering model at the same number of HiFi simulations. In C_2 , the clustering-based model achieves stable accuracy significantly faster than the model without clustering. Specifically, only 8 HiFi simulations are required for the clustering-based model, no-clustering requires 15 simulations to reach a comparable stability. Moreover, the no-clustering MF-ROM does not match the accuracy of the clustering-based approach until 52 HiFi simulations (red dotted line in Figure 7.9), demonstrating that the clustering-based model is approximately six times more efficient in achieving its optimal accuracy.




Fig. 7.9 Average NRME of QoIs ($\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$) of test cases $T1\&T2$ (minimum error is shown with the red dotted line).

The pronounced non-monotonicity in $\mathcal{E}_{\text{NRMS}}^{\text{QoIs}}$ for the MF-ROM with no clustering shown in Figure 7.8 and Figure 7.9 are primarily due to the regression model attempting to fit the training data with different behavior. The rapid error

reduction and stabilization, predominantly observed in C_1 and C_2 , highlight the effectiveness of clustering in improving the framework's performance.




Fig. 7.10 Temperature profiles for the test case $\phi = 0.75$, $T = 625K$ for 8 HiFi simulations for (a) MF-ROM with no clustering and (b) clustering-based MF-ROM.

Since global clustering has a stronger impact on MF-ROM performance when trained on C_2 , the flow field analysis is presented for the test case $\phi = 0.75$, $T = 625K$. The temperature profile comparison in Figure 7.10 demonstrates that with just 8 HiFi simulations, the clustering-based MF-ROM closely matches the CFD results, significantly enhances the prediction accuracy compared to the no-clustering model. This flow field representation highlights the substantial improvement achieved and underscores the critical role of incorporating clustering-based inputs in the MF-ROM framework.

The NO flow field predictions for the test case T2, shown in Figure 7.11, demonstrate the significant advantage of clustering. Since T2 lies within cluster C_2 , the MF-ROM with global clustering exhibits significantly improved NO predictions compared to the unclustered model, even with only 8 HiFi simulations. The $\mathcal{E}_{\text{NRMS}}^{\text{NO}}$ without clustering is obtained as 1.12, whereas with global clustering it decreases to 0.62, indicating a 50% improvement in NO prediction accuracy through clustering. This improvement is particularly notable because

the majority of training samples differ substantially from those in C_2 , resulting in poor prediction accuracy without clustering.




Fig. 7.11 NO field predictions of MF-ROM (*right*) built with 8 HiFi simulations for (a) no clustering, and (b) global clustering with respect to the CFD simulation (*left*) for T2.

7.3.2 Localized clustering for performance refinement

We investigated the potential for further improving the accuracy of the MF-ROM developed for C_1 , as the global clustering step indicated that the training dataset required additional separation. To achieve this, an additional unsupervised clustering was performed within the training simulations of C_1 , focusing on a more localized level. The cluster count was systematically varied between 2 and 6, and the corresponding DBI values were analysed, as shown in Table 7.2. According to the DBI values, C_1 was divided into two local clusters as illustrated in Figure 7.12. Given that Co-Kriging performs better for interpolation than extrapolation, only 3 of the 6 available test points within this region were selected to evaluate the model's performance. The resulting clusters consist of 23 training simulations in C_{11} and 36 in C_{12} .

A similar trend to what was observed in C_2 is also evident for the test case within C_{11} . Since the number of simulations in C_{11} is 50% less, the model

Table 7.2 DBI index values with increasing number of clusters for local clustering.

	2	3	4	5	6
DBI	0.81	0.92	0.94	0.99	1.03




Fig. 7.12 Cluster distribution after local clustering is applied to C_1 .

achieves higher accuracy with half the number of HiFi simulations. Specifically, the accuracy reached with 10 HiFi simulations in C_{11} is only achieved with 25 simulations in C_1 (Figure 7.13a). Additionally, the minimum error achieved in C_{11} at 12 HiFi simulations is 15% lower than the minimum error observed in C_1 at 58 HiFi simulations.

For the model accuracy within C_{12} , the test case at $\phi = 1.4$, $T = 730K$ is selected as a representative example. As illustrated in Figure 7.13b, local clustering achieves higher accuracy compared to both the no-clustering and global clustering approaches, enhancing model performance by up to 45% while utilizing 30% fewer HiFi simulations. However, the improvement is less pronounced than in C_{11} due to a higher concentration of simulations within this cluster—approximately 60% of the simulations in C_1 belong to C_{11} . Despite the slower improvement with local clustering for this test case, it still demonstrates potential for reducing prediction error.

The effects of clustering on the temperature field for T5 are illustrated in Figure 7.14 using 8 HiFi simulations for each condition. Although T5 belongs to




Fig. 7.13 Average NRMSE of temperature ($\mathcal{E}_{\text{NRMSE}}^T$) of test cases for local clustering.

cluster C_1 and is therefore not expected to be significantly influenced by global clustering, the improvement from Figure 7.14a to Figure 7.14b is substantial, both in terms of matching the temperature distribution and accurately capturing the flame length. For some applications, the accuracy achieved with global clustering alone may be sufficient. Nevertheless, as shown in Figure 7.14c, applying local clustering can further reduce the over- and underestimation of temperature, yielding a more precise prediction.

Beyond temperature, NO is a critical parameter in combustion technologies that warrants investigation. Therefore, the MF-ROM performance for NO predictions under both global and local clustering is evaluated at $\phi = 1.4, T = 730K$ and the flow fields are shown in Figure 7.15. The accuracy of NO prediction for




Fig. 7.14 Temperature field predictions of MF-ROM (*right*) built with 8 HiFi simulations for (a) no clustering, (b) global clustering, and (c) local clustering with respect to the CFD simulation (*left*) for T5.

the test case improved by 25% with local clustering compared to the global clustering. This highlights the potential of local clustering in enhancing prediction accuracy, not only for temperature but also for NO.

7.4 Summary

A clustering-based MF-ROM framework for ammonia combustion, addressing challenges in multi-fidelity modeling when dealing with varying combustion regimes is presented. The application of unsupervised clustering effectively identifies different flow characteristics, leading to a more accurate and efficient MF-ROM. By leveraging data-driven clustering, the framework optimizes the distribution of training samples, reducing the reliance on HiFi simulations while maintaining or improving predictive performance.

The findings highlight two key improvements:

- Global clustering enhances MF-ROM accuracy by segmenting the training dataset based on combustion regime differences, leading to improved performance with fewer HiFi simulations.




Fig. 7.15 Average NRMSE of NO ($\epsilon_{\text{NORMS}}^{\text{NO}}$) for global (left) and local (right) clustering. Data shown are for the test case with $\phi = 1.4$, $T = 730\text{K}$ for 30 HiFi simulations.

- Localized clustering further refines model accuracy within specific subsets up to 50%, demonstrating that targeted clustering strategies can improve interpolation capabilities and computational efficiency.

This study shows that unsupervised clustering accelerates multi-fidelity reduce order model convergence and reduces computational costs while preserving predictive accuracy. Future work could explore (1) integrating MF-ROM feedback (i.e. uncertainty estimates, eigenproperties) into the clustering algorithm to leverage physical information from the training dataset, potentially improving model accuracy without increasing computational costs, (2) extending the error threshold approach—which automates clustering quality assessment and data diversity—into a fully unsupervised approach.

Chapter 8

Multi-level multi-fidelity framework for methane-hydrogen combustion: integrating CFD and experimental data

This chapter addresses the following questions: (a) Can more than two fidelity levels be effectively combined in a multi-fidelity framework for combustion applications? (b) Does the inclusion of higher-fidelity data enhance model accuracy while reducing training cost? The proposed methodology employs sparse sensing for flow field reconstruction, integrating it as an additional fidelity level. The performance of the resulting multi-level multi-fidelity reduced-order model is further evaluated under varying training constraints derived from the dataset.

8.1 Case Study: ULB semi-industrial furnace

The test case considered in this chapter is the semi-industrial furnace, with the corresponding experimental setup described in detail in Chapter 3. Both experimental and numerical data are utilized to construct the training dataset. The parameter space, defined by fuel composition, equivalence ratio, and injector diameter, is illustrated in Figure 8.1. Numerical data from 2D and 3D simulations for 45 conditions were previously reported in Chapter 4. However, due to limitations in experimental coverage, this study uses only the numerical data corresponding to air injector diameters of 16mm and 25mm. The available experimental conditions are marked as black crosses in the figure.




Fig. 8.1 Design space composed of fuel composition, equivalence ratio, and injector diameter. Points are colored by the injector diameter value. The crosses represent the condition where three-fidelity data are available.

The experimental dataset comprises 12 operating conditions—3 with a 16 mm and 9 with a 25 mm air-injector diameter. The 16 mm cases, along with their corresponding numerical simulations, were detailed in the ULB furnace characterization study by Ferrarotti et al. [30]. The additional experimental configurations for the 25 mm and 16mm injectors are listed in Table 4.1. As

reported in Chapter 4, the numerical dataset is more extensive, consisting of 27 simulations (2D and 3D) for both injector sizes, as visualized in Figure 8.1. This broader numerical coverage enables the implementation of the multi-level MF-ROM approach, enhanced by the integration of experimental data.

8.1.1 Multi-level multi-fidelity (MLMF-ROM) configuration

The MLMF-ROM developed in this study is employed to reconstruct temperature fields using experimental measurements and CFD simulations of the ULB semi-industrial furnace. Experimental measurements serve as high-fidelity (HiFi) data, 3D RANS simulations as mid-fidelity (MidFi), and 2D RANS simulations as low-fidelity (LoFi). The HiFi dataset includes 12 conditions—3 for the 16 mm air-injector and 9 for the 25 mm—while the MidFi and LoFi datasets comprise 27 cases in total. Due to the limited availability of HiFi data, only two cases (T1 and T2 from Table 4.1) are selected for testing. The corresponding design parameters are summarized in Table 8.1. These cases are chosen in line with the literature, which suggests that Kriging-type algorithms yield more accurate predictions for interpolation scenarios [45].

	Dia [mm]	H ₂ [mole fraction]	ϕ [-]
TC1	16	0.30	0.78
TC2	25	0.45	0.85

Table 8.1 Design parameters for the test cases T1 and T2

As temperature is the sole quantity of interest (QoI), the dataset includes 36 probe locations for 16mm and 42 for 25mm in the HiFi data, with $n_{\text{cells, MidFi}} = 216360$ for MidFi and $n_{\text{cells, LoFi}} = 28683$ for LoFi simulations. Following the methodology in [77], an initial subset of 6 experimental and numerical cases is used. The critical role of the HiFi seed selection has been emphasized in [77, 78]. Of the 3 available HiFi conditions for the 16mm injector, 1 is reserved for testing (Table 8.1) and the remaining 2 are included in the initial HiFi-linked seed. The remaining 4 HiFi cases are selected from the 25mm injector experimental dataset,

positioned as close as possible to the boundaries of the experimental design space, as detailed in Table 8.2.

	Dia [mm]	H ₂ [mole fraction]	ϕ [-]
TC2	25	0.20	0.76
TC5	25	0.54	0.74
TC6	25	0.80	0.90
TC9	25	0.40	0.95

Table 8.2 Design parameters for the initial seed of HiFi dataset for 25mm air-injector diameter

The initial linked training dataset includes 6 HiFi cases, each paired with corresponding MidFi and LoFi data. Out of the total 10 available HiFi cases, the remaining 4 cases are used to build the first unlinked dataset with the MidFi data. The second unlinked dataset is constructed using the remaining 19 LoFi simulations — out of 27 total — after excluding the 2 test cases and the 6 cases used in the initial linked dataset.

8.2 Results

We first examine the sensitivity of sparse sensing reconstruction to the retained variance in the temperature field. Then, we compare the performance of the MLMF-ROM against the MF-ROM trained solely on CFD data to assess the impact of incorporating additional HiFi experimental measurements on prediction accuracy. Finally, we evaluate the model’s sensitivity to dataset constraints.

8.2.1 Performance of MLMF-ROM

We evaluated the prediction accuracy of the multi-level MF-ROM, which integrates experimental measurements with 3D and 2D RANS simulations, against the MF-ROM that relies solely on 3D and 2D simulations. This comparison aims to assess the impact of incorporating additional high-fidelity data. As detailed in Section 8.1.1, the test cases TC1 and TC2 (Table 8.1) are used to evaluate both

models. The initial training seed consists of six conditions: TC10 and TC12 for the 16mm injector, and TC2, TC5, TC6, and TC9 for the 25mm injector (Table 8.2).

The incremental sampling algorithm described in Section 2.4 is employed to sequentially select the next condition to add to the linked dataset from the remaining experimental cases (TC1, TC4, TC7, TC8) for the 25mm injector—these are the only experimental data not used in the initial seed or for testing. The same selection sequence is applied to both the MLMF-ROM and MF-ROM for consistency. Once all available experimental cases have been included, the MLMF-ROM evaluation is concluded, while the MF-ROM continues by selecting additional linked datasets from the remaining 17 numerical simulations. The resulting $\mathcal{E}_{\text{NRMS}}^T$, calculated with respect to the experimental data, values for both models across all sampling iterations are presented in Figure 8.2.




Fig. 8.2 NRMSE of temperature ($\mathcal{E}_{\text{NRMS}}^T$) averaged over T1 and T2 to compare MLMF-ROM and MF-ROM prediction accuracy.

The results are obtained using a 99.9% cut-off variance for the POD in both models, as recommended in Section 2.1. Throughout the process, the MLMF-ROM consistently delivers lower prediction errors compared to the MF-ROM. Specifically, the MLMF-ROM reduces the prediction error to around 1% after incorporating just 9 experimental measurements as HiFi data. In contrast, the MF-ROM requires 18 high-fidelity 3D RANS simulations to achieve a similar

level of accuracy. Despite adding more training points, the prediction error of the MF-ROM continues to fluctuate between 1.5% and 4.5%. This underscores the stabilizing role and performance gains achieved by incorporating high-fidelity experimental data into the multi-fidelity framework. By providing ground-truth reference points, experiments anchor the model, mitigate systematic CFD bias, and enable more accurate discrepancy corrections, ultimately yielding more reliable predictions across the design space, especially in regions where CFD alone is less dependable.

The temperature flow fields for test case TC2 are also compared in Figure 8.3 to visually assess model performance. The baseline MLMF-ROM configuration, using 6 linked training points, shows reasonable agreement in temperature magnitude but does not accurately capture the flame length. When 9 linked datasets are used in the MLMF-ROM, the model demonstrates excellent agreement with both the temperature distribution and the overall flow structure compared to the experimental observations. On the other hand, although the MF-ROM with 18 linked datasets achieves its minimum error—matching the error level of the MLMF-ROM with only 9 experimental cases, as shown in Figure 8.2—the MF-ROM still fails to accurately capture both the temperature field and the flame structure for TC2.




Fig. 8.3 Temperature profiles of CFD results and SpS reconstruction for T2 for (a) MLMF-ROM with 6 experimental data, (b) MLMF-ROM with 9 experimental data, and (c) MF-ROM with 18 3D RANS simulations.

8.2.2 Impact of dataset limitations on model performance

Although the MLMF-ROM consistently shows a steady decrease in prediction error with the addition of HiFi data—reaching errors as low as 1%—the fluctuating performance of the MF-ROM raises some concerns. This instability was not previously observed for the same problem presented in Chapter 6. The primary cause of this unexpected behavior is the constraints imposed on the selection of training data, which were necessary due to the limitations of the available dataset in this study.

The configuration established for the MLMF-ROM is also applied to the MF-ROM to ensure consistency in the comparison. To better understand the unexpected behavior observed in the MF-ROM, the following configurations are systematically investigated:

- **Constrained:** The initial set of simulations is fixed to cover the boundaries of the training dataset as much as possible. The selection of the next HiFi training point is limited to the four available conditions with experimental data, after excluding the test cases and the initial linked dataset.
- **Unconstrained:** The initial set of simulations remains fixed, but the model is now free to select the next HiFi training point from all available LoFi training points, without being restricted to the four experimental cases.
- **Unconstrained+Alt.Init.:** The initial set of simulations is redefined to achieve better coverage of the design space, and the model is free to select the next HiFi training point from the entire pool of available LoFi training points.

The change in the model's behavior is illustrated in Figure 8.4 in terms of prediction error. The instability improves slightly when the model is allowed to freely select the next HiFi training point without restriction. This is expected, as the model naturally aims to select the next point that minimizes its uncertainty. When the selection is constrained to a limited subset of conditions, it is highly unlikely that this subset contains the most informative point to reduce model uncertainty, which leads to the observed instability in performance.




Fig. 8.4 NRMSE of temperature (ϵ^T_{NRMS}) averaged over T1 and T2 for different constraints applied to MF-ROM.

The most significant improvement in the model's behavior is achieved when the initial linked subset of conditions is carefully selected to ensure better coverage of the entire design space. It has been previously shown in [78] that MF-ROM is particularly sensitive to the selection of the initial set of simulations. Due to the limitations of the experimental dataset—especially for the 16mm injector cases—the initial linked dataset included conditions that did not follow the commonly recommended strategy of selecting points at the boundaries of the design space. Specifically, for the 16mm injector, only three experimental conditions were available: one was used as the test case (the one located near the center), and the other two were included in the initial linked dataset. Unfortunately, these two remaining cases were also located near the center of the H_2 -ER plane for the 16mm configuration, making them far from the optimal boundary-focused selection suggested in the literature. When the initial linked dataset was restructured to include three 16mm and three 25mm cases specifically chosen at the edges of the design space, the instability in the model's predictive behavior with additional HiFi training points was notably reduced. Nevertheless, even after removing these constraints, the prediction error of the model did not exhibit the expected monotonic decrease. This remaining issue is likely due to the limited quantity of

available training data and the single QoI used in the model, which together may not provide sufficient information for CoKriging to achieve a robust fit.

8.3 Summary

We developed a Multi-Level Multi-Fidelity Reduced Order Model (MLMF-ROM) that integrates experimental measurements and numerical simulations at multiple fidelity levels to enhance prediction accuracy for combustion applications. The framework combines Sparse sensing for reconstructing the temperature field from limited experimental measurements, Proper Orthogonal Decomposition for dimensionality reduction, Procrustes Manifold Alignment to align the different fidelity datasets, and CoKriging to predict the POD coefficients for reconstructing the full-state temperature profiles at the highest fidelity. A semi-industrial furnace capable of operating under MILD combustion conditions was used as the benchmark case. The high-fidelity dataset includes experimental measurements at 12 design conditions within a three-parameter design space, with 10 conditions used for training and 2 reserved for testing. The mid- and low-fidelity datasets consist of 27 simulations using 3D and 2D RANS models, respectively. Based on the normalized root mean square error (NRMSE) and the reconstructed temperature profiles at the test conditions, the following conclusions can be drawn:

- Incorporating additional high-fidelity data improves the stability of the model and reduces the overall prediction error.
- The MLMF-ROM achieves its minimum prediction error with approximately 50% fewer simulations in the training dataset, effectively lowering the overall training cost.
- Although MLMF-ROM and MF-ROM could reach to similar NRMSE values when averaged over the test cases, the full-state field predictions from MLMF-ROM demonstrate much better agreement with the high-fidelity data compared to those from MF-ROM.
- Constraints imposed on the MLMF-ROM, driven by the limitations of the available training dataset, negatively impact the selection of the initial

linked dataset and restrict the effectiveness of the incremental sampling algorithm in reducing model uncertainty.

- The performance of the multi-fidelity framework is highly sensitive to the selection of the initial set of simulations and to the choice of unlinked simulations used to iteratively reduce model uncertainty.

In summary, the proposed framework effectively integrates multi-level training data to enhance prediction accuracy while significantly reducing training costs in replicating the behavior of a complex combustion system. Although the results are promising, further investigation is required by incorporating additional quantities of interest (QoIs) and a broader range of experimental measurements. This would help overcome current limitations related to the design space coverage and provide more flexibility for the incremental sampling strategy to reduce model uncertainty.

Chapter 9

Conclusions and future work

The final chapter of this thesis is dedicated to discussing the initial objectives and the key results achieved. It highlights the main accomplishments of this PhD work and concludes with a discussion on future directions and potential improvements.

9.1 Conclusions

The global energy crisis is compelling society to rethink and transform how we approach energy systems. It is crucial to meet the rising energy demands while addressing pressing environmental concerns, such as climate change and air pollution. Traditional combustion technologies, which have been the dominant energy source for decades, are among the first systems that must be adapted.

This challenge is particularly acute for *hard-to-abate* sectors, such as transportation and manufacturing, where electrification is not yet a viable solution. Reducing greenhouse gas emissions in these sectors remains one of the most significant obstacles facing researchers today. Designing innovative combustion systems is a key element of the energy transition. However, developing environmentally friendly combustion technologies that can accommodate fuel flexibility while ensuring stability, efficiency, and robustness still presents considerable challenges.

Despite substantial efforts to develop new combustion solutions, the design process remains expensive and time-consuming. This is largely because achieving a thorough physical understanding and creating robust designs for complex, multi-scale, and multi-physics combustion systems require advanced tools capable of exploring and accurately predicting a wide range of operational conditions and geometries.

With the growing interest in data-driven approaches, reduced-order models (ROMs) have attracted significant attention from both the scientific and industrial communities, particularly for combustion applications. ROMs aim to approximate expensive full-order models by leveraging available numerical and/or experimental data to predict system behavior under new operating conditions. These models are typically trained using a limited number of high-fidelity (and computationally expensive) data across a broad design space, embedding the critical parameters necessary for real-time predictions or use in outer-loop applications.

However, combustion problems are inherently complex, highly non-linear, and often involve multiple design parameters, which substantially increases the need for high-fidelity training data. Acquiring such comprehensive datasets is not only extremely costly but, in many cases, practically impossible. Multi-fidelity

reduced-order models (MF-ROMs) offer a promising solution by combining a small number of expensive high-fidelity simulations with a larger set of more affordable low-fidelity simulations. This approach enables a balance between prediction accuracy and computational efficiency, making it feasible to develop reliable models within constrained budgets.

The objective of this PhD thesis was to develop a multi-fidelity framework for reduced-order models which can be useful to digital twins of combustion systems that can achieve acceptable prediction accuracy while operating within a limited computational budget.

The key activities and accomplishments of this work are summarized below:

- **Chapter 5 – Development of the multi-fidelity reduced order model:** A multi-fidelity framework was developed and integrated with various incremental sampling strategies to effectively balance prediction accuracy and training cost. The trade-off between accuracy and computational savings was evaluated for each sampling method, and the strategy focused on minimizing model uncertainty demonstrated the best performance. The framework was validated successfully in a semi-industrial furnace capable of operating under MILD conditions with a methane-hydrogen fuel blend.
- **Chapter 6 – Adaptability of the MF-ROM to different type of combustion systems:** After demonstrating the successful operation of the MF-ROM, it was essential to ensure that its performance was not system-specific. To this end, the framework was further validated on a stagnation point reverse-flow combustor operating with ammonia as the working fuel, where both interpolation and extrapolation prediction accuracy were carefully assessed, showing reliable and robust results and confirming its broader applicability.
- **Chapter 7 – Clustering-based multi-fidelity reduced order model for ammonia combustion:** The multi-fidelity framework was improved with a clustering approach to detect distinct characteristics within the training data and improve corresponding prediction accuracy. Hierarchical clustering accurately identified different combustion regimes and integrated

physical information about the training dataset, enabling the framework to consistently deliver reliable predictions within each cluster while significantly reducing training costs.

- **Chapter 8 – Multi-level multi-fidelity framework for methane-hydrogen combustion: integrating CFD and experimental data:** A multi-level multi-fidelity framework was developed by integrating experimental and numerical data. The results demonstrated that, even with limited experimental measurements, this approach could provide more accurate predictions while requiring less training data compared to a multi-fidelity framework trained only on numerical simulations. Additionally, the framework showed sensitivity to the selection of the initial linked dataset and the design conditions of the added high-fidelity data.

9.2 Future work

Multi-fidelity reduced order models present an effective solution for problems where data at varying levels of fidelity are available, even in limited quantities. The successful application of this framework to combustion problems has been demonstrated, and its capabilities have been enhanced to address several key challenges. Nevertheless, many challenges remain, and there are still numerous opportunities to further advance and refine the current model.

In light of the multi-level multi-fidelity framework presented in Chapter 8, the experimental dataset should be expanded to better cover the boundaries of the parameter space, particularly for the 16 mm injector cases. Additionally, since 2D and 3D CFD simulations for the semi-industrial furnace with a 20 mm injector diameter are already available, future experimental efforts could focus on acquiring measurements for these configurations as well.

From the numerical perspective, the optimal C_{mix} value determined for the 16 mm cases does not appear to be suitable for the 25 mm cases. A validation study should therefore be conducted to determine the appropriate C_{mix} for the 25 mm configurations by comparing numerical simulations with experimental measurements, following the same methodology used for the 16 mm cases. If

additional experimental data for the 20 mm injector becomes available, the same validation process should be applied to identify the optimal C_{mix} for these simulations, ultimately enhancing the performance of the multi-level multi-fidelity ROM.

The CoKriging regression model used in the framework is known to have hyperparameters significantly influence the prediction accuracy and require careful calibration to enhance the framework's performance. While the model has demonstrated good predictive capability for certain species and temperature profiles, there remains substantial room for improvement, particularly in the prediction of minor species and challenging targets such as OH, which exhibit highly localized distributions.

The framework developed in this thesis is fundamentally data-driven, with no explicit incorporation of physical inputs or constraints. Although the application of the unsupervised clustering algorithm implicitly captures certain physical characteristics, the baseline methodology remains detached from the physics of the problem. However, physical information or laws or constraints can be integrated into the framework at various stages. In dimensionality reduction, physics-weighted correlation matrices or modes consistent with conservation laws may be used. For manifold alignment, alignment can be guided by physical quantities such as energy distributions or temperature gradients rather than purely geometric features. The greatest flexibility lies in regression, where CoKriging can be extended with physics-informed kernels or constrained formulations that embed governing equations.

CoKriging is not the only viable regression model for multi-fidelity frameworks. Exploring alternative approaches such as multi-fidelity Polynomial Chaos Expansion (MF-PCE), multi-fidelity Support Vector Regression (MF-SVR), multi-fidelity Bayesian Optimization (MFBO), and multi-fidelity Neural Networks (MF-NN) could offer valuable insights. The choice of regression model should carefully account for factors such as the size of the available training dataset, the underlying physics of the problem, and the computational resources available.

Over the next decade, multi-fidelity reduced-order models in combustion are expected to undergo transformative advancements. By combining physics-informed modeling with expansive databases of combustors, fuels, and fidelity

levels, and integrating AI-driven techniques for adaptive sampling and prediction, these models will achieve unprecedented accuracy with minimal computational cost. This will enable autonomous exploration of complex design spaces, robust predictions across diverse combustion systems, and real-time optimization, positioning MF-ROMs as a central tool for next-generation industrial combustion design and control.

List of publications

1. **Aysu Özden**, Alberto Procacci, Riccardo Malpica Galassi, Francesco Contino, Alessandro Parente, "Incremental sampling methods for multi-fidelity surrogate modeling: Application on a furnace operating in MILD combustion conditions", *Applied Thermal Engineering* (2024), 122902, doi:10.1016/j.applthermaleng.2024.122902.
2. **Aysu Özden**, Matteo Savarese, Lorenzo Giuntini, Alberto Procacci, Riccardo Malpica Galassi, Axel Coussement, Francesco Contino, Alessandro Parente, "A multi-fidelity framework for developing digital twins of combustion systems from heterogeneous data: Application to ammonia combustion", *Proceedings of the Combustion Institute* (2024), 105608, doi:10.1016/j.proci.2024.105608.
3. **Aysu Özden**, Riccardo Malpica Galassi, Francesco Contino, Alessandro Parente, "Clustering-based data-driven multi-fidelity reduced order modeling of ammonia combustion", *Proceedings of Combustion Institute* (2025), *under review*.

References

- [1] Abdelwahid, S., Malik, M. R., Hammoud, H. A. A. K., Hernández-Pérez, F. E., Ghanem, B., and Im, H. G. (2023). Large eddy simulations of ammonia-hydrogen jet flames at elevated pressure using principal component analysis and deep neural networks. *Combustion and Flame*, 253:112781.
- [2] Adanez, J., Abad, A., Garcia-Labiano, F., Gayan, P., and De Diego, L. F. (2012). Progress in chemical-looping combustion and reforming technologies. *Progress in energy and combustion science*, 38(2):215–282.
- [3] Ankenman, B., Nelson, B. L., and Staum, J. (2010). Stochastic kriging for simulation metamodeling. *Operations research*, 58(2):371–382.
- [4] Aversano, G., Bellemans, A., Li, Z., Coussement, A., Gicquel, O., and Parente, A. (2019). Application of reduced-order models based on pca & kriging for the development of digital twins of reacting flow applications. *Computers & chemical engineering*, 121:422–441.
- [5] Aversano, G., Ferrarotti, M., and Parente, A. (2021). Digital twin of a combustion furnace operating in flameless conditions: reduced-order model development from cfd simulations. *Proceedings of the Combustion Institute*, 38(4):5373–5381.
- [6] Barricelli, B. R., Casiraghi, E., and Fogli, D. (2019). A survey on digital twin: Definitions, characteristics, applications, and design implications. *IEEE access*, 7:167653–167671.
- [7] Beér, J. M. (1994). Minimizing nox emissions from stationary combustion: reaction engineering methodology. *Chemical Engineering Science*, 49(24):4067–4083.
- [8] Bozzelli, J. W. and Dean, A. M. (1995). O+ nnh: A possible new route for nox formation in flames. *International journal of chemical kinetics*, 27(11):1097–1109.

- [9] Brunton, S. L. and Kutz, J. N. (2022). *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press.
- [10] Brunton, S. L., Noack, B. R., and Koumoutsakos, P. (2020). Machine learning for fluid mechanics. *Annual review of fluid mechanics*, 52(1):477–508.
- [11] Cafiero, M. (2022). *Formation of Nitrogen Oxides in the Combustion of H₂-rich Fuel Mixtures in the presence of Aromatics*. PhD thesis, Université libre de Bruxelles.
- [12] Cai, X., Qiu, H., Gao, L., Wei, L., and Shao, X. (2017). Adaptive radial-basis-function-based multifidelity metamodeling for expensive black-box problems. *AIAA journal*, 55(7):2424–2436.
- [13] Caliński, T. and Harabasz, J. (1974). A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1):1–27.
- [14] Cant, S. (2002). High-performance computing in computational fluid dynamics: progress and challenges. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 360(1795):1211–1225.
- [15] Cavaliere, A. and De Joannon, M. (2004). Mild combustion. *Progress in Energy and Combustion science*, 30(4):329–366.
- [16] Chen, J. H. (2011). Petascale direct numerical simulation of turbulent combustion—fundamental insights towards predictive models. *Proceedings of the Combustion Institute*, 33(1):99–123.
- [17] Chomiak, J. (1987). *Combustion: a study in theory, fact and application*.
- [18] Copernicus (2024). The 2024 global climate highlights. Available at <https://climate.copernicus.eu/sites/default/files/custom-uploads/GCH-2024/GCH2024-PDF-1.pdf>.
- [19] Czakiert, T., Krzywanski, J., Zylka, A., and Nowak, W. (2022). Chemical looping combustion: a brief overview. *Energies*, 15(4):1563.
- [20] Davies, D. L. and Bouldin, D. W. (1979). A cluster separation measure. *IEEE transactions on pattern analysis and machine intelligence*, (2):224–227.
- [21] Demo, N., Tezzele, M., and Rozza, G. (2023). A deeponet multi-fidelity approach for residual learning in reduced order modeling. *Advanced Modeling and Simulation in Engineering Sciences*, 10(1):12.

- [22] Diaz, F. and Metzler, D. (2007). Pseudo-aligned multilingual corpora. In *IJCAI*, pages 2727–2732.
- [23] Donato, L., Galletti, C., and Parente, A. (2024). Self-updating digital twin of a hydrogen-powered furnace using data assimilation. *Applied Thermal Engineering*, 236:121431.
- [24] Draper, N. (1998). *Applied regression analysis*. McGraw-Hill. Inc.
- [25] Dudek, A. (2020). Silhouette index as clustering evaluation tool. In *Classification and Data Analysis: Theory and Applications 28*, pages 19–33. Springer.
- [26] D'Alessio, G., Attili, A., Cuoci, A., Pitsch, H., and Parente, A. (2020a). Analysis of turbulent reacting jets via principal component analysis. *Data Analysis for Direct Numerical Simulations of Turbulent Combustion: From Equation-Based Analysis to Machine Learning*, pages 233–251.
- [27] D'Alessio, G., Parente, A., Stagni, A., and Cuoci, A. (2020b). Adaptive chemistry via pre-partitioning of composition space and mechanism reduction. *Combustion and Flame*, 211:68–82.
- [28] Fenimore, C. P. (1971). Formation of nitric oxide in premixed hydrocarbon flames. In *Symposium (international) on combustion*, volume 13, pages 373–380. Elsevier.
- [29] Fernández-Godino, M. G. (2023). Review of multi-fidelity models. *Advances in Computational Science and Engineering*, 1(4):351–400.
- [30] Ferrarotti, M., Fürst, M., Cresci, E., de Paepe, W., and Parente, A. (2018). Key modeling aspects in the simulation of a quasi-industrial 20 kw moderate or intense low-oxygen dilution combustion chamber. *Energy & Fuels*, 32(10):10228–10241.
- [31] Forrester, A. I., Sóbester, A., and Keane, A. J. (2007). Multi-fidelity optimization via surrogate modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2088):3251–3269.
- [32] Frassoldati, A., Cuoci, A., Faravelli, T., Ranzi, E., Candusso, C., Tolazzi, D., et al. (2009). Simplified kinetic schemes for oxy-fuel combustion. In *1st International conference on sustainable fossil fuels for future energy*, pages 6–10.
- [33] Gastpar, M., Dragotti, P. L., and Vetterli, M. (2006). The distributed karhunen–loève transform. *IEEE Transactions on Information Theory*, 52(12):5177–5196.

- [34] Glaessgen, E. and Stargel, D. (2012). 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*, page 1818.
- [35] Goel, T., Hafkta, R. T., and Shyy, W. (2009). Comparing error estimation measures for polynomial and kriging approximation of noise-free functions. *Structural and Multidisciplinary Optimization*, 38(5):429–442.
- [36] Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. (2016). *Deep learning*, volume 1. MIT press Cambridge.
- [37] Gratiet, L. L. and Garnier, J. (2014). Recursive co-kriging model for design of computer experiments with multiple levels of fidelity. *International Journal for Uncertainty Quantification*, 4(5):365–386.
- [38] Gray, J. S., Hwang, J. T., Martins, J. R. R. A., Moore, K. T., and Naylor, B. A. (2019). OpenMDAO: An Open-Source Framework for Multidisciplinary Design, Analysis, and Optimization. *Structural and Multidisciplinary Optimization*, 59:1075–1104.
- [39] Grieves, M. and Vickers, J. (2016). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In *Transdisciplinary perspectives on complex systems: New findings and approaches*, pages 85–113. Springer.
- [40] Guénot, M., Lepot, I., Sainvitu, C., Goblet, J., and Coelho, R. F. (2013). Adaptive sampling strategies for non-intrusive pod-based surrogates. *Engineering computations*, 30(4):521–547.
- [41] Guerrero, R., Ledig, C., and Rueckert, D. (2014). Manifold alignment and transfer learning for classification of alzheimer’s disease. In Wu, G., Zhang, D., and Zhou, L., editors, *Machine Learning in Medical Imaging*, pages 77–84, Cham. Springer International Publishing.
- [42] Guo, S., Silva, C. F., and Polifke, W. (2021). Robust identification of flame frequency response via multi-fidelity gaussian process approach. *Journal of Sound and Vibration*, 502:116083.
- [43] Hafeez, M. A., Procacci, A., Coussement, A., and Parente, A. (2024). Challenges and opportunities for the application of digital twins in hard-to-abate industries: A review. *Resources, Conservation and Recycling*, 209:107796.
- [44] Ham, J., Lee, D., and Saul, L. (2005). Semisupervised alignment of manifolds. In *International Workshop on Artificial Intelligence and Statistics*, pages 120–127. PMLR.

- [45] Han, Z.-H., Zimmermann, and Görtz, S. (2012). Alternative cokriging model for variable-fidelity surrogate modeling. *AIAA journal*, 50(5):1205–1210.
- [46] Hasidi, O., Abdelwahed, E. H., Qazdar, A., Boulaamail, A., Krafi, M., Benzakour, I., Bourzeix, F., Baïna, S., Baïna, K., Cherkaoui, M., et al. (2022). Digital twins-based smart monitoring and optimisation of mineral processing industry. In *International Conference on Smart Applications and Data Analysis*, pages 411–424. Springer.
- [47] Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A. (2006). Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 32:369–382.
- [48] Huang, X. and BAUMANN*, W. T. (2007). Reduced-order modeling of dynamic heat release for thermoacoustic instability prediction. *Combustion science and technology*, 179(3):617–636.
- [49] Iavarone, S. and Parente, A. (2020). No x formation in mild combustion: potential and limitations of existing approaches in cfd. *Frontiers in Mechanical Engineering*, 6:13.
- [50] Ihme, M., Chung, W. T., and Mishra, A. A. (2022). Combustion machine learning: Principles, progress and prospects. *Progress in Energy and Combustion Science*, 91:101010.
- [51] Incropera, F. P., DeWitt, D. P., Bergman, T. L., Lavine, A. S., et al. (2010). *Fundamentals of heat and mass transfer*, volume 6. Wiley, New York.
- [52] Jolliffe, I. T. (2002). *Principal Component Analysis for Special Types of Data*. Springer New York, New York, NY.
- [53] Kennedy, M. and O’Hagan, A. (2000). Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13.
- [54] Kimpton, L., Salter, J., Dodwell, T., Mohammadi, H., and Challenor, P. (2023). Cross-validation based adaptive sampling for multi-level gaussian process models.
- [55] Konnov, A., Colson, G., and De Ruyck, J. (2001). No formation rates for hydrogen combustion in stirred reactors. *Fuel*, 80(1):49–65.
- [56] Kritzinger, W., Karner, M., Traar, G., Henjes, J., and Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. *Ifac-PapersOnline*, 51(11):1016–1022.

- [57] Kumar, R. and Ezhilarasi, D. (2023). A state-of-the-art survey of model order reduction techniques for large-scale coupled dynamical systems. *International Journal of Dynamics and Control*, 11(2):900–916.
- [58] Kuya, Y., Takeda, K., Zhang, X., and Forrester, A. I. (2011). Multifidelity surrogate modeling of experimental and computational aerodynamic data sets. *AIAA journal*, 49(2):289–298.
- [59] Lavadera, M. L., Kamal, M. M., Sharma, S., Donato, L., Galletti, C., Coussement, A., and Parente, A. (2024). A combined experimental, numerical, and data consistency approach for the characterization of temperature distribution in a mild combustion furnace. *Applied thermal engineering*, 243:122625.
- [60] Le Gratiet, L. (2013). Bayesian analysis of hierarchical multifidelity codes. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):244–269.
- [61] Li, J. and Zhang, W. (2016). The performance of proper orthogonal decomposition in discontinuous flows. *Theoretical and Applied Mechanics Letters*, 6(5):236–243.
- [62] Liu, H., Ong, Y.-S., and Cai, J. (2018). A survey of adaptive sampling for global metamodeling in support of simulation-based complex engineering design. *Structural and Multidisciplinary Optimization*, 57:393–416.
- [63] Liu, X., Wan, D., and Lei, L. (2023). Multi-fidelity model and reduced-order method for comprehensive hydrodynamic performance optimization and prediction of jbc ship. *Ocean Engineering*, 267:113321.
- [64] Luo, B. and Hancock (1999). Feature matching with procrustes alignment and graph editing. In *Image Processing And Its Applications, 1999. Seventh International Conference on (Conf. Publ. No. 465)*, volume 1, pages 72–76 vol.1.
- [65] MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA.
- [66] Malik, M. R., Isaac, B. J., Coussement, A., Smith, P. J., and Parente, A. (2018). Principal component analysis coupled with nonlinear regression for chemistry reduction. *Combustion and Flame*, 187:30–41.
- [67] Malpica Galassi, R., Ciottoli, P. P., Valorani, M., and Im, H. G. (2022). Local combustion regime identification using machine learning. *Combustion Theory and Modelling*, 26(1):135–151.

- [68] Malte, P. and Pratt, D. (1975). Measurement of atomic oxygen and nitrogen oxides in jet-stirred combustion. In *Symposium (international) on combustion*, volume 15, pages 1061–1070. Elsevier.
- [69] Manohar, K., Brunton, B. W., Kutz, J. N., and Brunton, S. L. (2018). Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns. *IEEE Control Systems Magazine*, 38(3):63–86.
- [70] McBride, B. J., Gordon, S., and Reno, M. A. (1993). Coefficients for calculating thermodynamic and transport properties of individual species. Technical report.
- [71] McKay, M. D., Beckman, R. J., and Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61.
- [72] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092.
- [73] Miller, J. A. and Bowman, C. T. (1989). Mechanism and modeling of nitrogen chemistry in combustion. *Progress in energy and combustion science*, 15(4):287–338.
- [74] Morris, M. D. and Mitchell, T. J. (1995). Exploratory designs for computational experiments. *Journal of statistical planning and inference*, 43(3):381–402.
- [75] Nony, B., Rochoux, M. C., Jaravel, T., and Lucor, D. (2023). Reduced-order model for microscale atmospheric dispersion combining multi-fidelity les and rans data. In *ECCOMAS Proceedia, UNCECOMP 2023*.
- [76] Otomo, J., Koshi, M., Mitsumori, T., Iwasaki, H., and Yamada, K. (2018). Chemical kinetic modeling of ammonia oxidation with improved reaction mechanism for ammonia/air and ammonia/hydrogen/air combustion. *Int. J. Hydrog. Energy*, 43(5):3004–3014.
- [77] Özden, A., Procacci, A., Galassi, R. M., Contino, F., and Parente, A. (2024a). Incremental sampling methods for multi-fidelity surrogate modeling: Application on a furnace operating in mild combustion conditions. *Applied Thermal Engineering*, 246:122902.
- [78] Özden, A., Savarese, M., Giuntini, L., Procacci, A., Galassi, R. M., Coussement, A., Contino, F., and Parente, A. (2024b). A multi-fidelity framework

- for developing digital twins of combustion systems from heterogeneous data: Application to ammonia combustion. *Proceedings of the Combustion Institute*, 40(1-4):105608.
- [79] Parente, A. and Sutherland, J. C. (2013). Principal component analysis of turbulent combustion data: Data pre-processing and manifold sensitivity. *Combustion and flame*, 160(2):340–350.
- [80] Peherstorfer, B., Willcox, K., and Gunzburger, M. (2018). Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review*, 60(3):550–591.
- [81] Pellegrini, R., Iemma, U., Leotardi, C., Campana, E. F., and Diez, M. (2016). Multi-fidelity adaptive global metamodel of expensive computer simulations. In *2016 IEEE Congress on Evolutionary Computation (CEC)*, pages 4444–4451.
- [82] Perdikaris, P., Venturi, D., Royset, J. O., and Karniadakis, G. E. (2015). 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):20150018.
- [83] Perron, C., Rajaram, D., and Mavris, D. (2020). Development of a multi-fidelity reduced-order model based on manifold alignment. In *AIAA Aviation 2020 Forum*, page 3124.
- [84] Perron, C., Sarojini, D., Rajaram, D., Corman, J., and Mavris, D. (2022). Manifold alignment-based multi-fidelity reduced-order modeling applied to structural analysis. *Structural and Multidisciplinary Optimization*, 65(8):236.
- [85] Pinnau, R. (2008). Model reduction via proper orthogonal decomposition. In Schilders, W. H. A., van der Vorst, H. A., and Rommes, J., editors, *Model Order Reduction: Theory, Research Aspects and Applications*, pages 95–109. Springer Berlin Heidelberg, Berlin, Heidelberg.
- [86] Procacci, A., Donato, L., Amaduzzi, R., Galletti, C., Coussement, A., and Parente, A. (2024). Parameter estimation using a gaussian process regression-based reduced-order model and sparse sensing: Application to a methane/air lifted jet flame. *Flow, Turbulence and Combustion*, 112(3):879–895.
- [87] Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical science*, 4(4):409–423.

- [88] Salzmann, R. and Nussbaumer, T. (2001). Fuel staging for no x reduction in biomass combustion: experiments and modeling. *Energy & Fuels*, 15(3):575–582.
- [89] Sarker, I. H. (2021). Machine learning: Algorithms, real-world applications and research directions. *SN computer science*, 2(3):160.
- [90] Savarese, M., Giuntini, L., Malpica Galassi, R., Iavarone, S., Galletti, C., De Paepe, W., and Parente, A. (2023). Model-to-model bayesian calibration of a chemical reactor network for pollutant emission predictions of an ammonia-fuelled multistage combustor. *Int. J. Hydrog. Energy*.
- [91] Saxena, A., Prasad, M., Gupta, A., Bharill, N., Patel, O. P., Tiwari, A., Er, M. J., Ding, W., and Lin, C.-T. (2017). A review of clustering techniques and developments. *Neurocomputing*, 267:664–681.
- [92] Scheffknecht, G., Al-Makhadmeh, L., Schnell, U., and Maier, J. (2011). Oxy-fuel coal combustion—a review of the current state-of-the-art. *International Journal of Greenhouse Gas Control*, 5:S16–S35.
- [93] Serani, A., Pellegrini, R., Wackers, J., Jeanson, C.-E., Queutey, P., Visonneau, M., and Diez, M. (2019). Adaptive multi-fidelity sampling for cfd-based optimisation via radial basis function metamodels. *International Journal of Computational Fluid Dynamics*, 33(6-7):237–255.
- [94] Shafto, M., Conroy, M., Doyle, R., Glaessgen, E., Kemp, C., LeMoigne, J., and Wang, L. (2010). Draft modeling, simulation, information technology & processing roadmap. *Technology area*, 11:1–32.
- [95] Smith, T., Shen, Z., and Friedman, J. (1982). Evaluation of coefficients for the weighted sum of gray gases model. *Journal of Heat Transfer*, 104(4):602–608.
- [96] Tong, S., Ma, Y., Guo, M., Tian, Y., Song, W., Wang, H., Le, J., and Zhang, H. (2023). Optimization of aero-engine combustion chambers with the assistance of hierarchical-kriging surrogate model based on pod downscaling method. *Advances in Aerodynamics*, 5(1):20.
- [97] TRACE, C. (2024). Sectors. Available at <https://climatetrace.org/sectors>.
- [98] Trespi, S., Nicolai, H., Debiagi, P., Janicka, J., Dreizler, A., Hasse, C., and Faravelli, T. (2021). Development and application of an efficient chemical reactor network model for oxy-fuel combustion. *Energy & Fuels*, 35(9):7121–7132.

- [99] UN (2022). The sustainable development goals report 2022. Available at <https://unstats.un.org/sdgs/report/2022>.
- [100] United Nations Framework Convention on Climate Change (2016). Paris agreement. Available at <https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement>.
- [101] VanDerHorn, E. and Mahadevan, S. (2021). Digital twin: Generalization, characterization and implementation. *Decision support systems*, 145:113524.
- [102] Vitali, R., Haftka, R. T., and Sankar, B. V. (2002). Multi-fidelity design of stiffened composite panel with a crack. *Structural and Multidisciplinary Optimization*, 23(5):347–356.
- [103] Wang, C. and Mahadevan, S. (2008). Manifold alignment using procrustes analysis. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, page 1120–1127, New York, NY, USA. Association for Computing Machinery.
- [104] Wang, C. and Mahadevan, S. (2009). A general framework for manifold alignment. In *2009 AAAI Fall Symposium Series*.
- [105] Wang, C. and Mahadevan, S. (2011). Heterogeneous domain adaptation using manifold alignment. In *IJCAI proceedings-international joint conference on artificial intelligence*, volume 22, page 1541. Citeseer.
- [106] Williams, C. K. and Rasmussen, C. E. (2006). *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.
- [107] Wünnig, J. A. and Wünnig, J. G. (1997). Flameless oxidation to reduce thermal no-formation. *Progress in energy and combustion science*, 23(1):81–94.
- [108] Yamashita, H., Hayakawa, A., Okafor, E. C., Colson, S., Somarathne, K. K. A., Tsujimura, T., Ito, S., Uchida, M., Kudo, T., and Kobayashi, H. (2024). Optimum primary equivalence ratio for rich-lean two-stage combustion of non-premixed ammonia/methane/air and ammonia/hydrogen/air flames in a swirling flow. *Fuel*, 368:131598.
- [109] Yang, H., Hong, S. H., Wang, G., and Wang, Y. (2023). Multi-fidelity reduced-order model for gpu-enabled microfluidic concentration gradient design. *Engineering with Computers*, 39(4):2869–2887.

- [110] Yang, X. I., Zhang, W., Abkar, M., and Anderson, W. (2024). Computational fluid dynamics: its carbon footprint and role in carbon emission reduction. *arXiv preprint arXiv:2402.05985*.
- [111] Yin, C. and McKay, A. (2018). Introduction to modeling and simulation techniques. In *Proceedings of ISCHIA 2018 and ITCA 2018*. Leeds.
- [112] Zdybał, K., Armstrong, E., Sutherland, J. C., and Parente, A. (2022). Cost function for low-dimensional manifold topology assessment. *Scientific Reports*, 12(1):14496.
- [113] Zdybał, K., Sutherland, J. C., and Parente, A. (2023). Manifold-informed state vector subset for reduced-order modeling. *Proceedings of the Combustion Institute*, 39(4):5145–5154.
- [114] Zeldvich, Y. B. (1946). The oxidation of nitrogen in combustion and explosions. *J. Acta Physicochimica*, 21:577.
- [115] Zhang, R. and Alemazkoo, N. (2024). Multi-fidelity machine learning for uncertainty quantification and optimization. *J. Mach. Learn. Model. Comput*, 5:77–94.
- [116] Zhou, L., Song, Y., Ji, W., and Wei, H. (2022). Machine learning for combustion. *Energy and AI*, 7:100128.
- [117] Zhukov, V. P. and Kong, A. F. (2018). A compact reaction mechanism of methane oxidation at high pressures. *Progress in reaction kinetics and mechanism*, 43(1):62–78.

Appendix A

Prediction and error metrics

The multi-fidelity reduced order model in this work is obtained as a functional relation of the kind:

$$\{\mathbf{x}_{pred}, \sigma(\mathbf{x}_{pred})\}(\mathbf{d}^*) = \mathcal{M}[\mathbf{X}, \mathbf{Y}](\mathbf{d}^*), \quad (\text{A.1})$$

where the model $\mathcal{M}$ depends on the employed HiFi and LoFi training data $\mathbf{X}$ and $\mathbf{Y}$, as well as on the CoKriging hyper-parameters, and returns a predicted flowfield $\mathbf{x}_{pred}$ and its point-wise uncertainty $\sigma(\mathbf{x}_{pred})$ in an unseen design condition $\mathbf{d}^*$.

In this work, the ROM prediction accuracy is quantified using the Normalized Root Mean Square (NRMS) error, which is a common metric to assess model performance. Specifically, NRMS error is used to evaluate the accuracy of the predictions compared to test points, namely CFD simulation results for selected test cases. The NRMS error is obtained by dividing the Root Mean Square (RMS) error by the data mean:

$$\mathcal{E}_{\text{NRMS}}^{\text{QoI}}(\mathbf{d}^{\text{test}}) = \frac{1}{\mu[\mathbf{x}(\mathbf{d}^{\text{test}})]} \sqrt{\frac{\sum_{i=1}^n (x_i(\mathbf{d}^{\text{test}}) - x_{i,pred}(\mathbf{d}^{\text{test}}))^2}{n}}, \quad (\text{A.2})$$

where $x_i(\mathbf{d}^{\text{test}})$ and $x_{i,pred}(\mathbf{d}^{\text{test}})$ are the test simulation data and the ROM predicted data respectively for a given quantity of interest (QoI), e.g. temperature or a chemical species, and n is the number of cells of the flow field. More synthetic

measures of error are also employed by averaging the NRMS error over a set of N_{test} test conditions:

$$\mathcal{E}_{NRMS}^{QoIs} = \frac{\sum_t^{N_{test}} \mathcal{E}_{NRMS}^{QoIs}(\mathbf{d}^t)}{N_{test}}, \quad (\text{A.3})$$

and in turn over a set of N_{QoIs} QoIs:

$$E_{NRMS} = \frac{\sum_q^{N_{QoIs}} \mathcal{E}_{NRMS}^q}{N_{QoIs}}. \quad (\text{A.4})$$

In addition to global accuracy measures, the prediction quality is assessed by employing topological and combustion-specific observables, such as the wall temperature, the OH radical peak position and value, and the flame length. In such cases, relative errors are used:

$$\varepsilon_{\text{obs}} = \frac{\text{obs}(\mathbf{d}^{test}) - \text{obs}_{pred}(\mathbf{d}^{test})}{\text{obs}(\mathbf{d}^{test})}. \quad (\text{A.5})$$