



HAL
open science

Multi-fidelity Bayesian inference of hypersonic flow free-stream conditions and heterogeneous chemistry model parameters

Michele Capriati

► **To cite this version:**

Michele Capriati. Multi-fidelity Bayesian inference of hypersonic flow free-stream conditions and heterogeneous chemistry model parameters. Physics [physics]. Institut Polytechnique de Paris, 2024. English. NNT: 2024IPPAX003 . tel-04651406

HAL Id: tel-04651406

<https://theses.hal.science/tel-04651406>

Submitted on 17 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Multi-fidelity Bayesian inference of hypersonic flow free-stream conditions and heterogeneous chemistry model parameters

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée au CMAP, École Polytechnique et
au Département d'aéronautique et
d'aérospatiale de l'Institut von Karman de dynamique des fluides

École doctorale n°626 École Doctorale de l'IP Paris (ED IP Paris)
Spécialité de doctorat: mécanique des fluides et des solides, acoustique

Thèse présentée et soutenue à Palaiseau, le 24/01/2024, par

Michele Capriati

Composition du Jury :

Didier Lucor Directeur de Recherche, CNRS, Université Paris-Saclay	Président du jury
Pierre-Henri Maire Directeur de recherche, CEA-CESTA	Rapporteur
Taraneh Sayadi Chargé de Recherche, CNRS, Sorbonne Université	Rapporteur
Giuseppe Pascazio Professor, Politecnico di Bari	Examineur
Alejandro Alvarez Laguna Chargé de Recherche, CNRS, École Polytechnique	Examineur
Pietro M. Congedo Directeur de recherche, Inria, École Polytechnique	Co-Directeur de thèse
Thierry E. Magin Professor, École Polytechnique de Bruxelles	Co-Directeur de thèse

All' amico mio.

Abstract

The characterization of the interaction between a spacecraft entering a planetary atmosphere at hypersonic speeds and its surrounding gas is a challenging task, requiring accurate experiments and high-fidelity simulations. Numerical predictions strongly depend on modeling and experimental uncertainties, as well as numerical errors, which accumulate during the inference process. In this context, uncertainty quantification methods offer a powerful framework to account for several sources of uncertainty. High-fidelity predictions for aerospace applications combine: I) physical models sufficiently complete to account for complex flow features, leading to expensive and hard-to-perform numerical simulations, II) rigorous uncertainty quantification permitting the use of experimental data to improve the prediction of quantities of interest. Because of the large number of simulations requested to conduct uncertainty quantification studies, efficient low-fidelity representations are appealing to reduce the computational effort. However, the simplifications contained in the low-fidelity models can lead to reduced accuracy, potentially deteriorating the outcome of an inference problem.

In this thesis, our objective is to develop tools and methodologies to perform accurate predictions using high-fidelity solvers and state-of-the-art experimental data. This approach involves the use of the US3D CFD solver and an overall UQ framework to solve inference problems, which permits to include mesh error and to employ multi-fidelity strategies. The first contribution of this thesis concerns the production of high-fidelity solutions for each phenomenon of interest with the US3D solver. We coupled US3D to the open-source Mutation++ physicochemical library, which we expanded to incorporate a state-of-the-art ablation model. The second contribution concerns a study about the influence of the mesh error on the convergence of high-fidelity simulations under uncertainty. We constructed an efficient surrogate model by balancing the grid's numerical errors and the problem-related uncertainties. We applied this methodology to the propagation of model uncertainties to characterize the pressure and heat flux experienced by a re-entry vehicle. Accurate results were obtained with a coarse mesh automatically aligned to the shock for each training point. The third contribution concerns the development of a multi-fidelity formulation to alleviate the computational cost associated with the construction of the surrogate model for the high-fidelity solver and its use in an inference problem. In particular, we defined a methodology to characterize an under-expanded high-enthalpy jet obtained in the von Karman Institute Plasma-tron facility, for which no standardized rebuilding procedure for the free-stream

conditions existed to date. The description of such a flow required axisymmetric simulations. The analysis allowed us to characterize the flow conditions at the entrance of the nozzle and the nitrogen catalytic recombination coefficient of the probe used to measure the heat flux and pressure at the stagnation point. The characterized uncertainties were then propagated through the numerical solver yielding an uncertainty-based high-fidelity representation of the supersonic flow structure variability. In the last application, we devised a methodology for the calibration and assessment of a finite-rate chemistry gas-surface interaction model for ablation. Specifically, we inferred the rate coefficients of the elementary reactions occurring between a carbon surface and a nitrogen gas from both molecular beam-surface scattering and Plasmatron experiments. The analysis highlighted that both experimental data sets are compatible with the same calibrated model.

In conclusion, we proposed powerful stochastic tools, encompassing one or more fidelity levels, to infer hypersonic flow free-stream conditions and heterogeneous chemical model parameters.

Résumé

La caractérisation de l'interaction entre un vaisseau spatial entrant dans une atmosphère planétaire et le gaz environnant est une tâche difficile, nécessitant des expériences précises et des simulations haute-fidélité. Les prédictions numériques dépendent fortement des incertitudes de modélisation et expérimentales, ainsi que des erreurs numériques, qui s'accumulent pendant le processus d'inférence. Les méthodes de quantification de l'incertitude offrent une structure puissante pour tenir compte de plusieurs sources d'incertitude. Les prédictions haute-fidélité pour les applications aérospatiales combinent des modèles physiques suffisamment complets pour prendre en compte des caractéristiques d'écoulement complexes, et une quantification rigoureuse de l'incertitude. En raison du grand nombre de simulations requises pour mener des études de quantification de l'incertitude, des représentations efficaces à basse fidélité sont attrayantes pour réduire l'effort de calcul. Cependant, les simplifications contenues dans les modèles à basse fidélité peuvent conduire à une précision réduite, détériorant potentiellement le résultat d'un problème d'inférence.

Dans cette thèse, notre objectif est de développer des outils et des méthodologies pour effectuer des prédictions précises à l'aide de solveurs haute-fidélité et de données expérimentales de pointe. Cette approche implique l'utilisation du solveur CFD US3D et d'une structure globale UQ pour résoudre des problèmes d'inférence, ce qui permet d'inclure l'erreur de maillage et d'utiliser des stratégies multi-fidélité. La première contribution de cette thèse concerne la production de solutions haute-fidélité avec le solveur US3D pour chaque phénomène d'intérêt. Ce dernier a été couplé à la librairie Mutation++, que nous avons étendue pour incorporer un modèle innovant d'ablation. La deuxième contribution concerne une étude sur l'influence de l'erreur de maillage sur la convergence des simulations haute-fidélité sous incertitude. Nous avons construit un modèle de substitution efficace en équilibrant les erreurs numériques associées au maillage et les incertitudes liées au problème. Nous avons appliqué cette méthodologie à la propagation des incertitudes du modèle pour caractériser la pression et le flux de chaleur subis par un véhicule de rentrée. Des résultats précis ont été obtenus avec un maillage grossier automatiquement aligné avec le choc pour chaque point d'entraînement. La troisième contribution concerne le développement d'une formulation multi-fidélité pour alléger le coût de calcul associé à la construction du modèle de substitution du solveur haute-fidélité et à son utilisation dans un problème d'inférence. En particulier, nous avons défini une méthodologie pour caractériser un jet hypersonique sous-détendu obtenu dans le

moyen d'essai Plasmatron de l'Institut von Karman, pour laquelle aucune procédure standard de reconstruction des conditions en amont n'existait à ce jour. L'analyse nous a permis de caractériser les conditions d'écoulement à l'entrée de la tuyère et le coefficient de recombinaison catalytique de l'azote de la sonde utilisée pour mesurer le flux de chaleur et la pression au point d'arrêt. Les incertitudes caractérisées ont ensuite été propagées à travers le solveur numérique, fournissant une représentation haute-fidélité basée sur l'incertitude de la variabilité de la structure de l'écoulement supersonique. Dans la dernière application, nous avons élaboré une méthodologie pour l'étalonnage et l'évaluation d'un modèle chimique d'interaction gaz-surface à taux de réaction finis pour l'ablation. Plus précisément, nous avons déduit les taux de coefficients des réactions élémentaires se produisant entre une surface de carbone et un gaz d'azote à partir d'expériences de faisceau moléculaire-surface de diffusion, ainsi que Plasmatron. L'analyse a montré que les deux ensembles de données expérimentales sont compatibles avec le même modèle étalonné.

En conclusion, nous avons proposé des outils stochastiques puissants, englobant un ou plusieurs niveaux de fidélité, pour déduire des conditions d'écoulement libre hypersonique et de paramètres de modèles chimiques hétérogènes.

Acknowledgments

As I look back at the path that brought me here, a famous quote, generally attributed to J. A. Shedd, comes to my mind:

'A ship in harbour is safe, but that is not what ships are built for.'

This quote has a dual meaning to me.

Firstly, it relates to the beginning of this journey: I left Bari, my *safe harbour*, to venture to Brussels for my master's thesis. Back then, I did not imagine how pivotal such a decision was going to be. Yet, six years later, here I am, writing the last piece of my thesis!

I am very thankful to all the people that I met during this journey. A special thanks go to Prof. Pascazio for his inspiring fluid dynamic classes and for encouraging me to participate in the STP program at the von Karman Institute. It was there that I met Thierry, who introduced me to the hypersonic world and has since supervised my research activities. Thank you for the opportunity to embark on this journey and for guiding me in developing my expertise. During my stay, I also met a fantastic group of people: Andrea A., Andrea C., Andrea F., Carmelo, Chiara, David, Domenico, Francisco, Lorenzo, Maria F., Maria R., Onur, Piero, Riccardo, Ruben, Sahadeo, Simon, Stefano, Utkan, Vincent, and Zoe. I am extremely grateful for all the *chill* time we shared together. The following year, I decided to stay and participate in the RM program. During this, I had the fortune of meeting many amazing friends. Among others, Dániel, Diana, Federica, Giacomo, Ludovico, Luca, Marco, Matilde, Pedro, Riccardo, and Salvatore, thank you for making so enjoyable such an intense year! During the RM, I had the opportunity to take an exam with Pietro, who, some months later, proposed to me to start a Ph.D. at VKI in collaboration with the INRIA team. Although the pandemic disturbed our initial plans, preventing me from spending much time in Paris, it did not prevent the collaboration from going smoothly. In fact, my only regret is that my French is still at the '*Ça va?*' level. I am very thankful to you for the opportunity to pursue my Ph.D. and for your guidance into the world of UQ! I also want to thank all the advisors who mentored me during the STP and the RM. J. B., George, Bruno, Bernd, and Alessandro, thank you for all your suggestions and for sharing with me your knowledge.

Secondly, in the collective imagination, ships are often associated with the desire to depart from harbours to explore unknown shores. It is precisely this love for *discovery* that moves every Ph.D., and I would like to thank all the people I met, who fueled my curiosity. I want to express my gratitude to Prof. Candler, Prof. Schwartzenuber, Prof. Minton, and their teams for all the scientific discussions we had and for hosting me at the University of Minnesota, where I had such an intellectually stimulating and *freezing* time. I would also like to thank all the friends, colleagues, and professors, at VKI and outside, with whom I had the pleasure to interact and collaborate: Alessandro, Anabel, Angela, Antonio, Ata, Chiara, Christos, David, Davide, Elena, Enrico, Federico, Giuseppe, Gonçalo, Guillaume, İpek, Joseph, Jure, Khalil, Krishna, Lilla, Maria Teresa R., Maria Teresa S., Meryem, Miguel, Olivia, Olivier C., Olivier L.M., Pierre, Pietro, Rebecca, Sanae, Sander, Thanos, Tommaso, Will, Zachary, and Zuheyr. The love for discovery is second only to the joy of sharing it, and I am extremely grateful to the students I collaborated with and had the pleasure to mentor: Alberto, Csaba, Dairo, Gregoire, Greyson, Matthias, Stefano, and Vittorio. Best of luck for your future achievements!

I express also my sincere gratitude to the Ph.D. jury, Dr. Pierre-Henri Maire, Dr. Taraneh Sayadi, Prof. Giuseppe Pascazio, Dr. Didier Lucor, Dr. Alejandro Alvarez Laguna, for accepting to review my Ph.D. work and for providing me valuable insights and comments.

A big *'thank you'* goes to my family for the support they constantly provide me, no matter *what* and *where*. Thanks to my parents and brother, for being my *safe harbour* in every storm, and for always encouraging me in every single step of my life. Thanks to my grandparents, for the timeless lessons you taught me, among which *'Vogliamoci bene!'*. Thanks to my aunts and uncles, as, being first-family-child, I felt a little like everybody's son. Thanks to my cousins, with which I grew up, making tons of mistakes, and one thing right. Thanks to my Colombian family, for all the *positive energy* you transmit me. Thanks to my friends in Belgium, for being my family in Brussels. Thanks to my lifelong friends - Ambrogio, Anna, Domenico M., Domenico P., Giuseppe D. F. , Giuseppe D., and Michele - as I sincerely do not understand why you are still friends of mine, but I am very happy of it.

Dulcis in fundo. The end of this challenging journey was not easy, not for an overthinker like me. I got anxious and stressed. I could have lost the direction of the navigation. But I did not. I did not also because of the continuous support of my positive-thinking wife, who, day after day, put back the helm in the correct position. Hence, my biggest acknowledgments go to you, *mi señora*. Thank you for backing me up in the worst moments. Thank you for sharing with me the best ones. Thank you for expecting the best of me. And, thank you for being the best of me.

Brussels, February 2024
Michele Capriati

Table of Contents

Abstract	i
Résumé	iii
Acknowledgments	v
Publications	i
1 Introduction	1
1.1 Atmospheric entry flows	2
1.2 Flow and material characterization in hypersonic applications . . .	3
1.3 Objectives and outline	10
2 Physical modeling	17
2.1 Overview	18
2.2 Governing equations	20
2.2.1 Homogeneous finite-rate chemistry	21
2.2.2 Multi-component thermodynamics	24
2.2.3 Multi-component transport	30
2.2.4 Gas-surface interaction	33
2.3 Gas surface interaction boundary condition in MUTATION ⁺⁺ . . .	40
2.3.1 MUTATION ⁺⁺ description	40
2.3.2 Nitridation model implementation and verification	41
2.4 Experiments in ground testing facilities	42
2.4.1 VKI Plasmatron	43
2.4.2 Molecular beam-surface scattering apparatus	47
2.5 Summary	49
3 Uncertainty quantification methods	51
3.1 Overview	51
3.2 Forward propagation	53
3.3 Stochastic inverse problem	56
3.3.1 MCMC algorithms	57
3.3.2 Chain convergence	64
3.4 Surrogate models	66
3.4.1 Kriging model	67

3.4.2	Hierarchical Kriging	71
3.4.3	Adaptive sampling	73
3.5	UQLab	78
3.6	Summary	78
4	Multi-fidelity and multi-level numerical simulations	81
4.1	Overview	81
4.2	CFD Solvers	83
4.2.1	High-fidelity: US3D	85
4.2.2	Medium-fidelity: stagnation-line	86
4.2.3	Low-fidelity: supersonic nozzle 0D model	87
4.3	Mesh numerical error	88
4.4	US3D-Mutation ⁺⁺ coupling	90
4.4.1	Description	90
4.4.2	Verification test-cases	92
4.5	Balancing numerical and epistemic uncertainties in an atmospheric entry flow	95
4.5.1	EXPERT atmospheric entry	97
4.5.2	Numerical uncertainty on deterministic simulations	97
4.5.3	Surrogate model construction and forward propagation	102
4.6	Summary	107
5	Holistic characterization of an under-expanded high-enthalpy jet under uncertainty	109
5.1	Overview	110
5.2	Deterministic simulations	111
5.2.1	Flow structure	113
5.2.2	Numerical uncertainty	115
5.3	Uncertainty-based results	117
5.3.1	Surrogate model construction and forward propagation	119
5.3.2	Characterization of the flow conditions and of the material catalytic response	123
5.3.3	Flow structure variability	124
5.4	Summary	127
6	Bayesian calibration of a finite-rate nitridation model from molecular beam and plasma wind tunnel experiments	129
6.1	Overview	130
6.2	Deterministic simulations of a Plasmatron experimental campaign	132
6.3	Stochastic inverse problem definition	134
6.4	Results of the calibration	139
6.4.1	Posterior marginal distributions	139
6.4.2	Model posterior predictions	143
6.5	Summary	149

7	Conclusions	153
A	US3D-Mutation⁺⁺ Coupling Verification Test Cases	159
A.1	0D Reactor	159
A.2	1D Diffusion Problem	160
A.3	1D Catalytic Diffusion Problem	161
A.4	Nitrogen flow over an ablative surface	161
B	Transport matrices	165

Publications

Journals articles

1. **M. Capriati**, A. Cortesi, T. Magin, and P. Congedo, Stagnation point heat flux characterization under numerical error and boundary conditions uncertainty. *European Journal of Mechanics / B Fluids*, 2022.
2. A. Baskaya, **M. Capriati**, A. Turchi, T. Magin, and S. Hickel, Assessment of immersed boundary methods for hypersonic flows with gas-surface interactions. *Computers & Fluids*, 2024.
3. **M. Capriati**, A. Turchi, P. Congedo, and T. Magin, Holistic characterisation of an under-expanded high-enthalpy jet under uncertainty. Under submission.
4. **M. Capriati**, A. del Val, T. Schwartzenruber, T. Minton, P. Congedo, and T. Magin, Stochastic calibration of a detailed nitridation model from molecular beam and plasma wind tunnel experiments. Under submission.

Conference proceedings

1. **M. Capriati**, A. del Val, T. Schwartzenruber, T. Minton, P. Congedo, and T. Magin, Bayesian calibration of a finite-rate nitridation model from molecular beam and plasma wind tunnel experiments, *EUCASS, 2023*, Lausanne, Switzerland.
2. **M. Capriati**, G. Bellas-Chatzigeorgis, A. Turchi, B. Helber, and T. Magin, Thermal non-equilibrium modeling for ablative gas-surface interaction, *HiSST, 2022*, Bruges, Belgium.
3. **M. Capriati**, A. Turchi, P. Cogedo, and T. Magin, Heat flux characterization of an under-expanded/supersonic plasma jet over a catalytic probe, *EUCASS, 2022*, Lille, France.
4. M. Geratz, **M. Capriati**, G. Grossir, and T. Magin, Influence of Physical Models on the Numerical Modeling of Hypersonic Nozzle Flow Expansion, *EUCASS, 2022*, Lille, France.

5. A. Baskaya, **M. Capriati**, D. Ninni, F. Bonelli, G. Pascazio, A. Turchi, T. Magin, and S. Hickel., Verification and Validation of Immersed Boundary Solvers for Hypersonic Flows with Gas-Surface Interactions, AIAA AVIATION, 2022, Chicago, IL & Virtual.
6. C. Garbacz, F. Morgado, M. Fossati, J. Scoggins, T. Magin, and **M. Capriati** Influence of Thermochemical Modelling of CO₂-N₂ Mixtures on the Shock Interaction Patterns at Hypersonic Regimes, AIAA AVIATION, 2021, Virtual.
7. **M. Capriati**, K. Prata, T. Schwartzentruber, G. Candler, and T. Magin, Development of a nitridation gas-surface boundary condition for high-fidelity hypersonic simulations, ECCOMAS, 2021, Virtual.

Book Chapter

1. **M. Capriati**, T. Magin, and P. Congedo, Determination of a stochastic finite-rate nitridation model consistent with molecular beam and plasma wind tunnel data. Review of the VKI Doctoral Research 2022-2023, Rhode-St-Genève, Belgium
2. **M. Capriati**, T. Magin, and P. Congedo, Multi-fidelity heat flux characterization of a supersonic plasma jet over a catalytic probe. Review of the VKI Doctoral Research 2021-2022, Rhode-St-Genève, Belgium
3. **M. Capriati**, T. Magin, and P. Congedo, Heat flux characterization of an under-expanded/supersonic plasma jet over a catalytic probe. Review of the VKI Doctoral Research 2020-2021, Rhode-St-Genève, Belgium

List of Figures

1.1	SpaceX's Dragon re-entering atmosphere. Credit: SpaceX.	1
1.2	Sketch of an entry hypersonic flow.	2
1.3	Sketch of a stochastic inference process.	6
1.4	Summary of the UQ works related to atmospheric entry flows reviewed in Section 1.2.	11
2.1	Sketch of the surface mass balance.	33
2.2	Sketch of the surface energy balance.	34
2.3	Sketch of atomic adsorption/desorption.	36
2.4	Sketch of ER reaction.	37
2.5	Sketch of LH reaction.	37
2.6	ACA model predictions of the efficiency of the nitridation and the recombination reactions at low pressure (solid lines) and high pressure (dashed lines) versus the experimental values. Adapted from Prata et al. [1].	39
2.7	Sketch of the surface balance solution.	41
2.8	Verification of the implementation of the reactions in Table 2.1 into Mutation ⁺⁺	42
2.9	Sketch of the Plasmatron facility, adapted from Fagnani [2].	44
2.10	Ablative carbon sample in a nitrogen flow in the left figure. Spectrometer output matrix, in the right one, with spectral emission on the horizontal axis, and radial position on the vertical axis. The spectrometer slit position is indicated with dashed red line in the left figure. Figure courtesy of Helber et al. [3].	46
2.11	Under-expanded air jet over the catalytic probe obtained in the VKI's Plasmatron facility.	47
3.1	Statistical convergence history for adaptive sampling: mean (solid line) plus and minus one standard deviation (shadow area).	76
3.2	Statistical convergence history for adaptive sampling: mean (solid line) plus and minus one standard deviation (shadow area).	77
4.1	Sketch of an under-expanded jet over a probe.	87

4.2	Comparison between the high-fidelity solution extracted along the stagnation line (US3D, in solid blue line), and the low-fidelity one (in dashed red line). The solution is well approximated by the low-fidelity solver.	88
4.3	Mach field: US3D+M ⁺⁺ solution on the top, US3D native solution on the bottom. The two solutions well agree with each other. . . .	94
4.4	Comparison of the QoIs along the stagnation line: US3D+M ⁺⁺ solution with solid line and US3D native solution with dashed line. The post-shock relaxation shows an excellent agreement, while slight discrepancies are observed in the BL.	94
4.5	Comparison of the QoIs around the surface: US3D+M ⁺⁺ solution with solid line and US3D native solution with dashed line. The experimental values are plotted with blue squares. The two pressure distributions well agree with each other and with the experimental data. The two heat flux distributions show some discrepancies. The US3D+M ⁺⁺ better predict the experimental data.	95
4.6	Mach contour at nominal condition: strong bow shock develops in front of the vehicle.	98
4.7	Temperature (left) and mass fractions (right) values along the stagnation line at nominal conditions. The jump in temperature drives the dissociation of the molecules. The resulting atoms diffuse through the BL and partially recombine because of the surface's temperature drop and the catalytic activity.	98
4.8	Top: Numerical grids used in the study. Bottom: zoom of the Mach contour on the mesh IV.	100
4.9	Angular distribution of the QoIs for the four meshes in Table 4.4 relative to the nominal conditions reported in Table 4.3. The solution obtained using the adapted mesh is plotted with dashed line, the one with a non-adapted mesh with solid line. The numerical error concentrates around the stagnation point, where the QoIs reach their maximum value.	101
4.10	Relative error on the QoIs as a function of the degree of refinement. The pressure, the force, and the heat load are in the monotonic convergence range. The observed order of convergence increases by employing grid adaption tools for the pressure and the force. The heat flux is outside the monotonic convergence range and the mesh adaptation improves the fit.	102
4.11	QoIs and relative numerical uncertainty as a function of the degree of refinement. The methodology is conservative: the uncertainty bars of a coarser grid consistently contain the finer grid ones; the only exception is the heat flux due to the used polynomial fit. Adaptation tools systematically reduce the numerical uncertainty. . . .	103
4.12	Projections of the surrogate models obtained for the adapted mesh II: no non-physical values are observed.	104
4.13	Heat flux obtained using mesh II.	104

4.14	QQplot obtained for the adapted mesh II: the values returned by the surrogate model are plotted against the ones obtain with CFD computations. The surrogate models well predict the verification points.	105
4.15	Statistical moments of the heat flux, obtained by propagating the uncertainties on the BCs, as a function of the degree of refinement of the meshes. They follow the same polynomial behavior observed in the nominal case.	105
4.16	Ratio between the average grid numerical uncertainty, $U_g(\phi)$, and the UQ-driven standard deviation, $\sigma(\phi, h_i/h_1)$, plotted with solid lines as a function of the degree of refinement. The standard deviation of the grid numerical uncertainty, $(S_g(\phi)/\sigma(\phi, h_i/h_1))$, divided by UQ-driven standard deviation is shown with opaque area. The numerical uncertainty is systematically lower when the mesh is adapted.	106
4.17	Numerical uncertainty as a function of the degree of refinement for a free-stream velocity of 5600 m/s. It is driven high when the carbuncle effect corrupts the solution on non-adapted grids.	107
5.1	Details of Mesh III of Table 5.1.	112
5.2	Numerical flow structure obtained with a total temperature of 7500 K and a total pressure of 16500 Pa imposed at the entrance of the sonic nozzle. Density gradient contours. Undisturbed flow on the top, probe-disturbed on the bottom. Structure: (a) Prandtl-Mayer expansion fan, (b) Jet boundary, (c) Barrel shock, (d) Mach disk, (e) Reflected shock, and (f) Detached shock. The triple point with a red circle.	114
5.3	Distribution of temperature along the jet central line: profile of the undisturbed flow in solid black line, probe-disturbed one in dashed red lines. The flow is nearly identical up to the edge of shock.	115
5.4	Temperature (top) and Mach (bottom) contours. The flow expands reaching a Mach number of around 4.5 before the shock. The temperature drops during the expansion and increases at the shock layer.	116
5.5	Mass fractions contours: atomic oxygen (top) and atomic nitrogen (bottom). The flow results chemically frozen through the whole expansion.	117
5.6	Temperature (left) and mass fractions (right) distribution along the jet center line, from the nozzle inlet to the probe stagnation point. The flow partially recombines in the BL, mostly driven by surface catalysis.	117
5.7	Dependency of the results versus the grid refinement for the three observables. Cross: CFD value; solid line: fit law according to Equation (4.9); bars: numerical uncertainty according to Equation (4.13).	118

5.8	Sketch of the adaptive/multi-fidelity surrogate model construction.	120
5.9	Iterative <i>NRMSE</i> for the three observables.	121
5.10	Added points: the majority of the points were sampled on the border of the hypercube.	121
5.11	QQplot: CFD response VS Surrogate model response. Points from the adaptive/multi-fidelity surrogate model on the 45-degree line. .	122
5.12	<i>NRMSE</i> as a function of the number of low-fidelity points associated with the α surrogate model (solid line), and with the δ model (dashed line). For the quantities: heat flux (in black), pressure (in red), mass flow rate (in blue).	122
5.13	Projection of the heat flux surrogate model on the T_0 and p_0 dimensions. No carbuncle effect observed.	123
5.14	Propagated uncertainties in red. Experimental uncertainties in black.	124
5.15	Prior and posterior marginal distributions for the four QoIs.	125
5.16	Joint distribution of γ_O and γ_N	125
5.17	Reconstructed points belonging to the jet boundary (plotted in yellow), the barrel shock (in green), and the detached shock (in red).	126
5.18	Prior and posterior uncertainty on the flow structure: jet boundary (in orange), barrel shock (in green), and detached shock (in red). Prior: mean of the propagation plotted with dashed line, bounded by dash-dotted lines indicating the 95% interval of confidence. Posterior: the continuous line represents the mean, while the shadow areas represent the uncertainty within the 95% of confidence. . . .	127
5.19	Prior (in red) and posterior (in black) uncertainties in the flow features.	128
6.1	Stagnation-point mass blowing rates obtained, for the four test conditions described in Section 2.4.1, with the stagnation-line code (black circle), and with the US3D solver (red cross). Experimental uncertainties indicated with black error bar. The two numerical solutions are in good agreement. However, both of them predict lower values of \dot{m} than the experiments and lack in capturing the upward trend of \dot{m} with respect to the surface temperature.	133
6.2	Probabilities of the nitridation (black) and recombination (red) reactions for the four test conditions described in Section 2.4.1, obtained with the stagnation-line code (circle). Uncertainties from the numerical-experimental rebuilding by Helber et al. plotted with error bars. High recombination probabilities are predicted. Nitridation probabilities in fair agreement with the numerical-experimental values.	134
6.3	<i>NRMSE</i> of the global surrogate model built on 14-dimensions with dashed red line. <i>NRMSE</i> of the 6-dimensions G6 surrogate model with solid black line.	138
6.4	Total Sobol indices for the mass blowing rates. The parameters at the BL edge and the surface temperature show poor influence on the mass blowing rate.	138

-
- 6.5 Probability densities of the efficiency and the activation energy of reaction 3. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al. 140
- 6.6 Probability densities of the efficiency and the activation energy of reaction 6. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al. 141
- 6.7 Probability densities of the efficiency and the activation energy of reaction 4. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al. 142
- 6.8 Probability densities of the efficiency and the activation energy of reaction 5. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al. 143
- 6.9 Probability densities of the activation energy of reaction 1 and total active site density. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al. (the three different lines corresponding to three distinct adsorption sites). 144
- 6.10 Predictions of the model calibrated with set A. The red and blue color indicates, respectively, the recombination and the nitridation efficiencies. The solid lines represent the mean of the prediction at molecular beam pressure (0.024 Pa), while the dashed lines at the Plasmatron pressure (1600 Pa). The shadow areas correspond to the prediction uncertainty (95% Confidence Interval (CI)), at low pressure on the left, and, on the right, at high pressure. The error bars with square as mean represent the molecular beam from Murray et al., while the error bars with circle as mean are the marginals obtained by del Val et al. 144

6.11	Predictions of the model calibrated with set B. The red and blue color indicates, respectively, the recombination and the nitridation efficiencies. The solid lines represent the mean of the prediction at molecular beam pressure (0.024 Pa), while the dashed lines at the Plasmatron pressure (1600 Pa). The shadow areas correspond to the prediction uncertainty (95% Confidence Interval (CI)), at low pressure on the left, and, on the right, at high pressure. The error bars with square as mean represent the molecular beam from Murray et al., while the error bars with circle as mean are the marginals obtained by del Val et al.	145
6.12	Model predictions of the surface coverage using the set B for the calibration. Solid line refers to the mean predictions at molecular beam pressure and the dashed line at Plasmatron pressure. The uncertainty corresponding to the 95% CI is shown, for both the pressure, with shadow area.	146
6.13	Mass blowing rate distributions. The continuous lines refer to the experimental uncertainty while the histograms show the propagated distributions. Blue, gray, red, and green colors indicate, respectively, the G4, G5, G6, and G7 case.	146
6.14	Mass blowing rate distributions obtained propagating the parameters posteriors calibrated with set C. The continuous lines refer to the experimental uncertainty while the histograms show the propagated distributions. Blue, gray, red, and green colors indicate, respectively, the G4, G5, G6, and G7 case.	147
6.15	Predictions of the model calibrated with set D. The red and blue color indicates, respectively, the recombination and the nitridation efficiencies. The solid lines represent the mean of the prediction at molecular beam pressure (0.024 Pa), while the dashed lines at the Plasmatron pressure (1600 Pa). The shadow areas correspond to the prediction uncertainty (95% Confidence Interval (CI)), at low pressure on the left, and, on the right, at high pressure. The error bars with square as mean represent the molecular beam from Murray et al., while the error bars with circle as mean are the marginals obtained by del Val et al.	148
6.16	Mass blowing rate distributions obtained propagating the parameters posteriors calibrated with set D. The continuous lines refer to the experimental uncertainty while the histograms show the propagated distributions. Blue, gray, red, and green colors indicate, respectively, the G4, G5, G6, and G7 case.	149
6.17	Joint distribution of the reaction probabilities.	150
A.1	Evolution in time of the mass fractions in an air mixture composed of five species in the 0D reactor case. The code-to-code comparison shows excellent agreement.	160

A.2	Comparison of the temperature (left) and mass fractions (right) along the tube for the 1D diffusion case. The solutions of the two codes perfectly agree with each other.	161
A.3	N ₂ mass fractions for different recombination coefficients γ for the 1D catalytic diffusion problem. The code solution perfectly agrees with its analytical counterpart.	162
A.4	Comparison of the QoIs along the stagnation line obtained with the stagnation-line code (dashed line) and the US3D solver (solid line). The two solutions agree well both for the temperature and the mass fractions profile in the BL.	163

List of Tables

1.1	Experiments and simulations: main differences.	5
2.1	Subset of the FRC reactions involving nitrogen in the ACA model.	38
2.2	Overview of Plasmatron test conditions for a nitrogen flow at 1500 ± 50 Pa over a 25 mm radius hemisphere carbonaceous sample: dynamic pressure p_d , generator power P_{el} , sample mean temperature T_w , mass blowing rate \dot{m} , and reconstructed BL edge conditions (temperature, T_e , and velocity, u_e). Conditions from Helber et al. [3].	46
4.1	Summary of the verification test cases.	92
4.2	Free-stream conditions for the 2D hypersonic flow over a probe.	93
4.3	Nominal conditions and associated uniform uncertainties for free-stream density, ρ_∞ , free-stream velocity, u_∞ , and recombination probability, γ	97
4.4	Numerical grids used in the study: tag of the mesh, number of nodes, number of cells, normalized characteristic length, h_i/h_1 , wall-normal distance of the first cell at the stagnation point, Δn , sonic Reynolds number based on the first physical cell, and time required to perform 5.000 iterations on 16 cores.	99
4.5	Orders of convergence on non-adapted (NA) and adapted (A) grids for the pressure, p_w , the force, F , and the heat load, Q	101
5.1	Details of the meshes used: index, number of cells, length of the first cell at the stagnation point (Δn), normalized characteristic mesh (h_i), and time to converge (t_{CPU}).	113
5.2	Results of the convergence study performed on the meshes in Table 5.1. For each observable: order of convergence, p , asymptotic value, ϕ_0 , numerical uncertainty of Mesh I, U_1 , and experimental uncertainty, σ_{exp}	116
5.3	Informative Gaussian experimental uncertainty.	119
5.4	Non-informative uniform uncertainty.	119

6.1	Non-informative uniform prior distributions on the parameters to calibrate, along with the values proposed by Prata et al. and Nieman et al. *The three values correspond to three different adsorption sites.	135
6.2	Matrix of the performed tests.	137
A.1	Initial conditions for the 0D reactor case.	160
A.2	Boundary and initial conditions for the 1D diffusion case.	160
A.3	BCs for the 1D catalytic diffusion case.	161

List of Acronyms

A

ACA	Air Carbon Ablation
AIES	Affine Invariant Ensemble Sampler

B

BC	Boundary Condition
BGK	Bhatnagar–Gross-Krook
BL	Boundary Layer

C

CFD	Computational Fluid Dynamics
CFL	Courant-Friedrichs-Lewy
CI	Confidence Interval

D

DG	Discontinuous Galerkin
DRNSE	Dimensionally Reduced Navier-Stokes Equations
DPLR	Data Parallel Line Relaxation
DSMC	Direct Simulation Monte Carlo

E

ER Eley-Rideal
EI Expected Improvement

F

FRC Finite-Rate Chemistry
FV Finite Volume

G

GP Gaussian Process
GSI Gas-Surface Interaction

H

HIAD Hypersonic Inatable Aerodynamic Decelerator

I

ICP Inductively Coupled Plasma

K

KEC Kinetic Energy Consistent

L

LES Large-Eddy Simulations
LH Langmuir-Hinshelwood
LHS Latin Hypercube Sampling
LHTS Local Heat Transfer Simulation

LOO Leave-One-Out
LTE Local Thermodynamic Equilibrium

M

MC Monte Carlo
MCMC Markov Chain Monte Carlo
MUSCL Monotonic Upstream-centered Scheme for Conservation Laws

P

PCE Polynomial Chaos Expansion
PDF Probability Density Function
PICA Phenolic Impregnated Carbon Ablator

Q

QoI Quantity of Interest

R

RANS Reynolds-Averaged Navier-Stokes
RHS Right Hand Side
RRHO Rigid Rotor Harmonic Oscillator

S

SMB Surface Mass Balance
SEB Surface Energy Balance
SCEBD Self-Consistent Effective Binary Diffusion
SST Shear Stress Transport

T

TOF	Time-Of-Flight
TPM	Thermal Protection Material
TPS	Thermal Protection System

U

UQ	Uncertainty Quantification
----	----------------------------

V

VKI	von Karman Institute
-----	----------------------

1

Introduction

*September 16, 2021, Kennedy Space Center: Falcon 9 launches the Dragon capsule and the Inspiration4's crew for the first private citizen's orbital spaceflight. A new era has come, where private ventures are taking the way opened sixty years ago by the space agencies. After three days of orbital flight, Dragon's thrusters perform the last burn, and the nose-cone closes. Dragon is ready to experience one of the most fascinating and challenging features of space flights: *the atmospheric entry.**



Figure 1.1: SpaceX's Dragon re-entering atmosphere. Credit: SpaceX.

1.1 Atmospheric entry flows

During the atmospheric entry from low Earth orbit, spacecraft approach a peak velocity of 8 km s^{-1} . It is a velocity 20 times higher than the speed of sound at standard conditions and the kinetic energy that they have to dissipate is enormous. While this is beneficial for space debris, such as meteors or satellites, as they get degraded through the atmosphere, spacecraft need to be protected for a safe return of both the crew and payload.

The blunt body design was proposed in 1951 by Allen to minimize the amount of energy transferred to the vehicle. Such a geometry provokes a strong detached shock ahead of the vehicle, as sketched in Figure 1.2. Across it, the kinetic energy of the flow is converted into particles' translational energy [4, 5]. Then, progressively, the internal energy is excited rotationally, vibrationally, and electronically. In the shock layer, the gas is outside chemical equilibrium and emits radiation because of the high temperatures reached. The chemistry kicks in and the molecules dissociate and ionize, driving down the temperature. When the shock is sufficiently far from the body, the particles reach a thermo-chemical equilibrium plateau before the Boundary Layer (BL) that develops in front of the vehicle. Within the BL, strong temperature and velocity gradients cause the chemistry to be again outside equilibrium: a chemically reactive gas diffuses toward the surface and interacts with it through heat and mass exchange. Such a phenomenon is generally referred to as Gas-Surface Interaction (GSI).

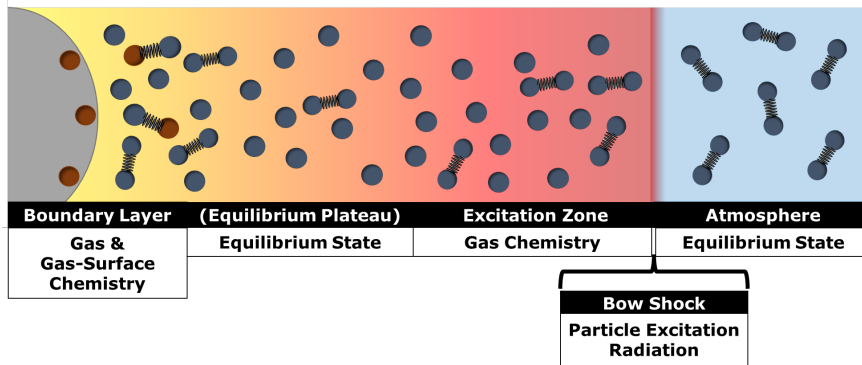


Figure 1.2: Sketch of an entry hypersonic flow.

Even if lowered by using a blunt geometry, the heat flux experienced by the spacecraft is still significant and they need to be protected by means of a Thermal Protection Material (TPM). There are two main classes of TPMs: ablative and reusable materials [6]. Ablative TPMs are generally made up of carbon fibers in a

phenolic resin. When they get heated by the hot gas, the virgin material degrades into a porous char and pyrolysis gases. Such a reaction is endothermic and adsorbs part of the incoming heat. Additionally, the charred material gets ablated by exothermic chemical reactions (mostly oxidation) and mechanical removal. Pyrolysis gases and ablative products are injected in the BL creating a barrier that further protects the vehicle from the incoming heat (*blowing effect*). Depending on the TPM, the surface might also promote exothermic recombination reactions. Because they experience significant mass loss, ablative TPMs have to be changed every mission. Contrarily, reusable TPMs, generally ceramic materials, can withstand multiple atmospheric entry missions. They do not experience significant degradation and cool down by re-radiating most of the incoming energy. The surface might promote catalytic reactions (recombination of atoms at the wall). Being this reaction exothermic, energy is released to the surface, increasing the heat flux. Hence, they are characterized by a low catalytic surface, low thermal conductivity, and high emissivity. These materials cannot withstand severe heat loads, and they are mostly used for orbital re-entry.

Depending on the type of the TPM, ablative and/or catalytic reactions can occur. Ablative reactions cause material degradation and shape change, which, in turn, alters the aerodynamic coefficients. Oxidation and catalytic reactions are strongly exothermic, increasing the surface temperature. In both cases, the thermo-chemical environment has to be correctly characterized for a robust simulation of the vehicle heat loads.

1.2 Flow and material characterization in hypersonic applications

As discussed above, atmospheric entry flows, and more generally hypersonic flows, are complex phenomena involving several physical features, such as high-temperature effects, finite-rate chemical processes, GSI, laminar-turbulent transition, turbulence, and radiation. A robust design of the hypersonic vehicles requires the characterization of the flow conditions and the TPM response. To this end, ground-testing facilities are widely used as they can duplicate the physical features encountered in flight and, compared to flight testing, they are less expensive and allow for better controlling the flow conditions [7]. To date, no facility is able to obtain a proper scaling of the phenomenon and the common practice is to decouple the high-velocity from high-temperature effects [7]. For instance, the aerodynamic forces experienced by hypersonic vehicles can be reproduced in a hypersonic wind tunnel, such as the Longshot facility at the von Karman Institute (VKI) [8], by jointly duplicating the in-flight Mach and Reynolds numbers. These facilities are

characterized by short testing time and low enthalpy of the gas, and the relevant heat loads cannot be reproduced. These can be obtained in a high-enthalpy, but low-speed, ground testing facility, such as the VKI Plasmatron [9], by duplicating the in-flight total enthalpy, stagnation pressure, and velocity gradient [10]. Such duplication ensures the reproduction of the same Damköhler number (ratio between the flow and the chemistry characteristic times) that characterizes the in-flight BL chemical non-equilibrium. The biggest drawback of the experimental approach is that an experimental campaign can be still rather time-consuming and expensive, making it prohibitive to analyze a large number of configurations.

Alternatively, Computational Fluid Dynamics (CFD) solvers can simulate several flow conditions in a cost-efficient manner. Among others, US3D [11] has been developed to robustly compute hypersonic flows. For example, it has been employed to simulate several aspects of the atmospheric entry, including I) steady-state configurations [12], II) ablative shape changing, coupling the code with a material response solver and mesh motion algorithm [13, 14], III) turbulent and unsteady wakes, leveraging high-order numerical schemes [15–17], and IV) dynamic stability [18]. In the case of numerical predictions, the biggest challenge arises from properly describing the complex physics involved in the hypersonic regime. In fact, CFD solvers embed several physical models for finite-rate chemistry, thermal relaxation processes, multi-component transport, and GSI, to name some. Correctly modeling them is not a trivial task. For instance, B' -tables have been widely used for predicting the ablative response of TPMs. This model assumes that the gas is in chemical equilibrium and overestimates the surface recession when this assumption is not valid [19, 20]. In these cases, finite-rate models should be used. The first of such models was proposed by Park [21]. It describes the macroscopic oxidation and nitridation reactions at the surface by means of Arrhenius formulations. Another widely used model was derived by Zhukov and Abe [22]. It includes detailed processes, such as the adsorption/desorption of atoms at the surface, and describes both the forward and backward mechanisms. Differently from the Park model, the latter does not account for the nitridation reaction, which was later added, in a modified version, by Alba et al. [23]. More recently, Poovathingal et al. [24] leveraged informative molecular beam data to derive a detailed ablation model, which was later improved by Prata et al. [1], supported by new experiments. It is important to remark that substantial differences were found in the predictions of these models [25–28], and that GSI characterization is, to date, an active area of research.

Experiments offer a natural source of data for the calibration and subsequent validation of the model parameters [7]. It is therefore evident that a tight complementarity between experiments and numerical simulations exists. The main

differences between the two approaches are reported in Table 1.1. On top of what was discussed above, experimental observations are affected by *uncertainty* in the *measurements*, while the outcome of the numerical simulations, excluding those relying on sampling strategies, is deterministic, it can be perfectly duplicated [29]. In fact, the uncertainty associated with numerical simulations is upstream, linked to its input. For instance, it can derive from the simulated geometries, the Boundary Condition (BC) imposed, the mesh used (and the relative numerical error), or, as we previously saw, from the *form* of the *models* and the choice of their *parameters*.

Experiments	Numerical Simulations
Time-consuming/Costly	Cost-efficient
The physics is intrinsic	The physics is modeled
Uncertainty in the measurements	Uncertainty in the inputs

Table 1.1: Experiments and simulations: main differences.

These uncertainties can severely undermine the accuracy of the numerical predictions, especially for atmospheric entry flows as they require the modeling of several physical features. Uncertainty Quantification (UQ) methods aim at identifying, characterizing, and potentially reducing the uncertainties involved in the studied phenomena, allowing one to obtain robust predictions. UQ analysis can be divided into two main categories. The first one concerns the *forward propagation* of the input uncertainties to estimate the one on the model output. The forward model can also be used to determine which input uncertainty affects the most the one on the output. Such a method is referred to as *sensitivity analysis*. The second category regards the *stochastic inverse* or *inference problem*, which aims at reducing the uncertainty on some model input, given the uncertainty on the model output. Stochastic inverse methods offer a robust bridge between experiments and simulations, as they allow for accounting for the different sources of uncertainty. The steps involved in such a stochastic approach are sketched in Figure 1.3: I) an experiment is designed to reproduce a Quantity of Interest (QoI), Y , which is measured with a given uncertainty; II) a model, $f(X)$, is derived to describe the experiment, input are the experimental conditions and the unknown model parameters; III) in the majority of the applications, such a model is solved in a discrete fashion, introducing a *discretization error*. In this case, one should *verify* that the mathematical structure of the model is correctly implemented in the solver and that the discretization error in the solution is negligible. IV) The unknown model parameters are inferred and validated by comparing the model output to the experimental observations. The analysis also allows for characterizing the uncertainty associated with the model parameters. Hence, V) when the model is exploited to reproduce the phenomenon

of interest, $f'(X)$, these uncertainties can be propagated to estimate the uncertainty associated with a given QoI.

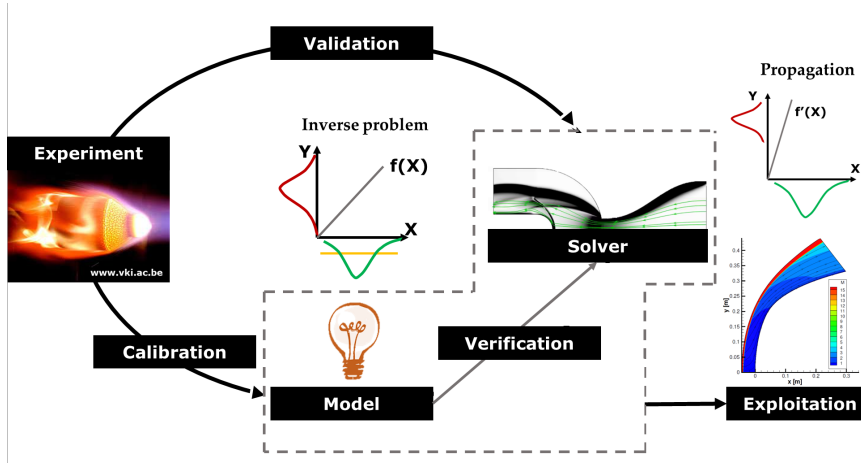


Figure 1.3: Sketch of a stochastic inference process.

A central element in both the uncertainty propagation and the inference problem is the model employed to represent the phenomenon of interest. It expresses the relationship between the input and the output. In the above-mentioned context, the input of the model would be the parameters that one intends to calibrate (e.g. the catalytic efficiency of a TPM surface), along with the conditions defining the experiment (e.g. the surface temperature), while the output would be a related experimental observation (e.g. the surface heat flux). Different models can be used to describe the phenomenon. Each of them embeds a given amount of assumptions and/or simplifications that define its degree of accuracy, or, said differently, its *fidelity*. Following the definition proposed by Peherstorfer et al. [30], ‘a model that estimates the output with the accuracy that is necessary for the task at hand’ is a *high-fidelity model*. Contrary, a *low-fidelity model* is ‘a model that estimates the same output with a lower accuracy than the high-fidelity model’. Lower fidelity could arise from a reduction of the problem dimension, the use of a coarser grid, or a simplification in the physical model, to name some. However, one should keep in mind that a low-fidelity model does not necessarily produce a low-fidelity solution. Let us imagine describing a plasma flow over a probe. Resorting to a finite-rate chemical model would have higher accuracy than assuming chemical equilibrium. However, if the flow is characterized by small velocity gradients, such that it has the time to relax towards its equilibrium condition, the assumption contained in the low-fidelity model is adequate, hence the results are accurate. That said, higher accuracy is

generally expected by models embedding fewer assumptions. In this manuscript, we will refer to a high-fidelity solver as one containing fewer assumptions and which is expected to produce the most accurate solution and to a low-fidelity solver as the one with a higher degree of assumptions.

Since UQ analysis generally requires a large number of computations, low-fidelity solvers are attractive as they are characterized by a pronounced efficiency. However, if they are used outside the range of validity of their underlying assumptions, or if the model inadequacy is not correctly accounted for, they may lead to erroneous results. In these cases, one should resort to high-fidelity simulations. Surrogate models represent a convenient method to restore the efficiency required to perform UQ studies when computationally expensive high-fidelity solvers are needed [31]. They are trained on a limited number of simulations (much lower than the ones required to perform the UQ study) and capture the underlying relationship between input and output. However, there might be situations where the cost associated with high-fidelity simulations is so high that it is not affordable even producing only the simulations requested to adequately train the surrogate model. To cope with this problem, multi-fidelity models have been proposed [30]. They allow for leveraging low-fidelity solutions to reduce the number of high-fidelity computations.

Different models, characterized by different degrees of fidelity, have been used to perform hypersonic-related UQ studies. In the following, we revise the state-of-the-art concerning the forward propagation first, and then the inverse problem, highlighting the fidelity of the solvers used in the analysis and the strategies employed to improve the efficiency of the method. We also distinguish works that focus on atmospheric entry flows from those addressing specific features of the flow, duplicated in on-ground facilities.

Pioneer works focused on the characterization of uncertainty related to the stagnation-point convective and radiative heat flux [32–34]. In these studies, the uncertainties on some hundred parameters (including chemical rate and transport coefficients, relaxations rates, and catalytic efficiencies) were propagated employing a Monte Carlo (MC) technique directly coupled to the numerical solver. The analysis required the evaluation of thousands of costly 2D axisymmetric high-fidelity CFD simulations. Forward propagation of the model uncertainties was also performed to characterize the ablative response of TPMs [35, 36]. In these studies, the analysis was conducted by means of a more efficient 1D material response solver. Similarly, 1D flow computations were employed by Ghaffari et al. [37] to quantify the uncertainties related to stagnation-point radiative heat flux. With respect to Bose et al. [34], the reduced computational cost associated with the dimension reduction, allowed them to increase the fidelity in the physical modeling, and therefore the robustness of the numerical predictions. Improvement in the method efficiency was

later achieved by approximating the computationally expensive high-fidelity solvers with cheap surrogate models. The number of computations requested to perform the UQ studies was reduced to some hundreds. For example, Hosder and Bettis [38] employed a Polynomial Chaos Expansion (PCE), constructed on 2D axisymmetric high-fidelity computations, for characterizing the uncertainty on the surface heat loads under atmospheric entry conditions. Other similar applications can be found in [39–42]. West et al. [43] coupled the surrogate model to a sensitivity-based dimension reduction strategy to further enhance the method’s efficiency. It was applied for characterizing the stagnation-point radiative heat flux on a Hypersonic Inatable Aerodynamic Decelerator (HIAD), whose flowfield was simulated through 2D axisymmetric high-fidelity computations. In a subsequent work [44], a subset of the training points required for a total-order expansion was used to further reduce the computational effort. The surrogate model accuracy was assessed by monitoring the evolution of the Sobol indices and model error, computed on verification points. The latter strategy was then employed by Brune et al. [45] to study coupled fluid-structure interaction for a HIAD, also simulated employing expensive high-fidelity computations. The characterization of the uncertainties related to the surface heating on HIADs was very recently tackled by Santos et al. [46, 47] in a multi-fidelity framework. They used a co-Kriging strategy that leverages low-fidelity simulations for the construction of the surrogate model, reducing the number of high-fidelity computations to those required to preserve accuracy. A co-Kriging method was also employed by Quinlan et al. [48] to approximate the high-fidelity response of a hypersonic flow over a blunted cone. It was used to construct an aerodynamic database for trajectory simulations.

Forward propagation studies addressed also the characterization of the flow and the material response in ground testing. Sorensen et al. [49] employed a MC technique to propagate the uncertainties on selected rates of a finite-rate catalytic model to assess their influence on the heat flux experienced by a cylinder immersed in a hypersonic flow. The MC strategy was directly coupled to a high-fidelity solver. Sanson et al. [50] estimated the uncertainty on the catalytic efficiency of a TPM, and identified the parameters contributing the most to its variability, from subsonic Plasmatron experiments. Experimental and modeling uncertainties were propagated through a PCE model constructed on 1D stagnation-line computations of the BL. A similar approach was employed by Turchi et al. [51, 52] with regards to ablation-related QoIs. Brune et al. [53] estimated the uncertainties associated with the stagnation-point pressure and heat flux experienced by a calibration probe exposed to a low- and a high-enthalpy supersonic flow obtained in arc-jet experiments. The forward analysis was performed by means of a PCE trained on high-fidelity 2D axisymmetric computations. Georgii and Volker [54] used a PCE built on inviscid quasi-1D computations of expanding high-enthalpy flows in nozzles to assess the sensitivity of the free-stream conditions to finite-rate processes and reservoir condi-

tions.

The works above mentioned concerned only the forward problem. The UQ literature also includes stochastic inferences related to hypersonic flows. They mostly concern the reconstruction of free-stream conditions, the calibration of GSI chemical models, and the characterization of gas-phase chemical rate coefficients. Tryoen et al. [55] attempted to reconstruct the free-stream conditions in a trajectory point of the EXPERT vehicle from the stagnation-point measurements of pressure and heat flux. A PCE was constructed from 2D axisymmetric simulations. The CFD solver turned out to be not robust under all the free-stream conditions, resulting in poor estimations of the heat flux. Consequently, the resulting surrogate model predicted non-physical negative values, and the analysis was performed with only stagnation-point pressure. This problem was later tackled by Cortesi et al. [56] rebuilding the free-stream conditions of a hypersonic flow over a cylinder, also from stagnation-point measurements. To improve the heat flux numerical predictions the mesh was automatically adapted to capture the shock at each training point. Furthermore, an active subspace was exploited to reduce the dimensionality of the input space. In a different study, Cortesi et al. [57] exploit the possibility of rebuilding the free-stream conditions, the angles of attack, and the catalytic efficiency of the surface in a trajectory point of the EXPERT vehicle. Both stagnation- and off-stagnation-point measurements provided by sensors located in the vehicle forebody were used in the Bayesian analysis. To this end, 2D axisymmetric simulations were needed to capture the angular dependency of the QoIs. Similarly, Ray et al. [58] employed a Bayesian methodology to infer the free-stream conditions in three experiments consisting of a hypersonic flow over a double cone. The pressure and heat flux measurements from those sensors located in the region where the flow was attached were used as trustworthy observations. The resolution of the complex shock interactions and of the separation regions required the computations of 2D axisymmetric simulations. However, the authors concluded that 3D simulations might be necessary to explain some of the departure observed in the simulated predictions from the experimental measurements. As a result, these 2D axisymmetric simulations should be considered low-fidelity predictions. Chowdhary et al. [59] also rebuilt the free-stream conditions from pressure and heat flux measurements on the surface of the HIFiRE-1 probe exposed to a hypersonic flow. Also in this case, 2D axisymmetric simulations were performed. Unlike the laminar applications previously mentioned, turbulence was expected in this test. It was described using a Shear Stress Transport (SST) model, whose parameters were inferred during the same analysis.

Regarding the calibration of chemical models for GSI, Upadhyay et al. [60] firstly employed a Bayesian framework to infer the probability of the nitridation reaction from a furnace-heated quartz tube experiment [27]. The flow in the tube was

reconstructed through a 1D approximation. Four different models, differing in the assumptions used to describe the nitridation surface reaction and the considered experimental uncertainties, were employed in the Bayesian analysis. The experimental data informed a temperature-dependent nitridation probability law for each model. Later, Sanson et al. [61] estimated the efficiency of the catalytic reaction on ceramic matrix composite samples from Plasmatron experiments. A 1D solver was used to describe the BL developing in front of the sample. The authors first recognized the importance of prescribing an epistemic uncertainty also to the catalytic efficiency of the calorimeter employed to measure the BL edge enthalpy. However, it turned out that the experimental data were not enough informative to learn the targeted efficiencies. In a following work, del Val et al. [62] upgraded the inference methodology by including in the process a coupled optimization procedure to estimate some nuisance parameters, such as the BL edge enthalpy. This approach led to an improved exploitation of the experimental data, resulting in more accurate posterior predictions of the catalytic efficiencies. Later, the same authors [63] characterized the efficiency of the nitridation reaction from Plasmatron experiments. Specifically, an Arrhenius law was calibrated considering jointly four experimental points at distinct surface temperatures. In a subsequent work [64], some of the model assumptions contained in the previous calibration, such as no surface recombination and thermal equilibrium, were relaxed by introducing a model-form uncertainty in the inference problem. An Arrhenius law was computed by means of a Bayesian model averaging strategy.

Stochastic inference methodologies have also been employed in the characterization of gas-phase properties. Miki et al. [65, 66] performed a Bayesian calibration of several parameters included in the finite-rate chemistry and in the radiative models from shock-tube spectrographic data. The plasma radiance was computed with a line-by-line radiation code fed with 1D inviscid computations of the shock layer. Recently, del Val and Chazot [67] investigated the possibility of learning gas chemical rates, along with the probes' catalytic efficiencies, from Plasmatron experiments. The BL developing in front of the sample was simulated by means of 1D computations.

1.3 Objectives and outline

Figure 1.4 is intended to summarize, for application and fidelity, the UQ works reviewed in the previous section. The studies resorting to geometrical assumptions or where the solver was not capable of correctly reproducing the phenomenon of interest have been classified as low fidelity. Those applications in which a surrogate model was built fusing low- and high-fidelity information were classified as multi fidelity, while the remaining works as high fidelity. We remark that such a

classification was performed purely by considering the solver assumptions. Thus, studies that have been performed with low-fidelity solvers might result in accurate predictions if the underlying assumptions are adequate for the task at hand.

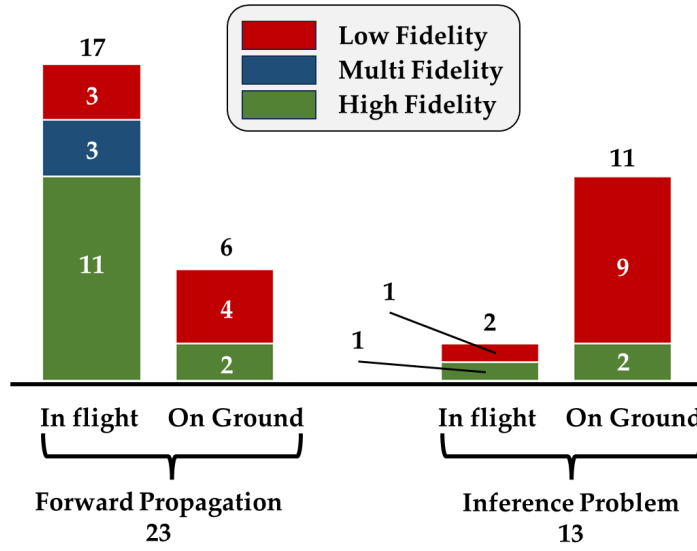


Figure 1.4: Summary of the UQ works related to atmospheric entry flows reviewed in Section 1.2.

The figure depicts that the biggest body of the hypersonic-related UQ literature concerns forward propagation studies. Fewer works focused on solving the inverse problem. Furthermore, the vast majority of these studies have been performed resorting to efficient low-fidelity solvers. This highlights a gap in the current literature concerning the use of high-fidelity models in hypersonic-related inverse problems.

Hence, the overarching purpose of this thesis is to **develop tools and methodologies to obtain accurate predictions of atmospheric entry conditions using high-fidelity simulations and state-of-the-art experimental data**. For this reason, an overall UQ framework, embedding a surrogate-based approach coupled to mesh-error estimations and multi-fidelity methods, is employed to solve inference problems. We considered both in-flight and on-ground testing conditions, spanning a wide range of flow regimes: from subsonic to hypersonic velocities. These flows might include high-temperature and multi-components effects, homogeneous and heterogeneous chemistry, GSI, strong detached shocks, shock-shock interactions, unsteadiness, and turbulence. High-fidelity simulations are then required to fully capture such complex and rich physics. The associated numerical solutions strongly

depend on the model parameters (that most of the time are affected by severe uncertainties) and on the mesh employed to discretize the physical domain. Furthermore, their computation is time-consuming, making them poorly efficient in a UQ framework, where the evaluation of many model outputs is required.

The main objectives of the thesis are the following ones:

- O1: to produce an accurate high-fidelity reference solution** for each phenomenon of interest. For what discussed above, it is essential to employ a high-fidelity multi-dimensional solver equipped with specific BCs to deal with chemically reacting surfaces, that is robust under hard-to-compute hypersonic conditions. Such a reference solution can be used to assess the accuracy of lower-fidelity solvers and determining whether they are adequate for performing UQ studies.
- O2: to investigate the benefits and the challenges associated with the systematic use of a high-fidelity solver in a UQ framework.** On the one hand, as the model fidelity increases, the resolution of the flow improves and we get access to more, and more reliable, features. This guarantees a **higher accuracy of the inference problem**. Furthermore, having access to more flow features is desirable as one can compare additional numerical predictions to experimental observations, **potentially reducing the posterior uncertainties** on these quantities that one intends to characterize. On the other hand, UQ analysis requires the computation of a large number of simulations. It is computationally expensive, especially when high-fidelity solutions are needed. Thus, we intend **to restore the efficiency of the method** by coupling surrogate-based techniques to mesh-error estimates and multi-fidelity strategies. Furthermore, a fixed mesh is generally used to perform all the simulations requested in the UQ analysis. This practice might be inaccurate in the presence of strong shocks, where grid misalignment can lead to numerical errors. For this reason, **we explore the benefits of systematically resorting to alignment tools** to automatically produce a shock-error-free solution for each condition.
- O3: to use (sometimes limited) experimental data to reduce the uncertainty on the prediction of the QoIs.** Bayesian methods have been proven to offer an optimal framework to account for all the different sources of uncertainty involved in the inference process. Hence, **we will employ a Bayesian methodology**, accelerated by the above-discussed surrogate-based strategy, to characterize selected model parameters leveraging state-of-the-art experimental data.

To investigate these objectives, we performed several actions. With regard to the first objective:

A1: We employed the US3D solver [11] to produce high-fidelity solutions of each phenomenon of interest. We coupled it to the open-source Mutation⁺⁺ library [68], which provides state-of-the-art thermodynamic and transport properties, as well as the solution for GSI balances, essential to describe chemical reacting surfaces. We included in Mutation⁺⁺ a recently developed detailed ablation model. It captures both the temperature and the pressure dependence of the surface response, enabling us to explain experiments characterized by different operating pressures with the same ablation model.

With respect to the second objective:

A2: We proposed a methodology to balance the grid's numerical errors and problem-related uncertainties to produce the optimal representation for a given computational budget. We applied it to characterize the surface pressure and heat flux experienced by a hypersonic entry object under uncertainty both in the free-stream conditions and in the catalytic efficiency of the TPM. It was simulated by means of axisymmetric computations performed with the US3D solver. Furthermore, we investigated the benefits of systematically resorting to a mesh-shock alignment tool.

A3: We developed an adaptive/multi-fidelity strategy to alleviate the computational cost associated with the construction of a surrogate model from high-fidelity simulations. We apply it to describe an under-expanded high-enthalpy jet recently obtained in the Plasmatron facility at the VKI. Particularly, we were interested in the computations of the stagnation-point pressure and heat flux, and of the mass flow rate in the nozzle. A low-fidelity model was derived to determine them. However, it underestimated their values compared to the high-fidelity counterpart. For this reason, 2D axisymmetric US3D computations were needed. Moreover, the grid required to accurately compute the QoIs was too expensive to be directly employed to construct the surrogate model, making necessary the use of the adaptive/multi-fidelity strategy. Furthermore, we investigated what additional information can be extracted by the high-fidelity simulations. These data can be potentially observed in future experiments to enrich the inverse analysis.

Finally, concerning the third objective:

A4: We characterized the previously-mentioned under-expanded high-enthalpy jet, for which no standardized rebuilding procedure existed to date. Specifically, we investigated whether the conditions at the inlet of the sonic nozzle, and the catalytic efficiencies of the probes, can be reconstructed by measuring the stagnation-point pressure and heat flux, and the mass flow in the

nozzle. To this end, we resorted to a Bayesian methodology, accelerated by the adaptive/multi-fidelity surrogate model.

A5: We inferred the rates of the elementary reactions of an ablation model from two distinct experiments. Specifically, we intend to verify whether the observations obtained in a molecular beam-surface scatter apparatus, at high-vacuum conditions, and in the VKI Plasmatron facility, at much higher pressure, can be compatible through the same ablation model. Compared to the phenomenological models used in the other two applications, **we employed a higher-fidelity GSI model to capture the ablative behavior of the surface at different temperatures and pressures**, as the ones encountered in the two experiments. On the other hand, 1D simulations were employed to accurately describe the Plasmatron conditions, while 0D computations to explain the molecular beam-surface scatter experiment.

The manuscript is structured as follows:

Chapter 1 introduced the atmospheric entry flows and the challenges associated with their duplication. After a review of previous UQ works, we outlined the objectives of the thesis.

Chapter 2 revises the governing equations used to describe the phenomena of interest, as well as the closure for the chemistry, the thermodynamic and the transport properties, and the GSI modeling. The Mutation⁺⁺ library is presented too. Finally, the experimental campaigns supporting the UQ analysis are introduced.

Chapter 3 introduces the UQ methodologies and tools. Specifically, we describe both the forward and the inverse problem, along with efficient algorithms for their solution, and the construction of (adaptive/multi-fidelity) Kriging surrogate models.

Chapter 4 describes the solvers employed in thesis. We present also the coupling that we have performed between the US3D code and Mutation⁺⁺, along with the verification test cases. Finally, we show the results of the characterization of the pressure and heat flux experienced by the EXPERT vehicle under both BC and numerical uncertainties.

Chapter 5 presents the methodology employed to characterize an under-expanded high-enthalpy jet obtained in the Plasmatron facility. Specifically, we resorted to a multi-fidelity adaptive Bayesian framework to characterize both the conditions at the entrance of a sonic nozzle and the catalytic efficiencies of the probe used to measure stagnation-point heat flux and pressure. The variability of the flow structure is also exterminated.

Chapter 6 describes the characterization of the rates of a finite-rate model for carbon ablation. Specifically, the rates were informed by means of both Plasmatron and molecular beam experiments.

Chapter 7 draws down the overall conclusions of this thesis and outlines future perspectives.

2

Physical modeling

This chapter provides an overview of the governing equations and of the physical models used to describe atmospheric entry flows. First, chemical reacting Navier-Stokes equations are presented along with the closure for transport fluxes and chemical reaction rates. The modeling of the thermodynamic properties of the gas in the presence of high-temperature effects is also discussed. An important aspect of this thesis is the description of the chemical reactions, either catalytic or ablative, between a gas and a surface. The description of two different closures for gas-surface interaction, phenomenological and finite-rate coefficient-based, for the chemical source terms is provided. A state-of-the-art air carbon ablation finite-rate model is presented and included in the Mutation⁺⁺ library, which is described next. Mutation⁺⁺ is also employed during the CFD simulations to compute both the thermophysical properties of the gas and to solve the gas-surface interaction balances. We finally introduce three different experimental campaigns whose data served as input for the analysis in the following chapters. The first two are subsonic and supersonic tests performed in the VKI Plasmatron facility. This facility allows for determining macroscopic aspects of the surface reactions, such as the mass blowing rates and heat fluxes. The third one is a molecular beam-surface scattering experiment, which provides more insights into the nature of the surface reactions.

2.1 Overview

The description of hypersonic flows, including atmospheric entries, is a challenging task, involving a multi-disciplinary description of the phenomena, sophisticated numerics, and complex experiments. Vehicles entering the atmosphere go through different flow regimes: free molecular, transitional, and continuum [69]. Typically, the Knudsen number, Kn , is used as a reference quantity to determine in which regime the gas is. It is defined as the ratio between the mean free path (average distance a particle has to travel between one and the consequent collision) and a characteristic length of the phenomenon.

The free molecular regime, $Kn > 10$, is encountered at altitudes higher than 100 km and during rapid expansions. The gas is extremely rarefied and is characterized by a negligible amount of particle collisions. The velocity distributions strongly depart from the Maxwellian equilibrium limit [70]. The Boltzmann equation, mostly solved with Direct Simulation Monte Carlo (DSMC) [71] and Bhatnagar–Gross–Krook (BGK) [72, 73] methods, is used to describe the flow. At lower altitude, the vehicle enters the slip flow regime, $0.1 < Kn < 10$, where a continuum formulation, supplied with velocity slip and temperature jump BCs, might be used to describe the flow [74]. At altitude lower than 70 km, the vehicle is in the continuum regime ($Kn < 0.1$): the density increases, driving up the rate of collisions [69].

This manuscript focuses on the characterization of entry flow conditions in the continuum regime. Transport phenomena distribute mass, momentum, and energy, and the equilibrium regime can be approached, allowing one to describe the flow using the Navier-Stokes equations. They describe the macroscopic behavior of the flow in terms of bulk properties (such as density, velocity, and temperature) and enforce the conservation of mass, momentum, and energy: the variation in time of a conservative variable in a given volume is a balance between fluxes on its boundary and internal production.

Both in-flight and on-ground applications are targeted in this manuscript. These are characterized by high temperatures (up to 10 000 K) and relatively low pressures (200–20 000 Pa). *High-temperature effects* need to be accounted for: particles' vibrational and electronic energy modes get excited, and their non-linear contributions have to be included in the thermodynamic modeling. Furthermore, the high temperatures promote the dissociation/ionization of the molecules and a *multi-component, chemically reacting* treatment of the gas should be employed. Although free electrons can be present in the applications considered, we will restrict our attention to quasi-neutral (net charge of the mixture equal to zero) and unmagnetized plasmas. Both the low pressures and the low residence times of the flow prevent the equilibrium of macroscopic quantities. The degree of such non equilibrium is generally estimated by means of the Damköhler number, Da , which is defined

as the ratio between the flow residence time and the characteristic chemistry time. For relatively low flow residence times, the gas chemistry is outside the equilibrium condition ($Da \approx 1$) and *Finite-Rate Chemistry (FRC)* has to be included in the modeling of the gas phase (homogeneous chemistry). Shorter flow residence times cause also the internal energy modes to deviate from equilibrium conditions. Moderate non-equilibrium can be described by means of multi-temperature models, such as the two-temperature model proposed by Park [5]. This model assumes that the translational and rotational modes of heavy species are characterized by a temperature which may be different from the one of the vibrational and electronic modes and the translational mode of the electrons, and that the population of each mode follows a Maxwell-Boltzmann distribution. When non-equilibrium becomes stronger, and the energy population deviate from the Maxwell-Boltzmann distribution, more complex models, for instance the state-to-state one [75], should be employed. State-to-state models consider as a pseudo-species the particles characterized by a specific energy level, whose conservation is governed by rate processes. Thus, it requires the solution of a large number of continuity equations. In this manuscript, we will consider the flow to be chemically reacting and in thermal equilibrium.

As said before, molecules dissociate because of the high temperatures involved in the hypersonic flows. The resulting atoms diffuse towards the surface of the object, reacting with it through ablative and catalytic reactions. Thus, another feature that requires to be modeled is the *Gas-Surface Interaction (GSI)*. In this regards, balances of mass and energy on reacting surfaces were derived to describe the material surface response [76].

Lastly, turbulence, transition to turbulence, and radiative energy transfer might have a significant influence on the computation of the Quantities of Interest (QoIs) in hypersonic flows. Turbulence is not expected in the applications targeted in this manuscript as they are characterized by high temperatures (and thus high viscosity), resulting in low Reynolds number (ratio between inertial and viscous forces, Re). On the other hand, the temperatures are not sufficiently high to make radiation transport processes relevant. Hence, their formulation will not be presented.

To summarize, we will consider flows that are I) in the continuum regime ($Kn < 0.1$), II) in thermal equilibrium, III) in chemical non-equilibrium ($Da \approx 1$), IV) unmagnetized and quasi-neutral, V) laminar (low Re), and do not include V) radiative phenomena. In Section 2.2, we revise the multi-component, chemically reacting, laminar Navier-Stokes equations in thermal-equilibrium for dilute gases, presenting the state-of-the-art modeling for the chemistry, thermodynamics, transport, and GSI phenomena. Efficient algorithms for their solution are included in the Mutation⁺⁺ library, which is introduced in Section 2.3.

The above-mentioned physical models are generally affected by uncertainties either

in their form and in their parameters. Experimental data are essential to reduce such an uncertainty. We are primarily interested in the stochastic inference of GSI parameters. Section 2.4 provides the description of the experiments that we employed for the GSI characterization. The first two are a subsonic and a supersonic experimental campaign performed in the VKI Plasmatron facility. The last one, is a molecular beam-surface scattering experiment.

2.2 Governing equations

The applications targeted in this manuscript can be fully described by means of the multi-component chemical reacting Navier-Stokes equations [4, 5, 77, 78]. In conservation form, they are:

$$\frac{\partial \rho_i}{\partial t} + \nabla \cdot (\rho_i \mathbf{u} + \mathbf{j}_i) = \dot{\omega}_i, \quad \forall i \in [1, \dots, n_s], \quad (2.1)$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + p \bar{\mathbf{I}} - \bar{\boldsymbol{\tau}}) = \mathbf{0}, \quad (2.2)$$

$$\frac{\partial \rho E}{\partial t} + \nabla \cdot (\rho \mathbf{u} H - \bar{\boldsymbol{\tau}} \cdot \mathbf{u} + \mathbf{q}) = 0. \quad (2.3)$$

External forces contributions are considered negligible in the treated applications. In fact, gravity does not play a significant role, while Lorentz forces are null as we restrict our attention to quasi-neutral and unmagnetized plasma flows.

Equation (2.1) represents the *species mass conservation*: symbol \mathbf{u} is the mass-averaged (bulk) mixture velocity, ρ_i the partial density of species i , \mathbf{j}_i its diffusive mass flux, $\dot{\omega}_i$ its chemical production/destruction rate, and n_s the number of species in the mixture. The *total continuity equation* can be obtained by summing Equation (2.1) over all the species, it reads:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (2.4)$$

as the diffusive mass fluxes and the chemical production rates sum to zero for mass constraints [78]. Symbol $\rho = \sum_{i=1}^{n_s} \rho_i$ is the mixture density.

Equation (2.2) represents the *conservation of momentum*: symbol p stands for the thermodynamic pressure of the mixture, $\bar{\mathbf{I}}$ the unit tensor, and $\bar{\boldsymbol{\tau}}$ the viscous stress tensor. Under the Stokes assumption, it reads:

$$\bar{\boldsymbol{\tau}} = 2\mu D = \mu \left[\nabla \mathbf{u} + (\nabla \mathbf{u})^T - \frac{2}{3} \nabla \cdot \mathbf{u} \bar{\mathbf{I}} \right], \quad (2.5)$$

where μ is the viscosity of the mixture (whose closure is given in Section 2.2.3), and D is the deviatoric part of the strain rate (symmetric part of the the velocity

gradient).

Finally, Equation (2.3) represents the *conservation of energy*: symbol $E = e + \frac{1}{2}\mathbf{u} \cdot \mathbf{u}$ is the total energy, sum of the thermodynamic contribution (e) and the kinetic one ($\frac{1}{2}\mathbf{u} \cdot \mathbf{u}$), $H = E + p/\rho$ is the total enthalpy, and \mathbf{q} the total heat flux. Neglecting the radiative component, it is composed of both a conductive and a diffusive contribution:

$$\mathbf{q} = -\lambda\nabla T + \sum_{i=1}^{n_s} \mathbf{j}_i h_i, \quad (2.6)$$

where λ is the gas thermal conductivity (whose closure is given in Section 2.2.3), and h_i is the specific enthalpy of the species i .

Given the relatively low pressures involved in the applications treated in this manuscript, no dense effects are encountered and the ideal law for dilute gas applies:

$$p_i = R_i \rho_i T. \quad (2.7)$$

Symbol $R_i = R/M_i$ is the species i gas specific constant, p_i its partial pressure, and M_i its molar mass, while R is the universal gas constant. Dalton's law does also apply:

$$p = \sum_{i=1}^{n_s} p_i. \quad (2.8)$$

2.2.1 Homogeneous finite-rate chemistry

When advection/diffusion phenomena are slow enough, the gas has the time to relax toward its chemical equilibrium condition: its chemical composition depends only on the thermodynamic state of the mixture. Generally, hypersonic applications deviate from this condition and the gas is in chemical non-equilibrium. The non-equilibrium degree can be estimated by means of the Damköhler number [5]. As we saw in Section 2.1, it is defined as:


$$\text{Da} = \frac{t^{\text{flow}}}{t^{\text{chemistry}}}, \quad (2.9)$$

where t^{flow} is the flow residence time (which can be associated with advection phenomena in most of the flow, or diffusive ones inside the BL), and $t^{\text{chemistry}}$ is the characteristic chemistry time. A very small Da number ($0 < \text{Da} \ll 1$) suggests *frozen chemistry* (no reaction has the time to occur), while a high number ($\text{Da} \gg 1$) a *chemistry in equilibrium*. In the middle ranges ($\text{Da} \approx 1$), we deal with *chemical non-equilibrium* [5].

Example Let us consider an adiabatic box containing oxygen, which is initialized with a temperature of $T_0 = 10\,000$ K and mass fractions equal to $Y_{\text{O}_2} = 1.0$ and $Y_{\text{O}} = 0.0$. The evolution in time of the mass fractions can be described by the conservation of mass and energy:

$$\frac{\partial \rho_i}{\partial t} = \dot{\omega}_i, \text{ with } i \in [\text{O}, \text{O}_2]; \quad \frac{\partial \rho E}{\partial t} = 0,$$

and by the dissociative reaction:



According to the law of mass action, the molecular oxygen chemical production rate reads:


$$\frac{\dot{\omega}_{\text{O}_2}}{M_{\text{O}_2}} = \sum_{m \in [\text{O}, \text{O}_2]} \left(-k_{f,m} \frac{\rho_{\text{O}_2}}{M_{\text{O}_2}} \frac{\rho_m}{M_m} + k_{b,m} \left(\frac{\rho_{\text{O}}}{M_{\text{O}}} \right)^2 \frac{\rho_m}{M_m} \right),$$

where M indicates a third body in the reaction, and $k_{f,m}$ and $k_{b,m}$ are, respectively, the forward and the backward reaction rate coefficients.

If we consider solely the forward direction of the reaction, and only the molecular oxygen as third body, the characteristic chemistry time is:

$$t^{\text{chemistry}} = \frac{\rho_{\text{O}_2}}{\dot{\omega}_{\text{O}_2}} \approx \frac{M_{\text{O}_2}}{k_{f,\text{O}_2} \rho_{\text{O}_2}} = \frac{RT}{k_{f,\text{O}_2}(T) p_{\text{O}_2}}.$$

Hence, we can expect $t^{\text{chemistry}}$ to be inversely proportional to the gas pressure. To highlight such a dependence, let us consider four different initial pressures of the box, specifically $p_0 = [10, 1, 0.1, 0.01]$ atm. The system is left free to time march towards equilibrium. The evolution of the molecular oxygen mass fraction versus the time is plotted in the figure below.




Given the high temperature, the oxygen dissociates with a speed that depends on the gas pressure. At 10 atm the chemistry takes 1×10^{-6} s to relax. As we decrease the pressure, fewer and fewer particles collide, and the chemistry slows: for a pressure of 0.1 atm, typical of entry applications, the characteristic chemistry time is 1×10^{-3} s.

Let us also assume a characteristic flow time representative of entry flows. When associated to advection phenomena, it can be computed as the inverse of the stagnation-point velocity gradient $t^{\text{flow}} = (du/dx)^{-1} \approx R/u_\infty$, where R is the effective radius of a hypothetical entry object and u_∞ its velocity. If we take $R = 2$ m and $u_\infty = 12\,000$ m/s, we obtain $t^{\text{flow}} = 1/6000$ s.

At this point, we can compute the Damköhler number. At pressures representative of entry flows, we had $t^{\text{chemistry}} = 1 \times 10^{-3}$ s, leading to a $\text{Da} = 1/6$, in the chemical non-equilibrium regime. By contrast, with a characteristic chemistry time of 1×10^{-6} s, relative to pressures much higher than those involved in the targeted applications, we would have $\text{Da} \approx 10^3$, in the equilibrium range.

The applications treated in this manuscript are characterized by low pressures and high speeds. Both contributes to having a Da ranging from non-equilibrium to frozen values. In these cases, homogeneous FRC should be accounted for.

All elementary reactions are reversible [78], and they can be written in compact form as:



where $\nu_{i,r}^f$ and $\nu_{i,r}^b$ are the forward and the backward stoichiometric reaction coefficients for the species i in the reaction r , and \mathcal{S}_i is the species with index i . The Maxwellian net chemical production rate of each species, $\dot{\omega}_i$, can be computed by means of the law of mass action [78]:

$$\dot{\omega}_i = M_i \sum_{r=1}^{n_r} (\nu_{i,r}^b - \nu_{i,r}^f) \left[k_r^f \prod_j \left(\frac{\rho_j}{M_j} \right)^{\nu_{j,r}^f} - k_r^b \prod_j \left(\frac{\rho_j}{M_j} \right)^{\nu_{j,r}^b} \right], \quad (2.11)$$

where n_r is the number of reactions considered, and k_r^f and k_r^b are, respectively, the forward and the backward reaction rate coefficients.

The forward reaction rate coefficients are generally given in the empirical Arrhenius form:

$$k_r^f(T) = A_r T^{n_r} \exp\left(-\frac{T_{a,r}}{T}\right), \quad (2.12)$$

where the pre-exponential, A_r , the exponential, n_r , and the activation energy, $T_{a,r}$, are usually calibrated on shock-tube experimental data. Example of rates coefficients for hypersonic applications can be found in Park et al. [5, 79], Olynick

et al. [80] (for Earth atmosphere) and Gökçen [81] (for Titan atmosphere).

In order to satisfy the chemical equilibrium condition and respect the second law of thermodynamics [78], the backward rate coefficients are obtained by means of the equilibrium rate:

$$k_r^b = \frac{k_r^f}{k_r^{\text{eq}}}, \quad (2.13)$$

defined as:

$$\ln k_r^{\text{eq}} = - \sum_{i=1}^{n_s} (\nu_{i,r}^b - \nu_{i,r}^f) \left[\frac{g_i(p_{\text{ref}}, T)}{R_i T} - \ln \left(\frac{p_{\text{ref}}}{RT} \right) \right], \quad (2.14)$$

where g_i is the Gibbs free energy of the species i and p_{ref} a reference pressure, which is arbitrary. Note that the equilibrium rate coefficient depends on the only temperature. Alternatively, the equilibrium rate coefficients can be obtained using empirical polynomials [79].

2.2.2 Multi-component thermodynamics

The thermodynamic properties of a mixture can be obtained by averaging the ones of the species composing it:

$$e = \sum_{i=1}^{n_s} Y_i e_i, \quad (2.15)$$

$$h = e + \frac{p}{\rho} = \sum_{i=1}^{n_s} Y_i h_i, \quad (2.16)$$

$$c_{p,\text{fr}} = \sum_{i=1}^{n_s} Y_i c_{p,i} = \sum_{i=1}^{n_s} Y_i \left(\frac{\partial h_i}{\partial T} \right)_p, \quad (2.17)$$

$$s = \sum_{i=1}^{n_s} Y_i s_i - \frac{nk_B}{\rho} \sum_{i=1}^{n_s} X_i \ln X_i, \quad (2.18)$$

$$g = \sum_{i=1}^{n_s} Y_i g_i + \frac{nk_B T}{\rho} \sum_{i=1}^{n_s} X_i \ln X_i, \quad (2.19)$$

with:

$$g_i = h_i - T s_i, \quad (2.20)$$

where $Y_i = \rho_i/\rho$ is the mass fraction of species i and $X_i = p_i/p$ its molar fraction. Symbols e_i , h_i , s_i , and $c_{p,i}$ stand respectively for the species' internal energy, enthalpy and entropy, and specific heat at constant pressure, while k_B is the Boltzmann constant. The second term on the Right Hand Side (RHS) of Equation (2.18) represents the mixture's mixing entropy. Symbol $c_{p,\text{fr}}$ is the mixture

frozen specific heat at constant pressure.

From a microscopic point of view, particles can store energy in different modes [4, 5]. They are:

1. *Translational mode* (e^T): kinetic energy associated with the motion of their center of gravity.
2. *Rotational mode* (e^R): kinetic energy associated with the rotation around their axis of inertia.
3. *Vibrational mode* (e^V): sum of potential and kinetic energies associated with their vibration.
4. *Electronic mode* (e^E): sum of a potential and kinetic contributions associated with the rotation of the electrons in a specific shell.

The translational mode has three thermal degrees of freedom, as the particles can move in three independent directions, while the rotational and the vibrational degrees of freedom depend on the structure of the particle.

Assuming that these modes are independent,¹ the energy of a molecule is the sum of the above-mentioned contributions:

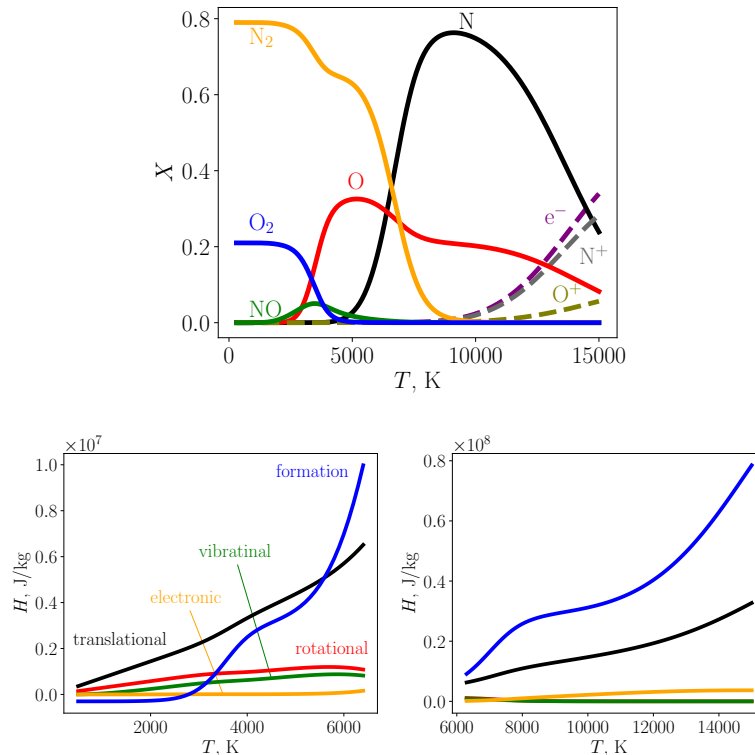
$$e_i = e_i^T + e_i^R + e_i^V + e_i^E + e_i^0, \quad (2.21)$$

and of the formation energy, e_i^0 , which accounts for the chemical energy. It is defined as the enthalpy adsorbed or released upon the formation of one mole of substance from its constituent elements, at standard conditions. Since in a reactor it is only possible to measure this difference in energy, the formation enthalpy cannot be determined in absolute way for each species. It is a convention to assume it to be zero for those elements in their standard states and to attribute the whole difference as formation energy to the other species.

Contrary to molecules, the energy of atoms does not include the vibrational and rotational contribution.

¹For gas with weakly interacting particles (that react only upon collision), it is safe to separate the translational by the internal (roto-vibronic) energy: $e_i = e_i^T + e_i^I$ [82]. Additionally, the use of the Born-Oppenheimer approximation allows for splitting the latter term in an electronic and a roto-vibrational contribution: $e_i^I = e_i^E + e_i^{RV}$. The rotational and vibrational mode are closely coupled as the rotation induces centrifugal forces that, in turn, affects the vibration, while the vibration changes the moment of inertia of the molecule, affecting the rotation [82]. Accounting for such a coupling is crucial to correctly model the thermodynamics at high temperatures, for example when dealing with radiation. However, the general practice in CFD simulations is to consider independent also these two modes because I) it is more efficient, and II) at those temperatures where the roto-vibrational coupling is important, if the gas is not too far from the equilibrium condition, the mixture is fully dissociated [83]. The same holds true for the computation of the entropy and Gibbs free energy.

Example Let us consider an air mixture of 11 species ($S = [\text{N}_2, \text{O}_2, \text{NO}, \text{N}, \text{O}, \text{N}_2^+, \text{O}_2^+, \text{NO}^+, \text{N}^+, \text{O}^+, \text{e}^-]$). The evolution of the equilibrium composition of the main species at atmospheric pressure is shown below against the gas temperature. It is followed by the contributions of each energy mode to the mixture enthalpy.



At room temperature, only translational and rotational modes are activated. Since they are linear in T the gas is *calorically perfect*. For temperatures above 1000 K the vibrational mode starts getting excited and, as its contribution is not linear in T , the gas is no longer calorically perfect. *High-temperature effects* have to be taken into account and the gas is said to be *thermally perfect*. Further increase in temperature drives molecule dissociation and the formation enthalpy sharply increases. Around 6000 K, the electronic mode starts to get excited and, at higher temperature, drives the ionization of the mixture. Around 9000 K, free electrons and ions appear in the mixture. By contrast, vibrational and rotational energy decreases at temperatures higher than 5000 K, as, because of the dissociation, the number of molecules that can vibrate and rotate decreases too.

For a thermally perfect gas, the thermodynamic properties of each species can be

computed either by means of the Rigid Rotor Harmonic Oscillator (RRHO) model or employing the NASA polynomials.

Rigid rotor harmonic oscillator

The RRHO is a theoretical model for computing thermodynamic properties. It assumes that I) the molecules rotate around their center of mass as a rigid body (the intermolecular distance does not change), II) the vibration can be described using a simple harmonic potential energy function (anharmonicity is neglected), and III) the rotational and vibrational modes are decoupled. Despite these strong assumptions, the RRHO is widely used in the hypersonic community as the mixture is dissociated at those temperatures where they are inadequate to describe the molecules [83]. The RRHO formulation can be found in classical textbooks [4, 82], following, we briefly present the main concepts.

As we saw before, particles can store energy in different modes. According to quantum physics, this energy is quantized: it can only have discrete values. However, we will present a semi-classical formulation, since the translation modes' energy levels are so packed that can be safely considered as a continuum. For the same reason, also the rotational mode is described with a continuous representation. When the mixture is in thermal equilibrium, assuming the gas to be non-degenerate (assumption valid for temperature higher than 5 K [4]), the population of the species i living in the mode energy level j (with energy $\varepsilon_i^{j,m}$) can be obtained by the Boltzmann distribution:

$$\frac{n_i^j}{n_i} = \frac{g_i^j \exp\left(-\sum_{m \in \mathcal{M}} \frac{\varepsilon_i^{j,m}}{k_B T^m}\right)}{Q_i}, \quad (2.22)$$

where n_i^j is the number density of species i with energy level j , such that the total number density of the i species is respected:

$$\sum_j n_i^j = n_i, \quad (2.23)$$

g_i^j is product of the modes degeneracy:

$$g_i^j = g_i^{j,T} g_i^{j,R} g_i^{j,V} g_i^{j,E}, \quad (2.24)$$

and the superscript m indicates the index for each energy modes, with $\mathcal{M} = [T, R, V, E]$. The quantity Q_i is the total partition function:

$$Q_i = Q_i^T Q_i^R Q_i^V Q_i^E, \quad (2.25)$$

where each component reads:

$$Q_i^m = \sum_j g_i^j \exp\left(-\frac{\varepsilon_i^{j,m}}{k_B T^m}\right). \quad (2.26)$$

The evaluation of the partition function allows for computing all the macroscopic thermodynamics properties, such as the internal energy, the pressure, the enthalpy, and the entropy:

$$e_i^m = R_i (T^m)^2 \left(\frac{\partial}{\partial T^m} \ln Q_i^m \right)_V, \quad (2.27)$$

$$p_i^m = n_i k_B T^m \left(\frac{\partial}{\partial V} \ln Q_i^m \right)_T, \quad (2.28)$$

$$h_i^m = e_i^m + \frac{p_i^m}{\rho_i}. \quad (2.29)$$

$$s_i^m = R_i \left(T^m \left(\frac{\partial}{\partial T^m} \ln Q_i^m \right)_V + 1 \right) + R_i \ln \left(\frac{Q_i^m}{n_i} \right) \quad (2.30)$$

For the translational mode, this leads to:

$$e_i^T = \frac{3}{2} T^T R_i, \quad (2.31)$$

$$p_i = n_i k_B T^T. \quad (2.32)$$

$$h_i^T = \frac{5}{2} T^T R_i, \quad (2.33)$$

$$s_i^T = \frac{h_i^T}{T^T} + R_i \ln \left(\frac{k_B T^T}{p_i} \left(\frac{2\pi m_i k_B T^T}{h^2} \right)^{3/2} \right), \quad (2.34)$$

Only the translational mode contributes to the pressure, recovering the ideal gas law.

For the rotational mode:

$$h_i^R = e_i^R = \frac{\mathcal{L}}{2} T^R R_i, \quad (2.35)$$

$$s_i^R = \frac{h_i^R}{T^R} + R_i \left(\frac{\mathcal{L}}{2} \ln \left(\frac{T^R}{\theta_i^R} \right) - \ln \sigma_i \right) \quad (2.36)$$

where \mathcal{L} is equal to 2 if the molecule is linear, 3 otherwise, θ_i^R is the characteristic rotational temperature, and σ_i is the steric factor, equal to 2 for symmetric molecules, and to 1 otherwise.

For the vibrational mode:

$$h_i^V = e_i^V = R_i \sum_v \frac{\theta_{i,v}^V}{\exp(\theta_{i,v}^V/T^V) - 1}, \quad (2.37)$$

$$s_i^V = \frac{h_i^V}{T^V} - R_i \sum_v \ln \left(1 - \exp \left(-\frac{\theta_{i,v}^V}{T^V} \right) \right), \quad (2.38)$$

Finally, for the electronic mode:

$$h_i^E = e_i^E = R_i \frac{\sum_e g_{i,e}^E \theta_{i,e}^E \exp(-\theta_{i,e}^E/T^E)}{Q_i^E}. \quad (2.39)$$

$$s_i^E = \frac{h_i^E}{T^E} + R_i \ln Q_i^E. \quad (2.40)$$

The quantities $\theta_{i,v}^V$ and $\theta_{i,e}^E$ are the characteristic vibrational and electronic temperatures. The sum in the Equation (2.39) is divergent and a cut-off value, e^{\max} , is usually imposed.

The energy and the enthalpy of the free electrons can be computed using, respectively, Equation (2.31) and (2.33) with the free electrons temperature, T^e . Their entropy is:

$$s^e = \frac{h^e}{T^e} + R_e \ln \left(\frac{k_B T^e}{p_e} \left(\frac{2\pi m_e k_B T^e}{h^2} \right)^{3/2} \right) + R_e \ln 2, \quad (2.41)$$

where the last term in the RHS is the spin contribution to the entropy of the free electrons.

In order to close the above equations, one need to specificity particles' degeneracy and characteristic temperatures. They can be found in thermodynamic tables [84, 85]. Furthermore, in this thesis we will assume that the thermal equilibrium condition holds in all the applications, such that:

$$T^T = T^R = T^V = T^E = T^e = T. \quad (2.42)$$

NASA polynomials

Alternatively, NASA polynomials can be used to compute the thermodynamic properties [86, 87]. They were obtained in the 200-20 000 K temperature range, and accounts for roto-vibrational coupling and anharmonicity. A seven-coefficients polynomial was chosen to fit the dimensionless heat capacity of each species:

$$\frac{c_{p,i}}{R_i} = a_{0i} T^{-2} + \frac{a_{1i}}{T} + a_{2i} + a_{3i} T + a_{4i} T^2 + a_{5i} T^3 + a_{6i} T^4, \quad (2.43)$$

The enthalpy can be obtained by integrating $c_{p,i}$ in T , while the entropy by integrating $c_{p,i}/T$ in T . They read:

$$\frac{h_i}{R_i T} = -a_{0i} T^{-2} + a_{1i} \frac{\ln T}{T} + a_{2i} + a_{3i} \frac{T}{2} + a_{4i} \frac{T^2}{3} + a_{5i} \frac{T^3}{4} + a_{6i} \frac{T^4}{5} + \frac{a_{7i}}{T}, \quad (2.44)$$

$$\frac{s_i}{R_i} = -a_{0i} \frac{T^{-2}}{2} - \frac{a_{1i}}{T} + a_{2i} \ln T + a_{3i} T + a_{4i} \frac{T^2}{2} + a_{5i} \frac{T^3}{3} + a_{6i} \frac{T^4}{4} + a_{8i}. \quad (2.45)$$

2.2.3 Multi-component transport

Particles move in the flow with a peculiar velocity, which, in general, is different both in module and direction from the one of the mixture. This allows them to diffuse, transporting momentum and energy. The kinetic theory allows for expressing transport fluxes as products between macroscopic properties gradients (i.e. ∇x_i , $\nabla \mathbf{u}$, and ∇T) and transport coefficients. The latter can be rigorously obtained by the first-order Chapman-Enskog expansion of the Boltzmann equation [78, 88].

The viscosity and the translational thermal conductivity of single species respectively are:

$$\mu_i = \frac{5}{16} \frac{\sqrt{\pi m_i k_B T}}{\bar{Q}_{ii}^{(2,2)}}, \quad (2.46)$$

$$\lambda_i^T = \frac{15}{4} R_i \mu_i, \quad (2.47)$$

while, the binary diffusion coefficient of species i diffusing in species j is:

$$D_{ij} = \frac{3}{16n\bar{Q}_{ij}^{(1,1)}} \sqrt{\frac{2\pi k_B T(m_i + m_j)}{m_i m_j}}, \quad (2.48)$$

where $\bar{Q}_{ij}^{(l,s)}$ is the reduced collision integral, which links macroscopic properties of the gas to the microscopic dynamics of the collisions [78, 89]. It is the average, over all the relative energies, of the relevant cross sections in the binary collision between the i and j particles. The quantities (l, s) represent the order of the Sonine polynomial employed in the spectral method used to solve the equation arising from the first-order Chapman-Enskog expansion.

Mass diffusion

Accurate diffusion velocities (\mathbf{V}_i) can be computed by solving the *Stefan-Maxwell* system [78, 89, 90]. In the absence of an electric field, and neglecting the pressure and the thermal driving forces, it reads:

$$\sum_{i=1}^{n_s} G_{ij}^V \mathbf{V}_i = -\frac{p}{nk_B T} \nabla X_i, \quad (2.49)$$

where G^V is the diffusion matrix, whose definition is given in Appendix B. The system is closed with the mass constraint:

$$\sum_{i=1}^{n_s} Y_i \mathbf{V}_i = 0. \quad (2.50)$$

The diffusion matrix is singular. Following Giovangigli [78], Magin and Degrez obtained a non-singular form by regularizing the diffusion matrix as $\hat{\mathbf{G}}^V = \mathbf{G}^V + \alpha \mathbf{y} \otimes \mathbf{y}$, where α is a constant of the same order of the matrix, for example $1/\max(D_{ij})$ [89]. Finally, the diffusive fluxes are obtained as:

$$\mathbf{J}_i = \rho_i \mathbf{V}_i. \quad (2.51)$$

Alternatively, Fick's law can be employed by using effective binary diffusion coefficients (diffusion of species i in the pseudo-species composed by the rest of the mixture, $S_{\setminus i}$). It is defined as:

$$D_i = \frac{1 - X_i}{\sum_{j \neq i} \frac{X_j}{D_{ij}}}. \quad (2.52)$$

Enforcing the mass constraint on the diffusive fluxes leads to the *Self-Consistent Effective Binary Diffusion (SCEBD)* formulation proposed by Ramshaw [91, 92]:

$$\mathbf{J}_i = -\frac{pM_i}{RT} D_i \nabla X_i + Y_i \sum_{l=1}^{n_s} \frac{pM_l}{RT} D_l \nabla X_l. \quad (2.53)$$

Momentum and energy transport

The multi-component mixture coefficients, $\mathcal{T} = [\mu, \lambda^T]$, are obtained by the average:

$$\mathcal{T} = \sum_{i=1}^{n_s} \alpha_i^{\mathcal{T}} X_i, \quad (2.54)$$

where the coefficients $\alpha_i^{\mathcal{T}}$ are solution of the transport linear system:

$$\mathbf{G}^{\mathcal{T}} \boldsymbol{\alpha}^{\mathcal{T}} = [X_1, \dots, X_{n_s}]^T. \quad (2.55)$$

Symbols \mathbf{G}^{μ} and \mathbf{G}^{λ^T} are the transport matrices, whose definition can be found in Appendix B.

The total thermal conductivity is then computed as:

$$\lambda = \lambda^T + \lambda^I + \lambda^e, \quad (2.56)$$

where the internal thermal conductivity is usually obtained through the Eucken correction:

$$\lambda^I = \lambda^R + \lambda^V + \lambda^E = \sum_{m \in \mathcal{M}_{\setminus T}} \lambda^m = \sum_{m \in \mathcal{M}_{\setminus T}} \sum_{i=1}^{n_s} \frac{\rho_i c_{p,i}^m}{\sum_{j=1}^{n_s} \frac{x_j}{D_{ij}}}, \quad (2.57)$$

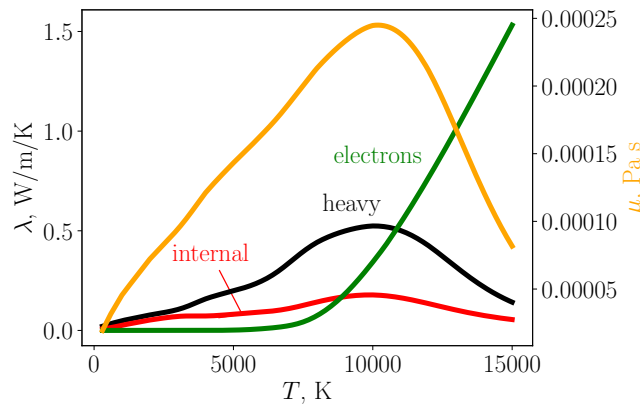
being $\mathcal{M}_{\setminus T} = [R, V, E]$, and the thermal conductivity of the free electrons reads:

$$\lambda^e = \frac{X_e^2 \Lambda_{ee}^{22}}{\Lambda_{ee}^{11} \Lambda_{ee}^{22} - (\Lambda_{ee}^{12})^2}. \quad (2.58)$$

The expressions for the $\Lambda_{ee}^{(\cdot)}$ terms can be found in [89, 90, 93]. On the other hand, due to their lower mass, electrons do not contribute to the viscosity [93].

Alternatively, *approximate transport coefficients* can be obtained by employing mixture rules [94–96].

Example Let us consider an air mixture of 11 species in chemical equilibrium ($S = [N_2, O_2, NO, N, O, N^+, O^+, N_2^+, O_2^+, NO^+, e^-]$) in the temperature range 200 - 15 000 K and atmospheric pressure. The temperature-dependent transport coefficients are plotted in the figure below.




The viscosity first increases with the temperature, in accordance with Equation (2.46), which shows a square root dependence on T . For higher temperatures, the mixture ionizes, resulting in an increase in the collision cross section and, consequently, in the collision integrals. For temperatures above 10 000 K the term $\bar{Q}_{ii}^{(2,2)}$ dominates on \sqrt{T} and the viscosity decreases.

At room temperature, only the translational and the rotational degree of freedom are excited, so only these components of the energy can be transported. Being the translational thermal conductivity linearly depended on the viscosity, Equation (2.47), it follows the same trend. The vibrational and electronic energies increase with the temperature. Consequently, their thermal conductivity component increases with T , until a point in which they start decreasing, first because of dissociation, then because of ionization. At this point, the mixture's thermal conductivity greatly increases because of the contribution of the free electrons. Note that the diffusive component of the heat flux


$(\sum_{i=1}^{n_s} j_i h_i)$ is an important contribution, but was not included in the example.

2.2.4 Gas-surface interaction

In all the applications targeted in this manuscript, a reacting flow impinges an object whose surface undergoes either catalytic or ablative reactions. In the first case, the surface (s) acts as a catalyzer for the recombination reaction of two gas-phase atoms, for example:



In the second case, a surface atom (\cdot_b) reacts with a gas-phase one. An example is the oxidation of a carbon surface:



Both reactions are exothermic and increase the surface heat flux. Additionally, the oxidation reaction ablates the material and generates a blowing velocity at the surface. Furthermore, they both affect the BL chemical composition.

To rigorously capture the physics of the material response, CFD and material computations should be coupled as in [97, 98]. This approach can be computationally expensive. Under some assumptions, one can model the material response only in terms of surface mass/energy balances on its surface [76, 99].

Gas-surface balances can be derived for any conservative variable, U , by means of the Reynolds transport theorem [100]. On the infinitesimal volume containing both the gas and the surface, under steady-state assumption ($\partial U / \partial t = 0$), the balance reads:

$$[\mathbf{F}_g - \mathbf{F}_s] \cdot \mathbf{n} = S_s, \quad (2.59)$$

where \mathbf{F}_g and \mathbf{F}_s respectively are the fluxes in/from the gas- and the solid-interface, S_s is the surface source terms, and \mathbf{n} is the normal to the surface, pointing the gas-phase.

Following, we will present the mass and energy balances. No balance is provided for the momentum, as its closure is given by the no-slip condition. Furthermore, phenomena like pyrolysis, spallation, or mechanical removal will not be considered.




Figure 2.1: Sketch of the surface mass balance.

When the surface temperature is known, one only needs to solve a *Surface Mass Balance (SMB)*, which is sketched in Figure 2.1: the surface chemistry generates a species gradient at the surface, driving diffusion. Assuming the recession velocity (\mathbf{u}_r) to be negligible compared to the blowing or advection one (\mathbf{u}_g), the SMB for each species i reads:

$$\dot{\omega}_i = (\rho_i \mathbf{u}_g + \mathbf{j}_i) \cdot \mathbf{n}, \quad \forall i \in [1, \dots, n_s]. \quad (2.60)$$

When reactions involve material removal, the material recedes, and the products are advected out from the surface (blowing effect). The blowing velocity is defined as:

$$\mathbf{u}_g = \frac{\dot{m}}{\rho}, \quad (2.61)$$

while the recession one as:

$$\mathbf{u}_r = \frac{\dot{m}}{\rho_s}, \quad (2.62)$$

where \dot{m} is the blowing mass rate:

$$\dot{m} = \sum_{i=1}^{n_s} \dot{\omega}_i, \quad (2.63)$$

which is greater than zero for ablative materials, $\rho = \sum_{i=1}^{n_s} \rho_i$, and ρ_s is the density of the material. These two velocities are null in the case of catalytic surfaces as $\sum_{i=1}^{n_s} \dot{\omega}_i = 0$.

Also the GSI chemistry is a rate process, thus, it can be characterized by a Damköhler number. It can be defined as:

$$\text{Da}^w = \frac{\omega_i}{\rho_i} \frac{\delta^{\text{BL}}}{D_{12}}, \quad (2.64)$$

where δ^{BL} is the BL thickness. A value of $\text{Da}^w \rightarrow 0$ suggests a *reaction limited regime* while $\text{Da}^w \rightarrow \infty$ a *diffusion limited one* [101].

When the surface temperature is not known, a coupled *Surface Energy Balance (SEB)* has to be solved. It is sketched in Figure 2.2.




Figure 2.2: Sketch of the surface energy balance.

The conductive heat flux toward the material is either neglected (for low thermal conductivity materials), or approximated by the steady-state ablation assumption [102]. The SEB balance reads:

$$\rho u_g h - \lambda \frac{\partial T}{\partial n} + \sum_{i=1}^{n_s} j_i h_i = S_{\text{rad}}, \quad (2.65)$$

where $S_{\text{rad}} = \sigma \epsilon T^4$ is the radiative energy flux, σ is the Stefan–Boltzmann constant, and ϵ the body emissivity.

The surface chemical source terms, $\dot{\omega}_i$, can be computed either by means of phenomenological approaches or by using FRC models.

Phenomenological approaches

Phenomenological approaches [103] are extensively used in the hypersonic community to model both catalysis and ablation [21, 99, 104]. They are rather simple methods as they only require the definition of the probability that a macroscopic reaction takes place. For this reason, they can be calibrated by means of macroscopic experiments, for example using Inductively Coupled Plasma (ICP) facilities or arc-jets data.

The probability or efficiency of the reaction is defined as:

$$\gamma_i^r = \frac{\mathcal{N}_i^r}{\mathcal{N}_i}, \quad (2.66)$$

where \mathcal{N}_i^r is the number flux of species i subject to the reaction r and $\mathcal{N}_i = n_i \sqrt{k_B T_s / (2\pi m_i)}$, the number flux of species i impinging the surface. A fully catalytic behavior is obtained when the probability approaches unity; whereas a non-catalytic behavior is reached when is equal to zero, in the between the surface is partially catalytic.

In this case, the chemical production rate reads:

$$\dot{\omega}_i = m_i \gamma_i^r \mathcal{N}_i. \quad (2.67)$$


Finite-rate chemistry models

Alternative formulations are FRC models [1, 22, 24, 105, 106]. They describe each elementary reaction occurring on the surface (adsorption/desorption, oxidation, nitridation, recombination, etc.). They have the advantage of capturing the pressure dependence of the chemical reactions and the surface coverage.


In this case, the chemical production rate reads:

$$\dot{\omega}_i = M_i \sum_{r=1}^{n_r} \nu_{i,r} k_r \prod_{j \in S} \left(\frac{\rho_j}{M_j} \right)^{\nu_{j,r}}. \quad (2.68)$$

The form of the rate coefficient, k_r , depends on the type of reaction [1, 106]. Let us consider a surface characterized by a total active site density B . Gas atoms can be absorbed on a surface free site:



Similarly, absorbed atoms can desorb from the surface:



as sketched in Figure 2.3. Symbol G stands for a gas-phase atom, s is a free spot on the surface, and the subscript \cdot_s indicates that the atom has been absorbed by the surface.




Figure 2.3: Sketch of atomic adsorption/desorption.

Adsorption rates are in the form:

$$k_{\text{abs},i} = \frac{s}{B} V_{T,i} \exp\left(-\frac{T_{\text{abs},i}}{T}\right), \quad (2.69)$$

where s is the selectivity of the adsorption, which lies between 0 and 1, $V_{T,i} = \sqrt{\frac{k_B T}{2\pi m_i}}$ represents a quarter of the thermal velocity of the species i , and $T_{\text{abs},i}$ indicates the activation energy of the absorption process.

According to the simple transition state theory assumption, desorption rates read:

$$k_{\text{des},i} = \frac{2\pi m_i k_B T^2}{N_A B h^3} \exp\left(-\frac{T_{\text{des},i}}{T}\right), \quad (2.70)$$

where N_A is the Avogadro's constant, h the Plank's one, and $T_{\text{des},i}$ stands for the activation energy of the desorption process.

Recombination reactions can be provoked by a gas-phase atom, which, striking a surface, recombines with an absorbed atom (Eley-Rideal (ER) reactions, sketched in Figure 2.4). Alternatively, two absorbed atoms can recombine. If they interact with enough energy, the molecule can escape the potential barrier of the surface (Langmuir-Hinshelwood (LH) reactions, sketched in Figure 2.5).



Figure 2.4: Sketch of ER reaction.

In the first case, the reaction rate is in the form:

$$k_{er,i} = \frac{1}{B} V_{T,i} \gamma_{er,i} \exp\left(-\frac{T_{er,i}}{T}\right), \quad (2.71)$$

where $\gamma_{er,i}$ is the probability of the ER reaction, with value between 0 and 1 and $T_{er,i}$ represents the activation energy of the reaction.



Figure 2.5: Sketch of LH reaction.

In the case of a LH reaction, the rate coefficient is:

$$k_{lh,i} = \sqrt{\frac{N_A}{B}} V_{T,2D,i} \gamma_{lh,i} \exp\left(-\frac{T_{lh,i}}{T}\right), \quad (2.72)$$

where $V_{T,2D,i} = \sqrt{\frac{\pi k_B T}{2m_i}}$ is the two-dimensional thermal velocity, and $T_{lh,i}$ indicates the activation energy of the reaction.

The same processes can yield to surface ablation when one of the reactants is a carbon surface atom. When the reaction is gas-dependent, the reaction rate reads:

$$k_{gd,i} = V_{T,i} \gamma_{gd}, \quad (2.73)$$

being γ_{gd} the probability of the reaction to take place.

When the reaction is gas-independent, the rate $k_{gi,i}$ has an Arrhenius-like formulation:

$$k_{gi,i} = A_{gi,i} T^{m_{gi,i}} \exp\left(-\frac{T_{gi,i}}{T}\right). \quad (2.74)$$

During the reactions, the total number of active sites is conserved, such that:

$$B = [G_s] + [s]. \quad (2.75)$$

Furthermore, in practical applications, one is not interested in the reactions transient and the steady-state solution is usually sought:

$$\frac{d}{dt}[s] = 0. \quad (2.76)$$

Air carbon ablation model

The Air Carbon Ablation (ACA) model proposed by Prata et al. [1] is a state-of-the-art FRC model. It comprises 20 elementary reactions occurring between a carbon surface and an air gas. As it will be discussed in Chapter 6, we used the subset of reactions involving atomic nitrogen to model the ablative behavior of a carbon surface exposed to a nitrogen plasma. Such a subset, as well as the corresponding rates, is reported in Table 2.1.

Index	Reaction	Reaction Rate	Rate Coefficient
1	$N + s \longrightarrow N_s$	$k_1[N][s]$	$\frac{V_{TN}}{B} \exp \frac{-2500}{T_s}$
2	$N_s \longrightarrow N + s$	$k_2[N_s]$	$\frac{2\pi m_N k_b^2 T_s^2}{N_A B h^3} \exp \frac{-73971}{T_s}$
3	$N + N_s + C_b \longrightarrow CN + N + s$	$k_3[N][N_s]$	$\frac{V_{TN}}{B} 1.5 \exp \frac{-7000}{T_s}$
4	$N + N_s \longrightarrow N_2 + s$	$k_4[N][N_s]$	$\frac{V_{TN}}{B} 0.5 \exp \frac{-2000}{T_s}$
5	$N_s + N_s \longrightarrow N_2 + 2s$	$k_5[N_s][N_s]$	$\sqrt{\frac{N_A}{B}} V_{T,2D,N} 0.1 \exp \frac{-21000}{T_s}$
6	$N_s + C_b \longrightarrow CN + s$	$k_6[N_s]$	$10^8 \exp \frac{-20676}{T_s}$

Table 2.1: Subset of the FRC reactions involving nitrogen in the ACA model.

Figure 2.6 shows the efficiencies of nitridation and the recombination reactions predicted by the ACA model, both at low and high pressure. Such predictions were compared both to high-pressure² and to low-pressure³ experimental data to infer the rate coefficients.

The model efficiencies can be computed as follows. According to the law of mass action, the set of surface production rates read:

$$\frac{d}{dt}[N] = -k_1[N][s] + k_2[N_s] - k_4[N][N_s], \quad (2.77)$$

$$\frac{d}{dt}[N_2] = k_4[N][N_s] + k_5[N_s]^2, \quad (2.78)$$

$$\frac{d}{dt}[CN] = k_3[N][N_s] + k_6[N_s], \quad (2.79)$$

$$\frac{d}{dt}[N_s] = k_1[N][s] - k_2[N_s] - k_3[N][N_s] - k_4[N][N_s] - 2k_5[N_s]^2 - k_6[N_s]. \quad (2.80)$$

Such a system is closed with the conservation of active sites, Equation (2.75), and the steady-state assumption, Equation (2.76). Further inputs to the model are the surface temperature, and the gas pressure and temperature, used to compute the

²ICP experiments from Lutz [107] and from Helber et al. [3]; the latter is introduced in Section 2.4.1.

³Molecular beam-surface scattering experiments from Murray et al. [108], described in Section 2.4.2.

concentration of the atomic nitrogen. For the high-pressure predictions, the gas pressure was set equal to 1600 Pa and the gas temperature equal to the surface one. For the low-pressure predictions, these two values were set to produce an equivalent flux close to the one experimentally observed ($p_{\text{beam}} = 0.024$ Pa and $T_{\text{beam}} = 1000$ K). Finally, the efficiencies of the macroscopic recombination ($\text{N} + \text{N} \rightarrow \text{N}_2$) and nitridation ($\text{N} + \text{C}_b \rightarrow \text{CN}$) reactions can be computed as:

$$\gamma_{\text{N}_2} = \frac{2(d[\text{N}_2]/dt)}{\mathcal{N}_{\text{N}}}, \quad (2.81)$$

$$\gamma_{\text{CN}} = \frac{(d[\text{CN}]/dt)}{\mathcal{N}_{\text{N}}}, \quad (2.82)$$

where \mathcal{N}_{N} is the number flux of particles of nitrogen impinging the surface.




Figure 2.6: ACA model predictions of the efficiency of the nitridation and the recombination reactions at low pressure (solid lines) and high pressure (dashed lines) versus the experimental values. Adapted from Prata et al. [1].

The adsorption process (reaction 1) was modeled as an activated process to reproduce the drop of the N_2 efficiency observed at low temperatures in the molecular beam experiments. The rate coefficient of the desorption process (reaction 2) was computed based on the transition state theory. A desorption energy equal to the bond energy of a double bond between C and N was chosen to describe the non-decreasing efficiency trend observed at high temperatures in the molecular beam experiments. Two nitridation reactions were included in the model to capture the significantly different experimental efficiencies at low and high pressure: a gas-phase dependent reaction (reaction 3) allows for capturing the high-pressure behavior, while the gas-independent reaction (reaction 6) for describing the behavior at low pressure. Lastly, both ER (reaction 4) and LH (reaction 5) mechanisms were included in the model to describe the nitrogen recombination. Although

molecular beam experiments only showed evidence of LH recombination, the ER recombination was necessary to describe the recombination at different pressures and temperature [1].

2.3 Gas surface interaction boundary condition in MUTATION⁺⁺

In this section, we briefly introduce the Mutation⁺⁺ library and its main features. Then, we describe the implementation of the FRC reactions reported in Table 2.1 into Mutation⁺⁺.

2.3.1 MUTATION⁺⁺ description

The MULTicomponent Thermodynamic And Transport properties for IONized gases library written in C++ (Mutation⁺⁺) was developed at the VKI to centralize physico-chemical models, algorithms, and data into a single software packages [68, 90, 99, 109], which may be shared among several CFD solvers. Examples of coupled codes are US3D [11], Argo [110], PATO [111], HDG [112], COOLFluid [113], CHESS [114], Inca [115], SU2-NEMO [116], and a solver developed by Margaritis et al. [117], that also support a surrogate version of the library [118].

Resorting to an external library presents several advantages: the equations implemented in the library can be verified separately, and the CFD code verification process is replaced by a verification of the coupling; being Mutation⁺⁺ open source, different CFD solvers share the same features, helping with code-to-code comparisons, as well regular updates of the library.

As discussed in the previous section, most of the gas properties depend solely on the thermodynamic state of the gas. The Mutation⁺⁺ library takes such a state as input to return:

1. *Thermodynamic properties.* They can be computed either using the RRHO model or by means of the NASA polynomials. Different options are available: thermochemical equilibrium, chemical non-equilibrium, and thermochemical non-equilibrium (Park's two temperature model).
2. *Transport coefficients and diffusion velocities.* Viscosity and thermal conductivity are computed by means of the Chapman-Enskog solution. Diffusive velocities can be computed either by using the full diffusive matrix or by means of the Stefan-Maxwell formulation. In the latter case, together with

the thermodynamic state of the gas, the diffusion driving force should be also passed to the code.

3. *Chemical production rates.* The user has to specify the set of gas-phase reactions along with the relative Arrhenius parameters. They are used to compute the forward reaction rate coefficients. The backwards coefficients are evaluated based on the equilibrium constant, which is determined by solving Equation (2.14).
4. *Transfer terms.* Energy modes exchanges are provided when working in a multi-temperature framework. Its description goes behind the scope of the thesis and is not given here.

2.3.2 Nitridation model implementation and verification

The Mutation⁺⁺ library was expanded by Bellas [99, 109] and Dias [119] to enable the solution of the GSI balances. It works as sketched in Figure 2.7. The user has to specify the set of surface reactions, the type of approach (phenomenological or FRC), and the surface emissivity when solving a coupled energy equation. At this point, the module can be called by the CFD solver passing a tentative value of the surface state (when an isothermal BC is used, the surface temperature is known and its value will not be updated), the gas state in the first cell close to the surface, and the distance of the cell-center/surface, for computing the transport fluxes. The GSI module then solves, through a Newton method the System (2.60) coupled, if necessary, to Equation (2.65), and returns to the CFD solver the surface state and the mass blowing rate. The Newton algorithm is terminated when the iterative error drops below a tolerance of 10^{-13} or when a maximum of 5 iterations is reached. Both phenomenological and FRC approaches are supported. In the latter case, a Newton sub-loop is performed to obtain the free sites steady-state solution, Equation (2.76).

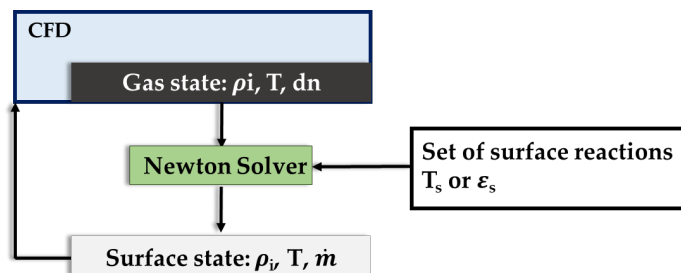


Figure 2.7: Sketch of the surface balance solution.

We expanded the GSI module of Mutation⁺⁺ to handle thermal non-equilibrium

phenomena at the surface [120] - whose description goes behind the scope of the manuscript - and to incorporate the FRC model in Table 2.1 [121]. The subset of the ACA reactions involving nitrogen was included in the library. Given the abstract architecture of the code, such an implementation is relatively straightforward. We wrote a GSI file containing the set of reactions, and included in the library the missing reaction rates laws. Furthermore, the default number of sub-iterations (5) to achieve the sites steady-state condition (Equation (2.76)) was increased to 50 to have a well converged solution. Such an increment was necessary because of the quadratic dependence in $[N_s]$ contained in Equation (2.80). The implementation was verified by comparing the solution at high pressure ($p = 1600$ Pa) from Prata et al. [1] with the one returned by Mutation⁺⁺, in the temperature range $T = [850, 2700]$ K. As one can see in Figure 2.8, the two solutions perfectly agree, verifying the implementation. The model predictive capabilities were then assessed, through CFD simulations, against experimental data. More details will be given in Chapter 6.




Figure 2.8: Verification of the implementation of the reactions in Table 2.1 into Mutation⁺⁺.

2.4 Experiments in ground testing facilities

Experiments in ground testing facilities can serve two distinct purposes. They can be directly used for the characterization of the QoIs. In the context of TPM thermo-chemical characterization, plasma wind tunnels can duplicate the in-flight response at precise trajectory points. ICP facilities [9, 122] and arc jets [123–125] are capable of sustaining a high-enthalpy flow for a relatively long time, typically several minutes or higher, which makes them suitable for the TPM characteriza-

tion. Nevertheless, when one is interested in collecting information at different trajectory points, this practice might result time-consuming. Alternatively, one can employ the experiments to inform physical models, which can be plugged in a CFD solver to reproduce more cost-efficiently different trajectory points. Plasma wind tunnel experiments have been widely used to characterize GSI models. In fact, the observed quantities, heat fluxes or recession rates, can be related to the macroscopic efficiencies of catalytic or ablative reactions. Nevertheless, along with the GSI phenomena that one is interested in characterizing, gas-phase phenomena (chemistry and diffusion) are also present. Therefore, the assumptions contained in the modeling of these processes directly affect the characterization of the surface reactions [26, 67].

When experiments are used to inform a model, one does not necessarily need to duplicate specific in-flight conditions and different type of experiments might be used. For instance, molecular beam-surface scattering experiments [126] can be performed to isolate the GSI phenomenon by bombarding a sample placed in a high-vacuum chamber with a beam of atoms or molecules. These experiments provide more insights on how the surface reactions do proceed, which are useful for calibration purposes. However, while these experiments manage to avoid the appearance of spurious gas-phase phenomena, we cannot fully trust that they are representative of hypersonic flight conditions because of the low pressures involved. Good practice would then be to calibrate the models based on informative molecular beam experiments, and to use plasma wind tunnels to assess their validity at atmospheric entry conditions.

In this section, we describe two types of facilities whose data were employed to inform the targeted GSI models. The first one is the Plasmatron ICP facility at VKI, and the second one is the molecular beam-surface scattering apparatus.

2.4.1 VKI Plasmatron

The VKI Plasmatron facility is a plasma wind tunnel that has been extensively used for the characterization of TPMs. In fact, when the Local Heat Transfer Simulation (LHTS) constraints are met, the Plasmatron can duplicate the stagnation-point heat flux and recession experienced by an object traveling at hypersonic speed by means of a subsonic flow [10, 127, 128]. Particularly, when the two edge conditions are identical in terms of total enthalpy, pressure, and radial velocity gradient in radial direction, the two BLs coincide. Although this feature is not strictly necessary for modeling purposes, it is desirable to ensure that the model is representative of atmospheric entry conditions.

The Plasmatron facility, sketched in figure 2.9, generates high-purity/high-enthalpy flows by means of a 160 mm diameter ICP torch powered by a high-frequency, high-power, high-voltage (400 kHz, 1.2 MW, 2 kV) generator [9]. The facility is started with a gas of Argon, as it is easy to ionize. Once the plasma is stabilized, the gas is switched to the operating one (air, N_2 or CO_2). The gas exits in the test chamber, where vacuum pumps maintain the pressure to the operating value. A holder injects the test sample on the centreline of the plasma jet.




Figure 2.9: Sketch of the Plasmatron facility, adapted from Fagnani [2].

For flow characterization, a copper calorimeter can be inserted in the plasma jet to measure the cold wall heat flux at the stagnation point. It can be swapped with a Pitot probe, which allows for determining the stagnation pressure. Both probes are water-cooled: the temperature is kept around 350 K. The heat flux is retrieved by using the known water mass flow along with the measured temperature difference in the water at the inlet and the outlet of the calorimeter, which is measured by two type-E thermocouples. The chain leads to an uncertainty of 10% of the measurement. The pressure is measured by a Validyne variable reluctance pressure transducer whose signal is amplified by a voltage demodulator. The uncertainty on the measurement is estimated to be 0.25% of the reading scale. The static pressure is assessed by means of an absolute pressure transducer, with uncertainty also of 0.25% of the reading scale. The mass flow rate is controlled by a rotameter, with an accuracy of 5%.

These measurements, along with the power supplied by the generator and its effi-


ciency, serve as an input for the flow rebuilding procedure developed at VKI [129]. First, the VKI ICP code [130, 131] simulates the Plasmatron torch and the test chamber by solving magneto-dynamic equations, under the Local Thermodynamic Equilibrium (LTE) assumption. In this way, it is possible to retrieve the flow conditions at the edge of the BL that develops in front of the probes. After that, the VKI BL code [104] solves the chemically reacting Navier-Stokes equations in the BL to compute the heat flux on the probe, assuming it to be catalytic. The BL edge temperature is then iterated until the numerical and the experimental heat flux do coincide.

Further flow characterization is possible by means of several windows, which allows for optical measurements [3, 132–134]. A spectrometer can detect the locally-resolved radiative signature, and, under LTE assumption, the gas temperature. A high-speed camera is used to measure the sample recession, while a pyrometer the surface temperature. Infrared thermography can be used to obtain 3D temperature maps of the sample.

The facility has been extensively used in subsonic regime for the characterization of catalytic and ablative materials [3, 134–138]. Nevertheless, a convergent or a convergent/divergent nozzle can be attached at the torch exit to accelerate the flow to supersonic conditions [139–141]. This allows for maximizing the stagnation-point heat flux and the shape change of ablative samples.

Subsonic campaign

An experimental campaign was performed at VKI to study the temperature-dependency of the nitridation reaction efficiency on carbon surfaces [3]:



The experiment consisted of a plasma nitrogen flow over a superfine-grain/high-density graphite ($\rho_s = 1760 \text{ kg/m}^3$), in-house machined to have a 25 mm radius hemispherical shape with a 25 mm cylindrical after-body. It was placed at a distance of 445 mm from the exit of the torch.

The experimental set-up included a digital camera for the determination of the stagnation-point recession, a two-color pyrometer for measuring the surface temperature of the sample, and a spectrometer to record the violet system of the CN molecule in the BL. It resolved the radial distribution (on the vertical line of the right Figure 2.10) versus the wavelength of the signal (on the horizontal line of the right Figure 2.10). Such a 2D matrix can be Abel-inverted to obtain the locally-resolved CN-emission. By post-processing it, it was possible to retrieve both the gas temperature and the CN concentration in a point belonging to the BL.

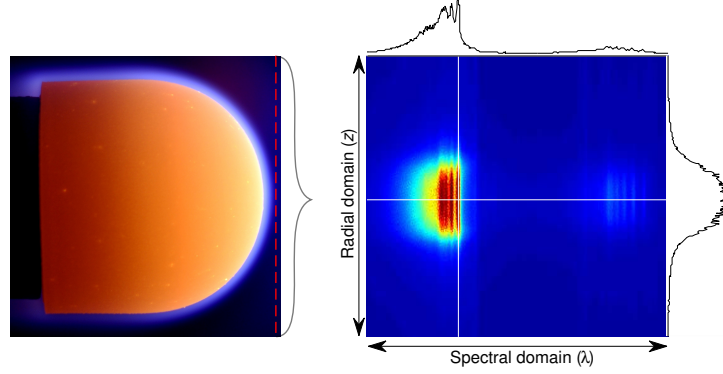


Figure 2.10: Ablative carbon sample in a nitrogen flow in the left figure. Spectrometer output matrix, in the right one, with spectral emission on the horizontal axis, and radial position on the vertical axis. The spectrometer slit position is indicated with dashed red line in the left figure. Figure courtesy of Helber et al. [3].

Seven runs were performed. From one to the other, the electrical power supplied to the generator was progressively increased to obtain a higher free-stream enthalpy, and, in turn, a higher surface temperature of the sample.

The BL edge conditions (T_e , u_e) were obtained through the rebuilding procedure [3] using the dynamic pressure and the heat flux measured, respectively, by a Pitot probe and a calorimeter, the chamber pressure, and the power supplied to the Plasmatron and its efficiency.

Three of the seven runs exhibited a strong thermal non-equilibrium and will not be considered for the calibration proposed within this manuscript. An overview of the conditions of the four considered runs is given in Table 2.2.

CASE	Experimental				Rebuilt	
	p_d [Pa]	P_{el} [kW]	T_w [K]	\dot{m} [g/m ² /s]	T_e [K]	u_e [m/s]
G4	231 ± 4.6	280	2225 ± 22	2.49 ± 0.91	10005	554
G5	268 ± 5.4	330	2410 ± 24	2.89 ± 0.97	10280	562
G6	312 ± 6.2	370	2535 ± 26	4.41 ± 0.80	11040	846
G7	330 ± 6.6	390	2575 ± 26	4.56 ± 0.70	10970	859

Table 2.2: Overview of Plasmatron test conditions for a nitrogen flow at 1500 ± 50 Pa over a 25 mm radius hemisphere carbonaceous sample: dynamic pressure p_d , generator power P_{el} , sample mean temperature T_w , mass blowing rate \dot{m} , and reconstructed BL edge conditions (temperature, T_e , and velocity, u_e). Conditions from Helber et al. [3].

Supersonic campaign

An experimental campaign was conducted at VKI using a supersonic jet to maximize the heat flux experienced by the sample, as well as its shape change. In the used configuration, the supersonic condition was achieved by accelerating the flow through a sonic nozzle, mounted at the exit of the plasma torch. The chamber pressure was lowered enough to generate a highly under-expanded jet.




Figure 2.11: Under-expanded air jet over the catalytic probe obtained in the VKI's Plasmatron facility.

Specifically, the chamber pressure was set to 5.5 hPa, against a reservoir pressure of 165 hPa, resulting in a total pressure ratio of $\beta_0 = p_0/p_c = 30$. The sonic nozzle has an exit diameter of 35 mm and the probe was placed at a distance of 75 mm from the exit of the nozzle, before the occurrence of the Mach disk, after which the flow would become subsonic. The mass flow entering the plasma torch, measured through a rotameter, was 6 g s^{-1} , and the Plasmatron was supplied with an electrical power of 600 kW. These conditions allowed for reaching a target heat flux of 4.5 MW.

The supersonic flow is shown in Figure 2.11: it expands out of the nozzle and a detached shock develops in front of the sample. The interaction between it and the barrel shock (developing from the exit of the nozzle) generates secondary shocks, that get reflected on the jet boundary.

Unlike the subsonic experiments, the chemistry in the supersonic flow is not expected to be in equilibrium and the VKI rebuilding procedure cannot be used. In fact, the VKI ICP code assumes an LTE condition. In Chapter 5, we will describe a novel methodology to characterize such a flow.

2.4.2 Molecular beam-surface scattering apparatus

Recent FRC models for describing carbon ablation [1, 24] were inferred from molecular beam-surface scattering experiments performed by Murray et al. [108,

126, 142, 143], as they provide microscopic insights on how surface reactions proceed. In fact, molecular beam-surface scattering is a technique that allows for assessing the dynamics of the interaction between a molecular beam (collisionless and well collimated plume of gas) and a surface [144].

The main components of the experimental apparatus are I) a beam source (for example a high-pressure radio frequency discharge source), II) a heated sample, over which the beam is directed, III) a rotatable mass spectrometer, and IV) a chopper wheel to modulate the scattered products. The sample is placed in a high-vacuum room, where pumps guarantee pressure of the order of 10×10^{-7} torr, essential to avoid gas-phase collisions. In addition, the sample temperature is measured by means of a calorimeter.

The rotatable mass spectrometer detects the number density distribution of the scattered products as a function of time, $N(t)$, at different final angle, θ_f , corresponding to an incident angle, θ_i . The relative curves are generally referred to as Time-Of-Flight (TOF) distributions. The flux of the scattered particles at a given final angle is computed as:

$$I(\theta_f) \propto \int_{t_i}^{t_f} \frac{N(t)}{t} dt, \quad (2.83)$$

where t_i is the time at which the chopper wheel slot passes in front of the detector, and t_f is the time at which the signal drops to the background level. The plot of the flux versus the final angle is referred to as angular distribution. By analyzing the shape of the TOF distributions, and of the angular and translational energy ($P(E_T) \propto t^2 N(t)$) distributions, insights on the type of mechanisms can be derived (impulsive scattering VS thermal desorption, LH VS ER mechanisms, etc.). For example, in a LH reaction the incident atom and the produced molecule does not thermalize with the surface. As a result, the products retain most of the incident atom's high translational energy. The relative TOF curves are characterized by a peak at short flight time and the angular distributions by a lobular shape [126].

As we will discuss in Chapter 6, we informed an FRC model by means of an experiment performed by Murray et al. [142]. Following, we present its main features. In the experiment, low-energy continuous beams were employed to study the reactions between a nitrogen, and an oxygen, gas and a vitreous carbon surface. Compared to the hyperthermal pulsed beams employed in previous studies by the same author [108, 143], low-energy and continuous beams ensured that the surface chemical processes were more representative of atmospheric entry conditions. Both oxygen and nitrogen beams were used. Being in this manuscript interested in calibrating the reactions involving nitrogen, we will consider only the results obtained using nitrogen beams.

The nitrogen beam was produced using a 95 mbar mixture (2.5% N₂ in He) discharged by a high-pressure radio frequency discharge source and expanded through a 0.48 mm diameter nozzle. It is then collimated by means of a skimmer with a diameter of 0.8 mm and after by a 0.4 × 0.4 mm² aperture. The resulting beam is composed by $x_N = 0.18$ and $x_{N_2} = 0.82$ with an average velocity of ≈ 2000 m/s. The data were collected at increasing temperatures (from 800 to 1873 K), annealing, before each test, the sample at 1873 K for around 1 h.

Results showed that the surface undergoes only recombination and nitridation reactions. Their reaction probabilities were determined as:

$$\gamma^r = \frac{\text{flux of } r \text{ product}}{\text{flux of N atoms onto the surface}} = \frac{n f(\text{product}_r)}{f(N + 2N_2 + CN)}, \quad (2.84)$$

where n is the number of nitrogen atoms in the specific product. The flux of products was computed by integrating the flux of the scattered particles over all the final angles.

This is the only application treated in this manuscript which is not described by means of the Navier-Stokes equations. Indeed, being the experiment conducted in high vacuum, one can neglect the gas collisions. We employed the OD model proposed by Prata et al. [1] to simulate the experiment. It was presented in Section 2.2.4 (*Air carbon ablation model* subsection).

2.5 Summary

This chapter aimed to present a comprehensive understanding of the equations that govern a hypersonic flow in the continuum regime and to describe the experimental facilities and campaigns that we will analyze in the following chapters.

The chapter begins by introducing the chemically reacting Navier-Stokes equations, which express the conservation of mass, momentum, and energy. In many hypersonic flows, the chemistry deviates from the equilibrium condition and homogeneous FRC needs to be accounted for. We presented the expression of the net chemical production rates, based on the law of mass action and empirical Arrhenius-like reactions rate coefficients. The chapter then moves into the descriptions of the thermodynamic properties of high-temperature gases. Two commonly used approaches are detailed: the RRHO model and the NASA polynomials. The chapter also covers the modeling of transport properties. The Chapman-Enskog formulation is presented to compute the gas viscosity and thermal conductivity, while the generalized Stefan-Maxwell equation to compute the diffusive velocities. Additionally, simplified mixture rules are described. The chapter then delves into the modeling of the GSI, needed to describe the interaction of the reacting gas and catalytic or ablative materials. We introduced two different approaches for the closure of

the surface chemical source term: the phenomenological and the FRC models. Both of them are available within the Mutation⁺⁺ library, which was described next. We implemented in the library the reactions involving atomic nitrogen of a state-of-the-art FRC model and we presented a verification test case. Finally, the chapter concludes by introducing a subsonic and a supersonic test performed in the VKI Plasmatron facility and a molecular beam-surface scattering experiment.

3

Uncertainty quantification methods

This chapter reviews the formulations behind the uncertainty quantification tools employed in this thesis. We first introduce the forward propagation problem. The Monte Carlo method is described, highlighting the differences among various sampling strategies. Next, we present the Bayesian formulation, used in this thesis to solve the inverse problem in a stochastic fashion. We show that a Markov Chain Monte Carlo can be employed for approximating the target posterior distribution. Examples of efficient algorithms used to construct such a chain are then described: Metropolis-Hasting, adaptive Metropolis, and the affine invariant ensemble sampler. Both the forward propagation and the inverse problem require the computation of a large number of numerical solutions. They can be efficiently performed through surrogate models. Their mathematical formulation is outlined at the end of the chapter, focusing on the Kriging method and on two strategies to enhance its efficiency and accuracy: multi-fidelity methods and adaptive sampling. Finally, we introduce the UQLAB software, that we employed in the thesis to perform the UQ studies.

3.1 Overview

Nowadays, it is common practice to rely on computer simulations for the analysis, design, and optimization of engineering systems [145]. In the context of aerospace applications, high-fidelity CFD simulations integrate ground testing experiments in the design of high-speed, access-to-space, and re-entry vehicles [146].

CFD predictions are the solution of an user-implemented algorithm, that translates the mathematical formulation of a physical model, solved on a discrete grid by marching in a discrete time. As pointed out by Roy and Oberkampf [147], they are thus prone to three main classes of uncertainties/errors:

1. *Coding mistakes*;
2. *Numerical approximations*: finite time-steps and volumes make the solution dependent on the degree of the discretization;
3. *Model uncertainty*: models are mathematical representations of the physics and they suffer from uncertainty both in their form and in their input parameters. Uncertainties on the simulation surrounding (for example boundary conditions), simulated geometries, and initial conditions also pertain to this category.

Ideally, coding mistakes are not present in robust solvers: they are identified and eliminated through the *code verification* and the *regression tests*, during which we make sure that *we solve the equations correctly*.

The majority of hypersonic simulations are performed using a second-order-space-accurate finite-volume formulation: a *discretization error* is - in the asymptotic limit - quadratically dependent on the mesh resolution. Further numerical errors may derive from the inability of the grid to well capture some of the flow features. For example, they can arise from mesh-shock misalignment or from cells extremely non-orthogonal to the wall, if specific treatments for dealing with such features are not included in the numerical methods. A broad class of applications seeks a steady-state solution. In these cases, *iterative errors* refer to the lack of convergence to steady-state. On the other hand, when one is interested in the solution time dependency, an error is dependent on the time step employed to time march. The magnitude of these errors is controlled during the *solution verification*: common practice is to have a drop in the residual of at least five orders of magnitude for the steady-state applications and a mesh-independent solution.

Model uncertainties derive from a lack of knowledge (they mostly are pure *epistemic uncertainties*). *Ad hoc* experimental campaigns can be designed to *calibrate* and, subsequently, *validate* the model. In this way, we make sure that *we solve the correct equations*.

Experiments are not exempt from uncertainty due to the variability of the system and the finite precision of the instruments. This kind of uncertainty is generally irreducible (called *aleatory*) and can be characterized by a Probability Density Function (PDF).

While predictions can rely on deterministic simulations, reducing the whole chain of uncertainties to a safety factor, stochastic frameworks allow for a robust

characterization of the output [148]. In this contest, UQ designates the set of methods aiming at dealing with these uncertainties by characterizing, propagating and, possibly, reducing them.

Following the framework outlined by Roy and Oberkampf [147], one should first identify the uncertainties: sensitivity analysis helps in recognizing the most relevant sources of uncertainty. The associated sensitive variables are then treated as random variables by assigning a PDF to each of them. PDFs are chosen based on expert judgment, previous characterization, or derive by experimental uncertainties. At this point, the uncertainties can be propagated through a model - provided that the numerical error does not exceed a certain threshold - using, for example, sampling techniques. Forward propagation problems are described in Section 3.2.

Alternatively, the uncertainty can be propagated backward and the experimental data can be exploited to characterize/reduce the uncertainty on the model inputs. Such an inverse problem is commonly referred to as *model calibration*. A Bayesian framework well suits this purpose as it allows for simultaneously accounting for prior knowledge, epistemic and experimental uncertainties, and inadequacies of the model [147, 149]. The Bayesian formulation will be discussed in Section 3.3, along with the algorithms used to solve the inverse problem.

Both forward and inverse problems are generally solved by means of sampling strategies. The statistical convergence of the output PDFs is achieved upon the use of thousands of points. In order to enhance the method's efficiency, it is common practice to approximate the numerical solver response with a computationally cheaper mathematical representation of it. Such representations are generally called *surrogate models*. Examples are Gaussian Process (GP) or Kriging [29, 150], PCE [151, 152], and artificial neural network [153]. In this thesis we resorted to GPs, whose mathematical formulation is presented in Section 3.4.

3.2 Forward propagation

In many engineering applications, one is interested in characterizing a QoI, y , that is not experimentally measurable or theoretically derivable. Auxiliary parameters, \boldsymbol{x}^* , may be used for its determination through a model $y = \mathcal{M}(\boldsymbol{x}^*)$. In general, the input vector \boldsymbol{x}^* contains both parameters that we know within a given uncertainty, \boldsymbol{x} , and parameters assumed to be perfectly known, \boldsymbol{t} . Since the latter are considered fixed in the forward problem, in the following, we will omit the model dependence on \boldsymbol{t} .

Forward propagation refers to the characterization of the stochastic model response given the stochasticity of the model input. The definition of a proba-

bilistic space¹ allows for expressing the model input and output as random vectors/variables². They are usually indicated with capital letters, respectively \mathbf{X} , and $Y = \mathcal{M}(\mathbf{X})$. It is then possible to evaluate the output's PDFs or statistical moments. Of main importance for engineering applications are the first (mean, μ_y) and the second (variance, σ_y^2) moment. They are defined as:

$$\mu_y = \mathbb{E}[\mathcal{M}(\mathbf{x})] = \int_{-\infty}^{\infty} \mathcal{M}(\mathbf{x})p(\mathbf{x})d\mathbf{x}, \quad (3.1)$$

$$\sigma_y^2 = \mathbb{E} \left[(\mathcal{M}(\mathbf{x}) - \mathbb{E}[\mathcal{M}(\mathbf{x})])^2 \right] = \int_{-\infty}^{\infty} (\mathcal{M}(\mathbf{x}) - \mathbb{E}[\mathcal{M}(\mathbf{x})])^2 p(\mathbf{x})d\mathbf{x}, \quad (3.2)$$

where $p(\mathbf{x})$ is the joint PDF of the random vector \mathbf{X} , and the operator $\mathbb{E}[\cdot]$ denotes the expectation.

The majority of engineering applications do not offer an easy analytical form of Equations (3.1) and (3.2) and numerical methods are generally employed for their computation. Two main classes of methods exist. *Intrusive methods* requires the modification of the numerical structure of the code to accommodate a stochastic representation of the problem, enabling it to output the response in terms of random vectors. On the other hand, *non-intrusive methods* do not require any code modification. The numerical solver, used as a black box, is called several times to sample realizations of the output. Non-intrusive methods have the advantage to be robust, as no code modification is needed, and easy to parallelize. A popular choice for uncertainty propagation is to resort to spectral methods, as the PCE [151], which express the model response in terms of a series of orthonormal polynomials. The forward problem can be then performed either analytically, exploiting the polynomial coefficients, or through sampling strategies, using the PCE model as surrogate model. A review of intrusive and non-intrusive methods is given by Le Maître and Knio [155]. Following, we will describe non-intrusive sampling strategies.

The MC technique [156, 157] is the most straightforward and popular non-intrusive method. It requires to quasi-randomly sample N_s independent realizations of \mathbf{X} , $\mathbf{x}^{(i)}$, according to its joint PDF. Then, the random variable Y can be inferred through the collection of the N_s independent model realizations, $\mathcal{M}(\mathbf{x}^{(i)})$. Specifically, the numerical approximations of the first two moments can be computed

¹A probability space is defined by a sample space Ω (set of all possible outcomes), an event space Σ (collection of subsets of Ω), and a non-negative probability measure P . The event space is a σ -field on Ω : it is a non-empty collection of subset of Ω , closed under complementation, union, and intersection [154]. An important σ -field is the collection of all the open intervals of \mathbb{R} , the Borel σ -field, \mathcal{B} .

²A random variable, X , assigns a real number to each possible outcome of a stochastic experiment. Considering the probability space (Ω, Σ, P) , an \mathbb{R} -valued random variable is a measurable function from (Ω, Σ) to $(\mathbb{R}, \mathcal{B})$. The same definitions hold in the case of a \mathbb{R}^n -valued function from (Ω, Σ) to $(\mathbb{R}^n, \mathcal{B}^n)$. In this case, \mathbf{X} is called random vector.

as:

$$\bar{\mu}_y = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathcal{M}(\mathbf{x}^{(i)}), \quad (3.3)$$


$$\bar{\sigma}_y^2 = \frac{1}{N_s} \sum_{i=1}^{N_s} \left(\mathcal{M}(\mathbf{x}^{(i)}) - \bar{\mu}_y \right)^2. \quad (3.4)$$

The law of large numbers ensures that the MC estimators converge to the expected values ($\bar{\mu}_y \rightarrow \mu_y, \bar{\sigma}_y \rightarrow \sigma_y$) for a large number of sample points ($N_s \rightarrow \infty$). The convergence rate of the MC method is $\mathcal{O}(N_s^{-0.5})$. It is the strongest and the weakest point of MC techniques. In fact, on the one hand, it is independent of the dimension of the problem, on the other, it requires the collection of an important number of points. For this reason, numerous strategies were developed to accelerate the method's convergence [156]. Among other methods, the Latin Hypercube Sampling (LHS) [158] ensures better coverage of the sampling space by dividing it into N_s equally probable subspaces and drawing a sample from each of them. The same is achieved by quasi-MC methods by employing a low-discrepancy sequence, for example, the Sobol's sequence [159].

Example Let us consider the Sutton Graves [160] correlation for stagnation-point heat-flux computation:

$$q = 1.7415 \cdot 10^{-4} \sqrt{\frac{\rho}{R}} v^3.$$

Let us assume that we perfectly know the atmosphere density $\rho = 2.88 \cdot 10^{-4}$ kg/m³, and the object radius $R = 0.3$ m and that we want to characterize the uncertainty on q given the uncertainty on the object velocity, v , assumed normally distributed, $V = \mathcal{N}(4900, 490)$. We randomly sampled $N_{MC\ 1} = 100$ and $N_{MC\ 2} = 10^5$ points with a MC technique, and $N_{LHS} = 100$ using a LHS method. We propagated them through the analytical model $q(v)$, obtaining the following probabilistic densities:



The distribution obtained using 10^5 MC points (MC 2) appears to be well converged and can be used as a reference. As one can see, such a distribution is approximated much better employing 100 LHS points, than 100 MC points (MC 1). In fact, although the LHS is a stochastic technique as the MC one (the sampled points are different every time we run the algorithm), it provides an improved filling of the sampling space.

3.3 Stochastic inverse problem

Inverse problems aim at inferring one or more model input parameters, \mathbf{x}^* , given N^{obs} auxiliary experimental observations, \mathbf{y}^{obs} . In *deterministic* inverse problems, the best estimate of \mathbf{x}^* is found by optimizing an error or likelihood function, $f(\mathbf{y}^{\text{obs}}, \mathcal{M}(\mathbf{x}))$. The outcome is a punctual realization of an underlying distribution which is not characterized. Contrary, in *stochastic* inverse problems, the input parameters that one intends to infer, \mathbf{x}^* , as well as all the other model parameters affected by uncertainty, \mathbf{x}' , are treated as random vectors [161]. Specifically, a Bayesian formulation allows one to obtain the posterior distributions of the parameters given the uncertainty on the observations.

Let us denote $\mathbf{x} = \{\mathbf{x}^*, \mathbf{x}'\}$, the random vector including all the unknown/uncertain parameters. According to Bayes' theorem, the posterior distribution of \mathbf{x} is:

$$\pi(\mathbf{x}|\mathbf{y}^{\text{obs}}) = \frac{\pi(\mathbf{x})\pi(\mathbf{y}^{\text{obs}}|\mathbf{x})}{\int \pi(\mathbf{x})\pi(\mathbf{y}^{\text{obs}}|\mathbf{x})d\mathbf{x}}, \quad (3.5)$$

where $\pi(\mathbf{x})$ is the prior distribution of \mathbf{x} , $\pi(\mathbf{y}^{\text{obs}}|\mathbf{x})$ indicates the likelihood of an observable quantity \mathbf{y}^{obs} given \mathbf{x} , and the integral at the denominator represents the marginal likelihood or evidence, which ensures the distribution to integrate to 1. Although the Bayesian approach may look straightforward, several crucial aspects

need to be correctly addressed.

The first aspect to consider is the definition of the *prior distribution* [161], which express our prior belief/knowledge about \boldsymbol{x} prior to observing the data. It can come from a previous calibration, from experiments, or from expert judgment. A correct specification of $\pi(\boldsymbol{x})$ is particularly important when the number of observations is limited (prior-dominated problem). In order to let the likelihood dominate the posterior, uniform *non-informative* prior should be specified. Log-uniform priors represent a preferable choice when the ignorance about the parameters spans several orders of magnitude. It is of course possible, when one has a strong belief over a quantity, to specify *informative* priors, for example, Gaussian experimental uncertainties. Of utmost importance is the definition of the range of the prior. Let us imagine to specify that $\pi(\boldsymbol{x} < \boldsymbol{a}) = 0$, then the posterior distribution for $\boldsymbol{x} < \boldsymbol{a}$ will always be equal to zero, whatever number of observations would indicate the opposite. This property can be used to constrain the posterior to physical values, for example to non-negative pressures or densities.

The second aspect concerns the definition of the likelihood function [162], which expresses how likely the realization $\boldsymbol{y}^{\text{obs}}$ can be obtained by the input \boldsymbol{x} through a numerical model $\mathcal{M}(\boldsymbol{x})$. A *discrepancy term* describes the departure of the model prediction from the observed quantity. It comprises both the model inaccuracy and the measurement noise. It is common practice to express it as an additive term:

$$\boldsymbol{\epsilon} = \boldsymbol{y}^{\text{obs}} - \mathcal{M}(\boldsymbol{x}). \quad (3.6)$$

Assuming that I) the model error is negligible with respect to the experimental uncertainty, σ_i , on the measurement y_i^{obs} , II) the experimental uncertainty is known and normally distributed around the mean ($\mathcal{N}(0, \sigma_i^2)$), and III) independence between the N^{obs} observations, the likelihood reads:

$$\pi(\boldsymbol{y}^{\text{obs}}|\boldsymbol{x}) = \prod_{i=1}^{N^{\text{obs}}} \mathcal{N}(y_i^{\text{obs}}|\mathcal{M}_i(\boldsymbol{x}), \sigma_i^2). \quad (3.7)$$

Lastly, in many engineering applications, the evidence of Equation (3.5) is difficult to compute analytically. For this reason, methods that allow for sampling directly from the posterior distribution are very attractive. Markov Chain Monte Carlo (MCMC) algorithms are an important class designed for this purpose.

3.3.1 MCMC algorithms

MCMC algorithms are powerful tools as they allow for approximating the posterior distribution of a Bayesian problem with the invariant distribution of a Markov

chain. An exhaustive explanation of the Markov chains can be found in the textbooks [161, 163]; next, we will present the main features.

A sequence of realizations $(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(t)})$ is a Markov chain on the state space \mathbb{R}^d if for all the iterations t and for all the measurable sets A , the *transition probability* to move to the state $\mathbf{x}^{(t+1)}$ only depends on the state $\mathbf{x}^{(t)}$. Such a condition is generally referred to as Markov property and is expressed as:

$$\mathcal{P}(\mathbf{X}^{(t+1)} \in A | \mathbf{x}^{(t)}, \mathbf{x}^{(t-1)}, \mathbf{x}^{(0)}) = \mathcal{P}(\mathbf{X}^{(t+1)} \in A | \mathbf{x}^{(t)}). \quad (3.8)$$

If the transition probability does not depend on the iteration, the Markov chain is said to be *time-homogeneous* or simply *homogeneous*.

A Markov chain is thus fully defined by an initial state, $\mathbf{x}^{(0)}$, and by the transition kernel, $\mathcal{K}(\mathbf{x}, \mathbf{x}')$. It is defined as:

$$\mathcal{P}(\mathbf{X} \in A | \mathbf{x}) = \int_A \mathcal{K}(\mathbf{x}, \mathbf{x}') d\mathbf{x}'. \quad (3.9)$$

The chain should be able to explore the whole support starting from an arbitrary initial state to eventually converge to the posterior distribution. It is essential that I) all the states communicate (the chain is said to be *irreducible*), and II) it is possible to return to any state at any time (the chain is said to be *aperiodic*). If these two conditions hold, the chain is *ergodic*. It is an important property as it ensures that the chain will converge to the *invariant* or *stationary* distribution whatever the initial point is. A distribution, $\pi(\cdot)$, is invariant if the global balance is respected:

$$\pi(\mathbf{x}^{(t+1)}) = \int_{\mathbb{R}^d} \mathcal{K}(\mathbf{x}^{(t)}, \mathbf{x}^{(t+1)}) \pi(\mathbf{x}^{(t)}) d\mathbf{x}^{(t)}, \quad (3.10)$$

meaning that the chain reached the stationary regime where $\pi(\cdot)$ does not depend on the iteration.

A chain is *reversible*, meaning that the probability to be in the state $\mathbf{x}^{(t)}$ and move to the state $\mathbf{x}^{(t+1)}$ is the same of being in the state $\mathbf{x}^{(t+1)}$ and move to the state $\mathbf{x}^{(t)}$, if the detailed balance condition is satisfied:

$$\pi(\mathbf{x}^{(t)}) \mathcal{K}(\mathbf{x}^{(t)}, \mathbf{x}^{(t+1)}) = \pi(\mathbf{x}^{(t+1)}) \mathcal{K}(\mathbf{x}^{(t+1)}, \mathbf{x}^{(t)}). \quad (3.11)$$

It can be shown that if a chain is reversible, it is also stationary. It is an important property because the condition in Equation (3.11) is easier to be assessed than the condition in Equation (3.10) [161].

Several algorithms have been proposed in literature to obtain such an invariant chain. Next, we will revise three of them.

Metropolis-Hastings

The Metropolis-Hastings algorithm, initially proposed by Metropolis et al. [164], successively extended by Hastings [165], is the first and the most popular MCMC

algorithm. Its implementation is very easy and straightforward. The main steps are summarized in the algorithm 1.

```

Initialize  $\mathbf{x}^{(0)}$ 
for  $t = 1$  to  $N$  do
  Sample  $\mathbf{x}^*$  from proposal distribution  $q(\mathbf{x}^* | \mathbf{x}^{(t-1)})$ 
  Compute acceptance probability:  $\alpha$ 
  Sample  $u$  from uniform distribution  $U(0, 1)$ 
  if  $\alpha \geq u$  then
    Accept state:  $\mathbf{x}^{(t)} = \mathbf{x}^*$ 
  else
    Reject state:  $\mathbf{x}^{(t)} = \mathbf{x}^{(t-1)}$ 

```

algorithm 1: Metropolis-Hastings algorithm

The chain is initialized at a state $\mathbf{x}^{(0)}$, sampled from the prior distribution. Successively, a new sample is generated according to a *proposal distribution*, $q(\mathbf{x}^* | \mathbf{x}^{(t-1)})$, which may depend on the state at the previous iteration. The sample is then accepted with an *acceptance probability*:

$$\alpha = \min \left(1, \frac{\pi(\mathbf{x}^* | \mathbf{y}^{\text{obs}}) q(\mathbf{x}^{(t-1)} | \mathbf{x}^*)}{\pi(\mathbf{x}^{(t-1)} | \mathbf{y}^{\text{obs}}) q(\mathbf{x}^* | \mathbf{x}^{(t-1)})} \right). \quad (3.12)$$

It depends both on the proposal and on the target distributions. According to Equation (3.5), the ratio between the target distributions at the proposed state, \mathbf{x}^* , and the previous state, $\mathbf{x}^{(t-1)}$, can be rearranged as:

$$\frac{\pi(\mathbf{x}^* | \mathbf{y}^{\text{obs}})}{\pi(\mathbf{x}^{(t-1)} | \mathbf{y}^{\text{obs}})} = \frac{\pi(\mathbf{x}^*)}{\pi(\mathbf{x}^{(t-1)})} \frac{\pi(\mathbf{y}^{\text{obs}} | \mathbf{x}^*)}{\pi(\mathbf{y}^{\text{obs}} | \mathbf{x}^{(t-1)})}, \quad (3.13)$$

as the evidence cancels out. The transition kernel is then independent of the hard-to-compute evidence, which is the strength of the algorithm.

It can be shown, that a chain produced with the Metropolis-Hastings algorithm is aperiodic and irreducible; furthermore, the transition kernel respects the detailed balance [163]. Thus, the chain can converge to the stationary distribution $\pi(\mathbf{x} | \mathbf{y}^{\text{obs}})$.

What remains to be defined is the proposal distribution. It should resemble the target distribution to guarantee a good convergence of the chain. However, most of the time, we do not have access to its shape.

In the original algorithm by Metropolis et al. [164], the proposal distribution was designed to be symmetric, such that $q(\mathbf{x}^{(t-1)} | \mathbf{x}^*) = q(\mathbf{x}^* | \mathbf{x}^{(t-1)})$. In this case,

the acceptance probability depends solely on the target invariant distribution:

$$\alpha^{\text{sym}} = \min \left(1, \frac{\pi(\mathbf{x}^* | \mathbf{y}^{\text{obs}})}{\pi(\mathbf{x}^{(t-1)} | \mathbf{y}^{\text{obs}})} \right). \quad (3.14)$$

Popular choices for the proposal distribution are uniform or normal distributions centered at $\mathbf{x}^{(t-1)}$. In the latter case:

$$\mathbf{x}^* = \mathbf{x}^{(t-1)} + \mathcal{N}(\mathbf{0}, \Sigma), \quad (3.15)$$


where Σ is the covariance matrix. This choice generates the so-called *random walk Metropolis*.

The choice of the covariance matrix, but more in general of the proposal distribution, significantly impacts the convergence rate of the chain. If the covariance matrix induces large changes in the state, most of the proposed states will be rejected. On the other hand, if the covariance matrix produces small changes in the state, most of the proposed states will be accepted, but the chain will move slowly. One should manually adjust the proposal distribution to have an acceptance rate close to the empirical optimal acceptance rate of 0.234 [166].

Example Let us consider again the Sutton Graves [160] correlation for stagnation-point heat-flux computation:


$$q = 1.7415 \cdot 10^{-4} \sqrt{\frac{\rho}{R}} v^3.$$

This time, we assume that we perfectly know the atmospheric density, $\rho = 2.88 \cdot 10^{-4} \text{ kg/m}^3$ and the object radius $R = 0.3 \text{ m}$, and that we want to characterize the uncertainty on the object velocity, v , given the uncertainty on q , $Q = \mathcal{N}(6 \cdot 10^5, 10^5)$. To illustrate the importance of the choice of the proposal distribution, we run the algorithm 1, using 3 similar distributions, $\mathcal{N}(0, \sigma_i)$, characterized by a different value of σ_i . Specifically: $\sigma_{\text{I}} = 10$, $\sigma_{\text{II}} = 1500$, and $\sigma_{\text{III}} = 25000$. The three random walks of v versus the iteration (referred to as trace plot) are shown in the figures below:



The first figure on the left is the trace plot of the chain built using a *small* standard deviation: many states are accepted ($\alpha^{\text{sym}} = 0.9878$), but the chain evolves slowly. It appears to be *poorly mixed*, and we can assume it will

need many iterations to converge to the stationary distribution. On the other hand, as one can see in the first figure on the right, a chain built with a *big* standard deviation leads to many points being rejected ($\alpha^{\text{sym}} = 0.0137$). Also in this case, the chain is poorly mixed, and the statistical convergence will be difficult. Finally, the trace plot in the middle was obtained by tuning the standard deviation to have an acceptance rate close to the optimal value, $\alpha^{\text{sym}} = 0.2259$. In this case, we can see how the chain is *well mixed*. The estimate of the posterior distribution on v is shown, for the three different chains, in the figure below on the left, while in the figure on the right, we can see the posterior predictive distribution on q . The imposed uncertainty on q is represented with a black solid line.



Contrary to the chains built using too small or too big values of σ , the chain built using the optimal value is characterized by smoother and more continuous histograms, which best approximates both the distribution of q , and the target posterior distribution.

Adaptive Metropolis


To avoid such a manual tuning, Haario et al. [167] proposed an adaptive variant of the Metropolis algorithm. It requires an initial run of t_0 iterations to estimate the covariance matrix of the invariant distribution, Σ' . The normal proposal distribution is then characterized by a covariance matrix which will depend on t :

$$\Sigma(t) = \begin{cases} \Sigma^0 & \text{if } t + 1 \leq t_0 \\ s_d \Sigma'(t) & \text{if } t + 1 > t_0 \end{cases} \quad (3.16)$$

where $s_d = 2.38^2/d$ and d is the problem dimension [168]. This formulation makes the chain dependent on all the previous steps and does not satisfy the Markov property. However, it was shown that the chain respects the correct ergodicity properties and thus can converge to the target distribution [167].

Example Considering the same case as before, we employed the adaptive Metropolis algorithm, with $\sigma^0 = 10$ and $t_0 = 1000$. In the figure below, the

trace plot of the chain is shown.



As one can see, after a slow convergence during the first 1000 iterations, where the original Metropolis algorithm is used, the chain is well mixed when the standard deviation is estimated on the chain sample points.

Affine invariant ensemble sampler

When the posterior distribution is characterized by high-dimensional and highly-correlated parameters, the performance of the Metropolis algorithms might deteriorate and important tuning of the proposal distribution is required to converge to the invariant distribution. To improve the convergence, Goodman and Wear [169] proposed to use an ensemble of MCMC chains, referred to as *walkers*, such that the target distribution is invariant to affine transformations of the parameters³. This property allows the algorithm to equally sample difficult-to-sample distributions and easy-to-sample transformations of them.

The Affine Invariant Ensemble Sampler (AIES) algorithm, outlined in algorithm 2, starts by sampling C walkers $\mathbf{X}^{(0)} = [\mathbf{x}_1^{(0)}, \dots, \mathbf{x}_C^{(0)}]$ from the prior distribution, with $C > d + 1$, where d is the dimension of the parameters. The proposal state for each i walker is sequentially obtained by means of all the current states. One common approach to achieve this is by using the symmetric *stretch move*:

$$\mathbf{x}^* = \mathbf{x}_i^{(t-1)} + z \left(\mathbf{x}_j^{(\hat{t})} - \mathbf{x}_i^{(t-1)} \right), \quad (3.17)$$

where $\mathbf{x}_j^{(\hat{t})}$ is a *conjugate walker*, randomly selected among all the walkers excluding i , and \hat{t} represents the current state of the j walker:

$$\hat{t} = \begin{cases} t & \text{if } j < i \\ t - 1 & \text{if } j > i \end{cases}. \quad (3.18)$$

³Given the affine transformation $\mathbf{y} = A\mathbf{x} + \mathbf{b}$, where A is a non-singular matrix and \mathbf{b} a scaling vector, the target distribution enjoy the affine invariance property if $\pi_{\mathbf{T}}(\mathbf{y}) = \pi_{\mathbf{T}}(A\mathbf{x} + \mathbf{b}) \propto \pi(\mathbf{x})$.

```

Initialize C chains  $\mathbf{X}^{(0)} = [\mathbf{x}_1^{(0)}, \dots, \mathbf{x}_C^{(0)}]$ 
for  $t = 1$  to  $N$  do
  Create temporary vector  $\mathbf{X}^{\text{tmp}} = \mathbf{X}^{(0)}$ 
  for  $i = 1$  to  $C$  do
    Randomly chose a conjugate walker for  $i$ :  $\mathbf{x}_j = \mathbf{X}_{(j \neq i)}^{\text{tmp}}$ 
    Sample  $z$  according to  $g(z) = 1/\sqrt{z}$ 
    Compute proposal state  $\mathbf{x}^* = \mathbf{x}^j = +z(\mathbf{x}_i^{t-1} - \mathbf{x}_j)$ 
    Compute acceptance probability:  $\alpha^{\text{AIES}} = \alpha^{\text{sym}} z^{d-1}$ 
    Sample  $u$  from uniform distribution  $U(0, 1)$ 
    if  $\alpha \geq u$  then
      Accept state:  $\mathbf{X}^{(t)}(i) = \mathbf{x}^*$ 
      Update temporary vector:  $\mathbf{X}^{\text{tmp}}(i) = \mathbf{x}^*$ 
    else
      Reject state:  $\mathbf{X}^{(t)}(i) = \mathbf{x}^{(t-1)}$ 

```

algorithm 2: Affine Invariant ensemble sampler algorithm

The variable z is randomly drawn from the PDF:

$$g(z) = \frac{1}{\sqrt{z}}, \text{ with } z \text{ in } [1/a, a], \quad (3.19)$$

and $a > 1$ is a tuning parameter, generally set equal to 2 [169]. The proposal state is then accepted with probability:

$$\alpha^{\text{AIES}} = \alpha^{\text{sym}} z^{d-1}. \quad (3.20)$$

The ensemble of walkers respects the detailed balance condition and can be used to approximate the posterior distribution.


One limitation of the algorithm is that the C walkers have to be computed sequentially since each update depends on the current state of all the other walkers. This sequential nature of the algorithm leads to a computational cost equal to $C \cdot t^{\text{alg}}$, where t^{alg} represents the cost of the algorithm if it were run in parallel on C processors.

Example Let us consider the highly stretched Rosenbrock function [170]:

$$\pi(x_1, x_2) \propto \exp\left(-\frac{(1-x_1)^2 + 100(x_2 - x_1^2)^2}{20}\right).$$

We created 20 chains of 10000 samples using the Metropolis algorithm setting $\Sigma = \begin{bmatrix} 0.02 & 0 \\ 0 & 0.02 \end{bmatrix}$ and 20 chains of 10000 samples using the AIES algorithm with $a = 2$. The former chains are shown in the figure below on the left, the

latter on the right.



As one can see the chains generated with a Metropolis algorithm slowly converge to the target distribution, whose analytical joint distribution kernel is plotted on the background with a gray scale. The convergence can be improved by better tuning the covariance matrix. On the other hand, the AIES makes a much better job at approximating the distribution without requiring any tuning process, as one can see in the figure on the right.

3.3.2 Chain convergence

An MCMC chain approximates the posterior distribution of the Bayesian problem when a sufficiently large number of points are sampled. It is then important to have a criterion to establish when a chain is converged.

As discussed earlier, the chains are initialized at arbitrary points, which might be far from the area characterized by the stationary distribution. It is especially true when little information is possessed about the target distribution and we have to use large priors. Consequently, there will be a *statistical transient* during which the chain approaches the stationary regime. It is important to discard the points belonging to the transit prior to using the chain. The fraction of discarded points is called *burn-in*. We revise now three methods to assess the chain convergence after the burn-in phase.

Acceptance rate. The acceptance rate gives a first indication of whether a chain is converged. It is defined as the ratio between the accepted points (points such that $\mathbf{x}^{(t)} \neq \mathbf{x}^{(t-1)}$) and the total number of points. A high acceptance rate suggests that the support has been poorly explored, while low values indicate that most of the proposal lay in a low probability region [162]. An optimal value of 0.234 is targeted. In practice, one avoids falling in the two extreme values of the interval.

Trace plots and histograms. Trace plots represent the evolution of a specific chain dimension over the iterations, while the histograms approximate the target

PDF. By visually inspecting the trace plots, we can determine if the chain mixes well, as shown above in the examples, and if it has converged to the stationary regime, where the chain shows an asymptotic behavior. Additionally, when independent chains, initialized at different points, follow similar paths, it indicates that the target area has been reached. One can also assess that the chain has reached the stationary regime by monitoring the histograms. For example, in the stationary regime, they are characterized by smooth variation, and they exhibit minimal changes with respect to the iteration.

Gelman-Rubin diagnostic. While the above-mentioned methods are qualitative, the Gelmen-Rubin diagnostic [171, 172] is a quantitative method that assesses the convergence by comparing C independent chains. It requires the computation of a *within-sequence variance*, \mathbf{W} , and of a *between-sequence variance*, \mathbf{B} . The former reads:

$$\mathbf{W} = \frac{1}{N} \sum_{i=1}^N \mathbf{W}_i, \quad (3.21)$$

where \mathbf{W}_i is the covariance matrix of each chain i :

$$\mathbf{W}_i = \frac{1}{N} \sum_{t=0}^N (\mathbf{x}_i^{(t)} - \bar{\mathbf{x}}_i) (\mathbf{x}_i^{(t)} - \bar{\mathbf{x}}_i)^T, \quad (3.22)$$

and $\bar{\mathbf{x}}_i$ its mean:

$$\bar{\mathbf{x}}_i = \frac{1}{N+1} \sum_{t=0}^N \mathbf{x}_i^{(t)}. \quad (3.23)$$

The second matrix is defined as:

$$\mathbf{B} = \frac{1}{C-1} \sum_{i=1}^C (\bar{\mathbf{x}}_i - \bar{\bar{\mathbf{x}}}) (\bar{\mathbf{x}}_i - \bar{\bar{\mathbf{x}}})^T, \quad (3.24)$$

where $\bar{\bar{\mathbf{x}}}$ is the mean of the means:

$$\bar{\bar{\mathbf{x}}} = \frac{1}{C} \sum_{i=1}^C \bar{\mathbf{x}}_i. \quad (3.25)$$

At this point, the *multivariate potential scale reduction factor* can be then computed as:

$$R = \frac{N}{N+1} + \frac{C+1}{C} \lambda, \quad (3.26)$$

where λ is the largest eigenvalue of the matrix $\mathbf{W}^{-1} \mathbf{B}$. A chain is well converged for $R \rightarrow 1$.

3.4 Surrogate models

As we saw before, UQ studies, either related to uncertainty propagation or inverse problems, require the evaluation of a large number of model realizations. Hence, they are generally efficiently performed through surrogate models. These are trained on numerical computations and mimic the response in terms of input-output relationship for a reduced computational cost [31]. PCE [151, 152] and Kriging models [29] are widely used for this purpose. In this thesis, we resorted to Kriging models, as they have been proven to be adequate in similar applications [57, 173]. Their formulation dates back to 1951 when Krige [174] proposed to use Gaussian processes for geostatistics applications. The method was later mathematically formalized by Matheron [150] and initially employed to efficiently represent computer experiments by Sacks et al. [29]. Since then, Kriging models have been extensively used over a wide range of applications and numerous techniques have been proposed to enhance their efficiency.

Practical applications might require the solutions of very expensive numerical models. In these cases, even if associated only with the surrogate model training, the total computational cost of the simulations can be high. A cure to this problem is represented by multi-fidelity methods, which allows for lowering the computational effort by leveraging cheaper lower-fidelity representations [30]. These methods are gaining increasing and rapid attention over a wide range of fields [175–177], including aerospace applications [47, 146, 178]. Kriging methods are very suitable to be applied in a multi-fidelity fashion: Co-Kriging [179–181] and Hierarchical Kriging [182] aim at combining different fidelity information to produce an accurate, but efficient, regression. Different fidelity can be obtained by employing models with different physics, from crude approximation to high-fidelity computations. They can also be obtained using differently refined meshes in a multi-level formulation. A review of multi-fidelity methods can be found in Peherstorfer et al. [30].

Efficiency can be also increased by reducing to the minimum the number of training points, and successively enlarging the set by means of adaptive strategies [183]. When dealing with optimization problems, one is interested to exploit the already existing surrogate model, and its associate Gaussian variance, to locally refine the model in points that are candidates to be *minima* or *maxima*. Among these *exploiting* methods, the Expected Improvement (EI), proposed by Jones et al. [184], and its variations [185] are widely used examples. Contrary to this, *exploring* strategies are used to globally increase the overall quality of the surrogate model. Gaussian processes are very suitable for this, as they are characterized by a process variance, which in a nutshell measures how uncertain we are on the prediction on a given untrained point. It thus makes sense to sample where this uncertainty is the

highest [183].

Practical applications might also involve a large number of input variables, which leads to an exponential increase of the number of training points, and of the computational cost, required to obtain an accurate surrogate model. This challenge, commonly referred to as ‘*curse of dimensionality*’, can be mitigated by resorting to dimensional reduction techniques. For instance, one can perform a sensitivity analysis to identify and eliminate the less important variables.

We revise the mathematical formulation of the Kriging in Section 3.4.1, and the one of the hierarchical Kriging in Section 3.4.2. Finally, in Section 3.4.3, we present the adaptive sampling methodology employed in this work to improve the accuracy of the surrogate models.

3.4.1 Kriging model

A Kriging model can be built for each observation y of interest. Let us assume that we have a model $\mathcal{M}(\boldsymbol{x})$ providing the response y as a function of some input quantity \boldsymbol{x} . Let also assume that we have n_t training points ($\mathcal{X} = \{\boldsymbol{x}^{(1)}, \dots, \boldsymbol{x}^{(n_t)}\}$), and the corresponding model realizations ($\mathcal{Y} = \{y^{(1)}, \dots, y^{(n_t)}\}$). A surrogate model can be trained on these pairs to provide the response at an untrained $\boldsymbol{x}^{(*)}$ point for a cheaper price.

The assumption behind the Kriging model is that the response $\mathcal{M}(\boldsymbol{x})$ is the realization of a Gaussian process⁴, such that:

$$\mathcal{M}(\boldsymbol{x}) \approx \mathcal{M}^K(\boldsymbol{x}) = \boldsymbol{\beta}^T \boldsymbol{f}(\boldsymbol{x}) + \sigma^2 Z(\boldsymbol{x}). \quad (3.27)$$

The first term in the RHS is a regression function, referred to as the mean or trend of the process. It captures the main features of the mathematical model $\mathcal{M}(\boldsymbol{x})$ and is characterized by the regression coefficients $\boldsymbol{\beta}$ and the basis functions $\boldsymbol{f}(\boldsymbol{x})$. The second term, $\sigma^2 Z(\boldsymbol{x})$, accounts for the departure of the model’s trend from the real response. The symbol $Z(\boldsymbol{x})$ represents a zero-mean unit-variance stationary Gaussian process with a constant variance σ^2 . Since the Gaussian process has zero-mean, $Z(\boldsymbol{x})$ is fully characterized by its covariance function, $R(\boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)})$. It is also known as correlation or kernel function and quantifies the correlation between each pair of infinite points [187].

Smoothness and stationarity are desirable properties of the Gaussian processes, and they can be obtained by opportunely modeling the covariance function. Stationary Gaussian processes are achieved by employing stationary kernels, which

⁴Rasmussen defines a Gaussian process as a ‘*collection of random variables, any finite number of which have a joint Gaussian distribution*’ [186]

have the form:

$$R(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = R(\mathbf{x}^{(i)} - \mathbf{x}^{(j)}). \quad (3.28)$$

In this case, the correlation is solely dependent on the distance between two points, and not on their specific values, guaranteeing the stationarity of the process. We employed kernels belonging to the Matérn family. The Matérn correlation function is defined as:

$$R(\mathbf{x}^{(i)}, \mathbf{x}^{(j)} | \theta, v) = \frac{1}{2^{v-1} \Gamma(v)} \left(2\sqrt{v} \frac{h}{\theta} \right)^v \mathcal{K}_v \left(2\sqrt{v} \frac{h}{\theta} \right), \quad (3.29)$$

where $h = |\mathbf{x}^{(i)} - \mathbf{x}^{(j)}|$ is the distance between the two points, θ and v respectively represent the length-scale and the shape parameters, and $\Gamma(v)$ and \mathcal{K}_v indicates the Euler's Gamma and the modified Bessel functions [187]. The relative Gaussian process is $\lceil v \rceil - 1$ times differentiable (being $\lceil \cdot \rceil$ the ceiling function). Thus, the parameter v can be used to control the level of smoothness of the process. Half integer values of v are widely used, as this choice leads the kernel to have the form of the product of an exponential function and a polynomial function of order $\lceil v \rceil$. We used a value of $v = 5/2$, which leads to a twice differentiable Gaussian process with a relatively smooth behavior.

In what was discussed above, the kernel was parameterized by means of $h/\theta = |\mathbf{x}^{(i)} - \mathbf{x}^{(j)}|/\theta$ and is thus invariant to all rigid motions. Such kernels are said to be *isotropic*. However, when the behavior of the response cannot be assumed to be the same across the various dimensions, one may use distinct length scales, θ_k , for each dimension. The kernel is said to be *anisotropic*. In this work, we used anisotropic ellipsoidal kernels, in which the scale parameter is computed as:


$$\frac{h}{\theta} = \sqrt{\sum_{k=1}^d \left(\frac{x_k^{(i)} - x_k^{(j)}}{\theta_k} \right)^2}. \quad (3.30)$$

Another important aspect to consider when building the Kriging model is the trend function. One may model it as a known regression function (simple Kriging), as an unknown regression constant (ordinary Kriging), or as a regression polynomial function (universal Kriging). In most cases, an *a priori* estimation of the trend function or of its form is rather complex, and one generally opts for ordinary Kriging, relying on the stationary process to capture the model behavior. Nevertheless, when a model hierarchy can be defined, an estimation of the trend function can be obtained by means of a lower-fidelity representation, as in the hierarchical Kriging formulation, which we will describe in the next section. In this work, we resorted to an ordinary Kriging formulation when no lower-fidelity representation was available. In the following, the unknown regression constant of the ordinary Kriging will be indicated as β_0 .

Example In the ordinary Kriging, the hyperparameters β_0 , σ , and θ_k respectively govern the main value, and the amplitude and frequency of the variation of the Gaussian process realizations with respect to change in \mathbf{x} .

Let us consider the case in which we model the Gaussian process with $\beta_0 = 0$, $\sigma = 1$, and a Matèrn kernel with $\nu = 5/2$ and three different values of θ_i , respectively $\theta_1 = 0.01$, $\theta_2 = 0.1$, and $\theta_3 = 0.5$.

Examples of functions sampled from the Gaussian process $\mathcal{GP}(0, \sigma^2 R(x, x' | \theta_i))$, for the three different values of θ are respectively plotted below with dashed lines. In the same plots, the trend function is plotted with a solid line, and the shadow represents the 95% confidence interval.



As one can see in the left figure, when small values of θ are employed, the points are correlated over a short length, inducing rapid variations in the model realizations. When the value of θ is increased, as in the central figure, the correlation scale enlarges, causing less variation in the process realizations. Ultimately, when the value of θ is very large, points are highly correlated over a large length, resulting in little variations of the response.

Regarding the other two hyperparameters, the effect of using a different β_0 is to shift the mean of realizations toward its value, while the one of increasing/decreasing σ is to stretch/contract the oscillations.

In the example above, no point was constrained to a specific value. Now, let us include in the analysis the n_t training points $(\mathcal{X}, \mathcal{Y})$. According to the Gaussian assumption, they are characterized by a multivariate joint distribution with any untrained point $(\mathbf{x}^{(*)}, y^{(*)})$. Once the Gaussian process is conditioned to the training points, one can compute the posterior mean of the process, $\hat{\mathcal{M}}^K(\mathbf{x}^{(*)})$, at any untrained point, and its standard deviation, $\hat{\sigma}^K(\mathbf{x}^{(*)})$.

When the response is noise-free, the Kriging posterior mean reads:

$$\hat{\mathcal{M}}^K(\mathbf{x}^{(*)}) = \mathbb{E}[Y^{(*)} | \mathcal{Y}] = \boldsymbol{\beta} \mathbf{f}^T(\mathbf{x}^{(*)}) + \mathbf{r}^T(\mathbf{x}^{(*)}) \mathbf{R}^{-1} (\mathcal{Y} - \boldsymbol{\beta} \mathbf{F}), \quad (3.31)$$

where \mathbf{F} is the matrix of the trends computed at the training points, $\mathbf{r}(\mathbf{x}^{(*)})$ is the vector of cross-correlations between $\mathbf{x}^{(*)}$ and each of the training points, and \mathbf{R} is the $n_t \times n_t$ correlation matrix between all the training points. Their components respectively are $F_{ij} = f_j(\mathbf{x}^{(i)})$, $r_i = R(\mathbf{x}^{(*)}, \mathbf{x}^{(i)})$, and $R_{ij} = R(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$. In the case of ordinary Kriging, the functions reduce to $\mathbf{f} = \mathbf{1}$ and $\mathbf{F} = \mathbf{1}$. As one can see in Equation (3.31), the evaluation of the Kriging predictor requires the

inversion of a $n_t \times n_t$ matrix. This represents the bottleneck of the Kriging model, as this computation has a cost of $\mathcal{O}(n_t^3)$, limiting the number of training points to a maximum of 10^4 for practical applications.

One of the reasons that make the Kriging model attractive over other kind of surrogate models is the possibility to estimate the level of trust associated with the model prediction. In fact, one can compute the process' variance as:

$$\begin{aligned} (\hat{\sigma}^K)^2(\mathbf{x}^{(*)}) &= \text{Var}[Y^{(*)}|\mathcal{Y}] \\ &= \sigma^2 \left(1 - \mathbf{r}^T(\mathbf{x}^{(*)}) \mathbf{R}^{-1} \mathbf{r}(\mathbf{x}^{(*)}) + \mathbf{u}^T(\mathbf{x}^{(*)}) (\mathbf{F}^T \mathbf{R}^{-1} \mathbf{F})^{-1} \mathbf{u}(\mathbf{x}^{(*)}) \right) \end{aligned} \quad (3.32)$$


where $\mathbf{u}(\mathbf{x}^{(*)}) = \mathbf{F}^T \mathbf{R}^{-1} \mathbf{r}(\mathbf{x}^{(*)}) - \mathbf{f}(\mathbf{x}^{(*)})$. When the response is noise-free, the variance falls to 0 at the training points.

Example Let us consider the Forrester function [181]:

$$\mathcal{M}(x) = (6x - 2)^2 \sin(12x - 4) \quad \text{with } x \in [0, 1]. \quad (3.33)$$

We trained a Kriging model on six training points, equally spaced between 0 and 1. For the model construction, we used $\beta_0 = 0$, a Matèrn kernel with $\nu = 5/2$ and $\theta_1 = 0.01$, $\theta_2 = 0.1$, and $\theta_3 = 0.5$. The value of σ was naively computed as the standard deviation of the realization points.

In the plots below, the true model's response is plotted with black solid line, and the Kriging predictor with gray solid lines. The shadow areas represent the Kriging uncertainty ($\pm 1.96 \hat{\sigma}^K(x)$). Examples of functions sampled from the joint multivariate distribution are plotted with dashed gray lines.



In the figure on the left, a small value of θ was used: the training points are correlated over a small range and the Kriging predictor approaches their values only when the points are very close. Away from them, the Kriging predictor is dominated by the trend function, resulting in large uncertainty in the prediction. Contrary, when a large value of θ is used, as in the figure on the right, all the points are very correlated, driving down the uncertainty in the prediction, which does not include the real function. The optimum behavior is obtained with a correct value of θ , as shown in the central figure. In this case,

all the points are well correlated, but not too much to cancel the prediction uncertainty.

As highlighted in the above example, the quality of the Kriging response is strongly determined by the choice of the hyperparameters β , σ^2 , and θ_k . Nevertheless, their optimum values are not known *a priori* and they have to be estimated based on the available training points solving an optimization problem. Maximum likelihood estimation or cross-validation methods can be used for this purpose. In the former case, the likelihood of the observed output is maximized:

$$\hat{\theta} = \arg \min \left[\frac{1}{2} \log(\det(\mathbf{R})) + \frac{n_t}{2} \log(2\pi\sigma^2) + \frac{n_t}{2} \right]. \quad (3.34)$$

In the latter case, one first splits the set training points into two subsets: $n_t - n_{cv}$ points are used to train the model, and n_{cv} points are retained to compute the model error. A total of $n_s = n_t + 1 - n_{cv}$ subsets can be created. When $n_{cv} = 1$, n_t pair of subsets can be created (Leave-One-Out (LOO) method). The optimum values of hyperparameters are then estimated by minimizing the error between the training point realizations and the Kriging predictor built on the $n_t - 1$ points:

$$\hat{\theta} = \arg \min \left[\sum_{i=1}^{n_t} \left(\mathcal{M}(\mathbf{x}^{(*)}) - \hat{\mathcal{M}}^K(\mathbf{x}^{(*)}) \right)^2 \right]. \quad (3.35)$$

3.4.2 Hierarchical Kriging

Now, let us assume to have N_l different solvers characterized by distinct fidelity, ranked from the lowest to the highest fidelity. In such a scenario, the Kriging surrogate model can be built in a hierarchical way to reduce to the minimum the numerical effort of sampling solutions from the highest-fidelity model.

The hierarchical Kriging, mathematically formulated by Han [182], proposes to exploit the $l - 1$ fidelity prediction ($\hat{\mathcal{M}}^{K,l-1}(\mathbf{x})$), scaled by a parameter β^l , as trend function for the l fidelity Kriging:

$$\mathcal{M}^{K,l}(\mathbf{x}) = \beta^l \hat{\mathcal{M}}^{K,l-1}(\mathbf{x}) + Z^l(\mathbf{x}). \quad (3.36)$$

The hyperparameter β^l captures the correlation between the $l - 1$ and the l fidelity. The l fidelity predictor, $\hat{\mathcal{M}}^{K,l}(\mathbf{x})$, and the associated variance, $(\hat{\sigma}^{K,l}(\mathbf{x}))^2$, are obtained conditioning the Gaussian process on the n_l training points, according to Equations (3.31) and (3.32), in which $\mathbf{f}(\mathbf{x}^{(*)}) = \hat{\mathcal{M}}^{K,l-1}(\mathbf{x}^{(*)})$, and $\mathbf{F} = [\hat{\mathcal{M}}^{K,l-1}(\mathbf{x}^{(1)}), \dots, \hat{\mathcal{M}}^{K,l-1}(\mathbf{x}^{(n_l)})]$.

The main advantages [176, 182] of the hierarchical Kriging formulation are I) its efficiency and implementation ease: it does not require a cross-correlation


function to be built, as for co-Kriging modeling; being one Kriging evaluated per fidelity, it demands the inversion of N_l correlation matrices of size $n_l \times n_l$ in spite of a bigger $n_t \times n_t$ matrix, which is advantageous for computer performing; II) training points independence, i.e., the set of lower fidelity training points does not require the inclusion of the high-fidelity ones; and III) ease of application of infill techniques, i.e., its Gaussian variance is well suited for adaptive sampling.

Example Let us consider the high-fidelity, \mathcal{M}^{HF} , and low-fidelity functions, \mathcal{M}^{LF} , proposed by Forrester to test the multi-fidelity model [181]:

$$\mathcal{M}^{\text{HF}}(x) = (6x - 2)^2 \sin(12x - 4) \quad \text{with } x \in [0, 1]. \quad (3.37)$$

$$\mathcal{M}^{\text{LF}}(x) = 0.5\mathcal{M}^{\text{HF}}(x) + 10(x - 0.5) - 5 \quad \text{with } x \in [0, 1]. \quad (3.38)$$

We trained three different Kriging models. The first, $\mathcal{M}^{\text{K,LF}}(x)$, has been constructed using ten equally-spaced low-fidelity points. Its response is plotted in orange in the figure below, and it correctly mimics the real low-fidelity function, plotted in red.



The second model, $\mathcal{M}^{\text{K,HF}}(x)$, has been constructed using four high-fidelity points. Its response is plotted in gray. The main trend does not correctly reproduce the real high-fidelity function, plotted in black, and the uncertainty in the prediction is large. The third model, $\mathcal{M}^{\text{K,MF}}(x)$, has been constructed using the same four high-fidelity points, but employing the low-fidelity function as trend function. Its response is plotted in green. As one can see, constructing the model using a multi-fidelity formulation dramatically increases the accuracy of the Kriging prediction.

3.4.3 Adaptive sampling

The accuracy of Kriging models can be increased by enlarging the set of training points through space-filling strategies. If they rely on the current version of the surrogate model, they are said to be adaptive.

When working in a multi-fidelity framework, one is not only interested in sampling in the stochastic space, but also in having a criterion to establish which fidelity level to refine. To this end, we slightly modified the strategy proposed by Zhang et al. [188]. Since the high-fidelity prediction is correlated to the low-fidelity by means of the β factor, the uncertainty in the high-fidelity prediction associated with the low-fidelity model is dependent on the same factor, such that:

$$(\hat{\sigma}^K)^2(\mathbf{x}, l) = \begin{cases} (\beta^{\text{HF}} \hat{\sigma}^{\text{LF}}(\mathbf{x}))^2 & \text{if } l = \text{low-fidelity} \\ (\hat{\sigma}^{\text{HF}}(\mathbf{x}))^2 & \text{if } l = \text{high-fidelity} \end{cases}. \quad (3.39)$$

Zhang's work concerns an optimization problem, thus, the EI was maximized:

$$(\mathbf{x}, l) = \operatorname{argmax} EI(\mathbf{x}, l). \quad (3.40)$$

Contrary to Zhang's work, we want to rely on more than two fidelity models and adaptively explore the high-fidelity one. Following Zhang's work, it makes sense to scale the l fidelity variance, by the products of the l scaling factors that separate it from the high-fidelity one:

$$(\hat{\sigma}^K)^2(\mathbf{x}, l) = \begin{cases} \left(\prod_{j=l+1}^{N_l} (\beta^j)^2 \right) (\hat{\sigma}^l(\mathbf{x}))^2 & \text{if } l \neq \text{high-fidelity} \\ (\hat{\sigma}^{\text{HF}}(\mathbf{x}))^2 & \text{if } l = \text{high-fidelity} \end{cases} \quad (3.41)$$

Furthermore, for exploring purposes, we want to train the model where the Kriging standard deviation is maximum:

$$(\mathbf{x}, l) = \operatorname{argmax} \hat{\sigma}(\mathbf{x}, l) \quad (3.42)$$

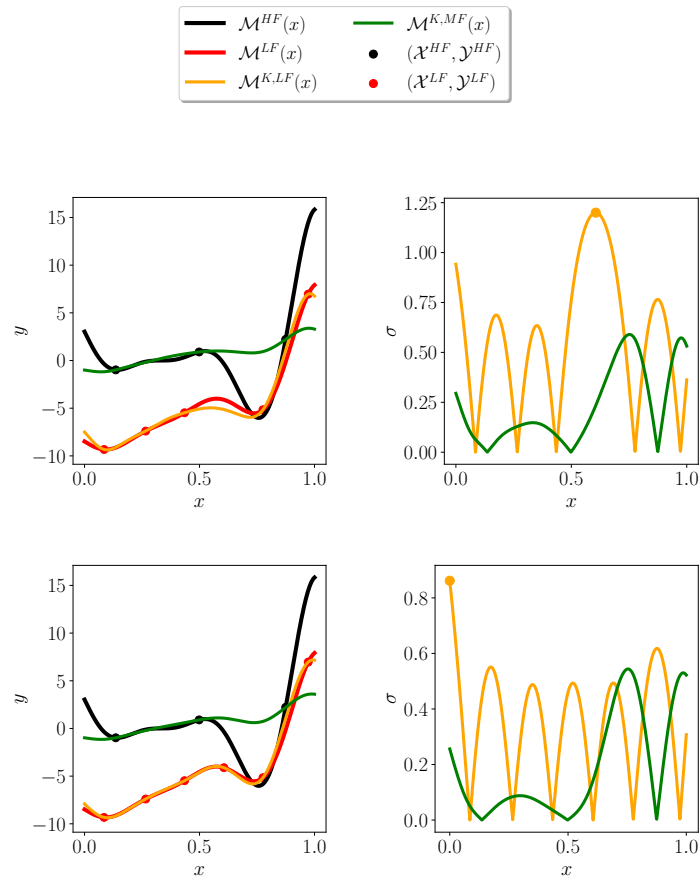
Example Let us consider again the Forrester functions [181]:


$$\mathcal{M}^{\text{HF}}(x) = (6x - 2)^2 \sin(12x - 4) \quad \text{with } x \in [0, 1]. \quad (3.43)$$

$$\mathcal{M}^{\text{LF}}(x) = 0.5\mathcal{M}^{\text{HF}}(x) + 10(x - 0.5) - 5 \quad \text{with } x \in [0, 1]. \quad (3.44)$$

We trained two different Kriging models. The first, $\mathcal{M}^{\text{K,LF}}(x)$, has been constructed using five points. Its response is plotted in orange in the figures below on the left, while its standard deviation, multiplied by $\beta^{\text{K,HF}}$, is shown in the figures on the right. The real low-fidelity function is plotted in red. The second model, $\mathcal{M}^{\text{K,MF}}(x)$, is a multi-fidelity model trained on three training points, using the low-fidelity Kriging predictor as trend function. Its mean and its standard deviation are plotted in the figures below in green. Both sets of training points have been generated using an LHS strategy.

At each iteration, a new training point is selected according to the adaptive sampling methodology. It is highlighted in the standard deviation figures with a circle marker. The predictions at each iteration are plotted consecutively. At the first iteration, the largest Kriging uncertainty is associated with the low-fidelity model, where the first adaptive training point is selected. The process is then iterated until the Kriging predictor well approximates the real function. A total of four points were added to the efficient low-fidelity model, and only one to the expensive high fidelity model to obtain a good representation.





1D Algebraic Verification Test Case

The Forrester function [181] is employed to test the hierarchical Kriging with the adapting sampling strategy above proposed. As seen in the examples above, it consists of two fidelity, namely $\mathcal{M}^\alpha(x)$, and $\mathcal{M}^\delta(x)$. Two further fidelity were added, $\mathcal{M}^\theta(x)$, and $\mathcal{M}^\gamma(x)$, to mimic the four-fidelity form that we intend to use in the application in Chapter 5. Furthermore, we assumed a fictitious computational cost, t_{CPU}^* , in line with the one observed for the corresponding CFD simulations.

The four functions are:

$$\begin{cases} \mathcal{M}^\alpha(x) = (6x - 2)^2 \sin(12x - 4) & t_{\text{CPU}}^* = 1600 \\ \mathcal{M}^\theta(x) = 1.9\mathcal{M}^\alpha(x) + 2.5 & t_{\text{CPU}}^* = 30 \\ \mathcal{M}^\gamma(x) = 4\mathcal{M}^\alpha(x) + 5.5 & t_{\text{CPU}}^* = 4 \\ \mathcal{M}^\delta(x) = 0.5\mathcal{M}^\alpha(x) + 10x - 10 & t_{\text{CPU}}^* = 1 \end{cases} \quad (3.45)$$

A normalized error between the high-fidelity prediction and the real function is computed as:

$$NRMSE = \sqrt{\frac{\sum_{i=0}^{N_v} (\mathcal{M}^\alpha(x_i) - \hat{\mathcal{M}}^\alpha(x_i))^2}{N_v} \frac{100}{\max(\mathcal{M}^\alpha(x_i)) - \min(\mathcal{M}^\alpha(x_i))}},$$

where N_v is the number of validation points, and $\hat{\mathcal{M}}^\alpha(x_i)$ is the multi-fidelity model prediction at the validation points. A total of 1000 equally-spaced points were used for validation. The initial surrogate model was built on six δ training points, and three training points for the remaining fidelity, sampled in the range $[0, 1]$ with a LHS technique; 15 more points were adaptively added. The whole experiment was run 200 times to avoid the verification being influenced by initial *lucky* or *unlucky* points. The results were averaged on the 200 repetitions.




Figure 3.1: Statistical convergence history for adaptive sampling: mean (solid line) plus and minus one standard deviation (shadow area).

The exercise showed that an average of 10.93 points were added in the efficient δ model, and 1.34 and 1.28 on the medium γ and θ fidelity. Only an average of 1.45 points were sampled from the α model. The statistical convergence history averaged on the 200 repetitions, is plotted in Figure 3.1: the normalized error dropped from a value of 19.7% to 0.32%. The difference in performance is evident when we compare this result to the one we would get using the same equivalent computational cost only on high-fidelity LHS points: we could have afforded only an average of around 4.55 LHS training points, which would lead to an $NRMSE = 17.48\%$. To

achieve a comparable error of 0.35% we would need 20 LHS high-fidelity training points, corresponding to a computational cost of more than four times higher.

4D Algebraic Verification Test Case

A second test was performed using a four-dimensional input space and four models characterized by a different fidelity. The high-fidelity function was built to mimic the stagnation-point heat flux computation⁵, while the lower fidelity to imitate the numerical error. Also in this case, we assumed a fictitious computational cost, equal to the one of the previous test. The four functions are:

$$\begin{cases} \mathcal{M}^\alpha(x, y, z, w) = x^3\sqrt{y} + y(0.2z + 0.8w) & t_{\text{CPU}}^* = 1600 \\ \mathcal{M}^\theta(x, y, z, w) = 1.9\mathcal{M}^\alpha + 2.5 & t_{\text{CPU}}^* = 30 \\ \mathcal{M}^\gamma(x, y, z, w) = 4\mathcal{M}^\alpha + 5.5 & t_{\text{CPU}}^* = 4 \\ \mathcal{M}^\delta(x, y, z, w) = 0.5\mathcal{M}^\alpha - 10 + 10x^2\sqrt{y} + y(0.1z + 0.2w) & t_{\text{CPU}}^* = 1 \end{cases}$$

A total of 10^4 LHS points were used for validation. The initial surrogate model was built on six δ training points, and three training points for the remaining fidelity, sampled in the range $[1, 5]$ with an LHS technique; 50 more points were adaptively added. The whole experiment was run 200 times and the results averaged on the 200 repetitions.




Figure 3.2: Statistical convergence history for adaptive sampling: mean (solid line) plus and minus one standard deviation (shadow area).

The exercise showed that an average of 26.335 points were added in the efficient δ model, and 7.3 and 6.865 on the medium γ and θ fidelity. Only an average of 9.5 points are sampled from the high-fidelity model. The statistical convergence

⁵It is the sum of a conductive and a diffusive part. The conductive part can be assumed to be dependent on the cube of the velocity and the square root of the density, while the diffusive contribution can be modeled as linearly dependent on the density and the recombination efficiencies of oxygen and nitrogen: $q = q^{\text{cond}}(v^3, \sqrt{\rho}) + q^{\text{diff}}(\rho_O, \gamma_O, \rho_N, \gamma_N) = q^{\text{cond}}(v^3, \sqrt{\rho}) + q^{\text{diff}}(0.2\rho, \gamma_O, 0.8\rho, \gamma_N)$.

history averaged on the 200 repetitions, is plotted in Figure 3.2: the normalized error dropped from a value of 10.19% to 1.26%. The difference in performance is evident when we compare this result to the one we would get using the same equivalent computational cost only on high-fidelity LHS points: we could have afforded only an average of around 12.73 LHS training points, which would lead to an $NRMSE = 3.31\%$. To achieve a comparable error of 1.27% we would need 28 LHS high-fidelity training points, corresponding to a computational cost of more than two times higher.

3.5 UQLab

All the UQ analyses performed in this thesis were conducted using the UQLab software [189] (Uncertainty Quantification in matLAB). It is an open-source Matlab-based software developed at ETH Zürich with the aim of providing a user-friendly, highly optimized, and extensible package for dealing with non-intrusive stochastic problems. The package includes modules for sampling, training surrogate models (PCE, Kriging, and Polynomial Chaos-Kriging techniques are available), performing sensitivity analysis, classification, and Bayesian inversion. Specifically, we used the modules for the Kriging training [187], implementing the hierarchical Kriging formulation, as outlined in Abdallah et al. [176], and for the Bayesian inversion [162].

3.6 Summary

In this chapter, we revised methods, algorithms, and tools for UQ. We first introduced the forward problem and the MC technique. It is used for computing the PDFs and the statistical moments associated with the output of a model by propagating the uncertainties on its input parameters. An example was used to highlight the differences in the statistical convergence of the method when a LHS sampling strategy, which improves the space filling, is employed over a MC one. Next, we presented the stochastic inverse problem. We shown that Bayesian methods allow for mapping the uncertainties in the model output back to the ones on the model input and that the target posterior distribution can be approximated by an MCMC. We also presented three algorithms to build such a chain: the Metropolis-Hasting algorithm, and its adaptive form, as well as the AIES. The latter improves the quality of the chain when dealing with highly-correlated parameters. After that, we described the formulation behind the Kriging surrogate models. Particularly, we introduced its ordinary and hierarchical form, which alleviates the computational cost associated to its construction by leveraging low-fidelity representations. We also showed how to increase the surrogate model accuracy by resorting to adaptive

sampling strategies. Two test cases were used to highlight the benefits of employing such a strategy. Finally, we introduced the UQLab software that we used in this thesis to perform the UQ studies.

4

Multi-fidelity and multi-level numerical simulations

In this chapter, after a brief introduction to finite volume methods for hypersonic applications, we describe the numerical solvers employed in this work. They are characterized by three different fidelity. The first one is a three-dimensional high-fidelity solver (US3D), the second is an in-house code (stagnation-line code) solving dimensionally-reduce Navier-Stokes equations, and the last is a low-fidelity solver that computes the pressure and the heat flux at the stagnation point of a probe exposed to a supersonic jet. We then present a methodology to estimate the numerical error associated with a given mesh. After that, we delve into the description of the coupling that we performed between the US3D solver and the Mutation⁺⁺ library. The coupling allows us to use Mutation⁺⁺ to compute both the transport and the chemistry gas properties, as well as to solve the GSI balances. We also introduce the test cases performed to verify the coupling. Finally, we present the result of a study that aims at balancing the numerical and the epistemic uncertainties in an atmospheric entry flow.

4.1 Overview

As discussed in the previous chapter, CFD simulations are integrated into the design phase of vehicles traveling at hypersonic speed. Despite progresses, challenges still remain in the accurate prediction of the aero-thermodynamic loads [190, 191].

Besides the complexity of the physical modeling, which was addressed in Chapter 2, difficulties also arise in devising well-suited numerical algorithms. In fact, hypersonic flows are characterized by I) a wide range of time scales, II) a strong discontinuity at the shock layer, which induces numerical instabilities, and III) a stagnation region characterized by strong gradients. These features impose a trade-off in the characteristics of the numerical algorithm. High dissipation is desirable at the discontinuities to produce a stable solution. On the other hand, low dissipation is essential in the BL to well capture the strong gradients developing in front of an object. Such strong gradients further require high grid resolution, leading to heavily stretched cells and limiting the explicit time step employable. It is thus essential to resort to implicit methods. They are also important for resolving all the different time scales [192].

Despite the progress in high-order methods, essential to deal with turbulent and unsteady flows, classical Finite Volume (FV) formulations with second-order limited upwind schemes on hexahedral elements and implicit time-integration still represent the most reliable combination for hypersonic steady-state applications.

An important source of numerical error can derive by the discretization associated with the FV representation of the physical domain. Such an error is inversely proportional to the grid density and can be estimated by using at least three geometrically similar meshes [145, 193]. Such a rigorous estimation is often replaced by the pragmatic practice to prove that the solution is grid-independent, i.e. that the discretization error is almost null. In fact, one verifies that the solutions obtained on a medium and on a fine grid are almost superimposed. While such a practice is customary for deterministic applications, in a UQ framework, it is possible to relax this constrain and to use a grid whose numerical error is lower than the experimental or propagated uncertainty on a given QoI. A rigorous estimation of the mesh numerical error is thus required.

In Section 4.2, we first describe the FV method and its main features for dealing with hypersonic flows, then we present the three different solvers employed in this work. In Section 4.3, the method proposed by Eça [145] to evaluate the numerical error associated with a given mesh is outlined. Section 4.4 is devoted to illustrating the coupling that we have performed between the US3D solver and the Mutation⁺⁺ library. The test cases carried out to verify the coupling are presented in the same section. Finally, in Section 4.5, we present a study in which we aim at balancing the deterministic numerical errors and the stochastic uncertainties. To this end, we performed a systematic quantification of the numerical error and the uncertainty-induced variability for the simulation of an atmospheric entry hypersonic flow. Specifically, a mesh-convergence study using grid adaptation tools is coupled with surrogate-based approaches to UQ.

4.2 CFD Solvers

In this section, we describe the solvers employed in this work. They span different fidelity. The lowest fidelity one is used to compute the pressure and the heat flux at the stagnation point of a probe exposed to a supersonic jet. Conservation laws allow for reconstructing the state of the gas at the exit of a convergent-divergent nozzle, just before the detached shock. From here, the flow is simulated by an in-house CFD solver, the stagnation-line code [194], which represents the medium-fidelity solver. We also resorted to the high-fidelity US3D solver [11]. Both US3D and the stagnation-line code are FV solvers.

When employing a FV formulation, the first step to obtaining the system to solve is to recast the governing equations (2.1), (2.2), and (2.3) into the integral form:

$$\frac{\partial \bar{\mathbf{U}}}{\partial t} + \frac{1}{V} \oint_S (\mathbf{F}^i + \mathbf{F}^v) \cdot \mathbf{n} \, dS = \bar{\mathbf{\Omega}}, \quad (4.1)$$

where the overline $\bar{\cdot}$ indicates a volume-average quantity, V is an arbitrary integration volume, and S represents its boundary with outward normal \mathbf{n} . The vector $\mathbf{U} = [\rho_1, \dots, \rho_{n_s}, \rho \mathbf{u}, \rho E]^T$ contains the conservative variables, $\mathbf{F}^i = [\rho_1 \mathbf{u}, \dots, \rho_{n_s} \mathbf{u}, \rho \mathbf{u} \otimes \mathbf{u} + p \bar{I}, \rho \mathbf{u} H]^T$ the inviscid fluxes, $\mathbf{F}^v = [\mathbf{J}_1, \dots, \mathbf{J}_{n_s}, \bar{\tau}, \bar{\tau} \cdot \mathbf{u} + \mathbf{q}]^T$ the viscous fluxes, and $\mathbf{\Omega} = [\omega_1, \dots, \omega_{n_s} \mathbf{0}, 0]^T$ the source terms.

After that, the physical domain is discretized in n_c cells of volume V_i characterized by n_f faces of surface area $A_{i,j}$. For cell-centered methods, the values of the flow quantities are stored at each cell center. Such a discretization allows for solving the system (4.1) for each cell i by time marching in the discretized time $t^{N+1} = t^N + \Delta t^N$. If the system is solved implicitly, it reads:

$$\delta \mathbf{U}_i^{N+1} + \frac{\Delta t}{V_i} \sum_{j=1}^{n_f} A_{i,j} \mathbf{F}_{i,j}^{N+1} \cdot \mathbf{n}_{i,j} = \Delta t \mathbf{\Omega}_i^{N+1}, \quad (4.2)$$

where, for simplicity, the overline has been dropped, and the inviscid and the viscous fluxes have been grouped, $\mathbf{F} = \mathbf{F}^i + \mathbf{F}^v$. The term $\delta \mathbf{U}_i^{N+1} = \mathbf{U}_i^{N+1} - \mathbf{U}_i^N$ is the update at the iteration $N + 1$. A steady-state condition is achieved for $\delta \mathbf{U}_i^{N+1} \rightarrow 0$. The fluxes and the source terms at the iteration $N + 1$ can be obtained by linearization:

$$\mathbf{F}^{N+1} \approx \mathbf{F}^N + \frac{\partial \mathbf{F}}{\partial \mathbf{U}} \delta \mathbf{U}^{N+1} = \mathbf{F}^N + \mathbf{A} \delta \mathbf{U}^{N+1}, \quad (4.3)$$

$$\mathbf{\Omega}^{N+1} \approx \mathbf{\Omega}^N + \frac{\partial \mathbf{\Omega}}{\partial \mathbf{U}} \delta \mathbf{U}^{N+1} = \mathbf{\Omega}^N + \mathbf{B} \delta \mathbf{U}^{N+1}. \quad (4.4)$$

The next step is to define the numerical fluxes, the source terms, and the Jacobians, \mathbf{A} and \mathbf{B} .

Given the complex structure of hypersonic flows, the computation of the inviscid fluxes is a challenging task [190]. As highlighted in Section 4.1, high dissipation is desirable at the shock layer to enhance the numerical stability, while low dissipation in the BL is requested to accurately compute the wall gradients without resorting to very refined grids. Upwind schemes, such as the modified Steger-Warming [195], are widely used in the hypersonic community. In a general form, the Steger-Warming numerical flux can be computed at each cell interface as:

$$F = F_L^+ + F_R^- = A_L^+ U_L + A_R^- U_R, \quad (4.5)$$

where F_L^+ is the flux entering from the left, whose Jacobian A_L^+ is characterized by positive eigenvalues. Contrary, F_R^- is the flux that exits toward the left, whose Jacobian A_R^- is characterized by negative eigenvalues. The stencil on which the left and right states are computed controls the degree of dissipation of the scheme, hence its numerical accuracy and stability.

The numerical dissipation associated with a constant reconstruction of the states is high. It is a preferable propriety in the shock region, but it leads to inaccurate predictions of the BL gradients. On the other hand, second-order accuracy is achieved by reconstructing the state linearly, for example through the Monotonic Upstream-centered Scheme for Conservation Laws (MUSCL) method, proposed by van Leer [196]. However, such a reconstruction causes numerical oscillations in the shock region. Thus, it is customary to use a second-order reconstruction, limiting it to the more dissipative first-order when a discontinuity is detected.

It is worth noticing that even limiting the reconstruction to the first-order, it may not completely prevent non-physical results across shocks if the mesh is not aligned with the discontinuity. Although manually or automatically aligning the mesh is the best solution, limiting the value of the eigenvalues outside the BL represents a partial cure to the problem [190].

The computation of the viscous fluxes is more straightforward, as they can be expressed as a function of the primitive variable gradients. On structured grids, such gradients can be computed based on grid metrics, whereas, for unstructured meshes, the Green-Gauss and weighted least square approaches are widely used. Their formulation has to be corrected when dealing with high aspect-ratio cells [192], as those in a hypersonic BL.

Source terms are also easy to be computed: they are simply evaluated at the cell center. When using an implicit method, particular attention should be devoted to the computation of the Jacobian, $\partial\Omega/\partial U$, as little numerical errors can severely deteriorate the rate of convergence of the algorithm [192].

4.2.1 High-fidelity: US3D

US3D is a high-fidelity three-dimensional FV CFD solver developed at the University of Minnesota in collaboration with the NASA Ames Research Center [11]. It is specifically designed for high-velocity/high-temperature applications.

One of the main features of US3D lies in the capability to work with unstructured grids, which is advantageous in generating meshes for complex geometry. It supports hexahedral, tetrahedral, prismatic, and pyramidal elements. Hexahedra are the most reliable choice to solve hypersonic flows, as they well capture wall gradients and they are relatively easy to be aligned to the shocks.

Different schemes for the computation of the numerical fluxes and for the time integration are available. For steady-state applications, the customary choice is the modified Steger-Warming scheme [195], with a MUSCL-limited reconstruction [196], and the Data Parallel Line Relaxation (DPLR) [197] implicit time integration. They ensure stable and rapid convergence to the steady-state solution. On the other hand, the time resolution of unsteady problems requires low dissipation. For these cases, unbiased second-, fourth-, and sixth-order numerical fluxes [198] may be used, adding numerical dissipation at discontinuities to prevent numerical instabilities. For the time integration, one may resort to second- or third-order Runge-Kutta scheme or to second-order implicit methods. For the computation of the gradients required for the viscous fluxes, both weight least-squares and Green-Gauss methods are available. A deferred correction approach is used to improve the accuracy of such a computation in highly stretched cells [199].

US3D also offers a tailoring routine to automatically align the grid to the shock, reducing the magnitude of the error generated at the discontinuity. The solution is first obtained on an initial mesh and the routine is run: it detects the shock and the mesh is aligned to it. The solution is then reconverged. Two tailoring steps are generally sufficient to obtain satisfactory accuracy.

Being designed for hypersonic applications, US3D solves multi-species/multi-temperature Navier-Stokes equations. Chemical relaxation is obtained through homogeneous FRC, while non-equilibrium phenomena are captured by means of the Park's two temperatures model with a Landau-Teller relaxation. Thermodynamic properties can be computed either by using the RRHO model or the NASA polynomials. Regarding the transport, various models are available. For example, accurate values of the multi-component viscosity and thermal conductivity are obtained by means of the Gupta-Yos mixture rule along with the Eucken correction for the internal modes. The computation of diffusion velocities relies on the SCEBD model. Finally, turbulence can be modeled using Reynolds-Averaged Navier-Stokes (RANS) or Large-Eddy Simulations (LES).

We finally remark that US3D is highly scalable and efficient. It is then possible to run the simulation in parallel, reducing the time requested to converge. It is an essential quality to perform UQ studies.

4.2.2 Medium-fidelity: stagnation-line

The governing equations (2.1), (2.2), and (2.3), can be simplified to a quasi-1D formulation when dealing with an axisymmetric flow over spherical/cylindrical bodies. Based on the Dimensionally Reduced Navier-Stokes Equations (DRNSE) proposed by Klomfass and Müller [200], Munafò [194, 201] developed the VKI stagnation-line code. It has been extensively used and extended to study stagnation-point ablation, catalysis, evaporation, and flow radiation [90, 99, 119, 128]. It has also been employed to efficiently perform UQ studies regarding the ablative response of materials exposed to Plasmatron conditions [51, 52, 173].

To achieve the dimension reduction, the set of equations is first expressed in a spherical coordinates system (r, θ, ϕ) , where r is the radial coordinate, and θ and ϕ are respectively the zenith and the azimuth angles. Exploiting the problem axisymmetry, the temperature and the mass fractions can be expressed as a function of the only radial coordinate:

$$T = \bar{T}(r), \quad Y_i = \bar{Y}_i(r),$$

while the radial and the azimuthal components of the velocity, as well as the pressure, depend also on the azimuthal angle:

$$u_r = \bar{u}_r(r) \cos \theta, \quad u_\theta = \bar{u}_\theta(r) \sin \theta, \quad \text{and } p - p_\infty = \bar{p}(r) \cos^2 \theta.$$

Then, the DRNSE formulation is obtained by letting $\theta \rightarrow 0$. In discretized form, it reads:

$$\delta \mathbf{U}_i^{N+1} + \frac{\Delta t_i}{\Delta r_i} \sum_{j=1}^2 \mathbf{F}_{i,j}^{N+1} \cdot \mathbf{n}_{i,j} = \Delta t_i \mathbf{\Omega}_i^{N+1} + \Delta t_i \mathbf{S}_i^{N+1}, \quad (4.6)$$

where $\Delta r_i = V_i/A = r_{i+\frac{1}{2}} - r_{i-\frac{1}{2}}$. In the case of DRNSE, $\mathbf{F} = \mathbf{F}^i + \mathbf{F}^v$, with $\mathbf{F}^i = [\rho_i u_r, \rho u_r^2 + p, \rho u_r u_\theta, \rho u_r H]^T$, $\mathbf{F}^v = [J_{i,r}, -\bar{\tau}_{rr}, -\bar{\tau}_{r\theta}, q_r - \bar{\tau}_{rr} u_r]^T$. The vector \mathbf{S} represents the numerical source term that arises from the dimension reduction. It reads:

$$\mathbf{S} = -\frac{1}{r} \begin{bmatrix} 2J_{i,r} \\ 2(\bar{\tau}_{rr} - \bar{\tau}_{\theta\theta} + \bar{\tau}_{r\theta}) \\ 3\bar{\tau}_{r\theta} - \bar{\tau}_{\theta\theta} \\ 2(q_r + \bar{\tau}_{rr} u_r - \bar{\tau}_{r\theta} u_r + \bar{\tau}_{\theta\theta} u_\theta) \end{bmatrix} - \frac{u_r + u_\theta}{r} \begin{bmatrix} 2\rho_i \\ 2\rho u_r \\ 3\rho u_\theta - 2\frac{p-p_\infty}{u_r+u_\theta} \\ 2\rho H \end{bmatrix}.$$

In the stagnation-line code, the inviscid fluxes can be computed either by using the Roe scheme [202], or the AUSM family ones [203], using a MUSCL-limited reconstruction [196]. The viscous fluxes and the source terms are evaluated using primitive variables ($\mathbf{P} = [\rho_i, u_r, u_\theta, T]^T$). Their value at the interface is calculated through a weighted average, while the relative gradients employing finite differences. The DRNSE numerical source term requires the computation of the primitive variable gradients at the cell center, which are also computed with finite differences.

Fast convergence to steady-state is achieved with an implicit method employing a local time step:

$$\Delta t_i = \frac{\text{CFL} \Delta r_i}{(|u_r| + c + \frac{1}{\Delta r} \max(\frac{4}{3} \frac{\mu}{\rho}, \frac{\lambda}{c_v}))_i}, \quad (4.7)$$

where CFL represent the Courant-Friedrichs-Lewy (CFL) number, c is the gas frozen speed of sound, and c_v is the frozen volumetric specific heat at constant volume.

The code is fully coupled to Mutation⁺⁺ which is called for the computation of thermodynamics, chemistry, and transport properties, as well as for the GSI BC. Further details about the stagnation-line code can be found in [119, 201].

4.2.3 Low-fidelity: supersonic nozzle 0D model

A low-fidelity solver was built to compute the pressure and the heat flux at the stagnation point of a probe exposed to a supersonic jet. It enhances the efficiency of the UQ study.

First, we compute the state of the gas just before the detached shock in front of the probe, as sketched in Figure 4.1. The geometrical throat area, together with




Figure 4.1: Sketch of an under-expanded jet over a probe.

the sonic condition at the exit of the nozzle, the conservation of total enthalpy and pressure, and frozen chemistry are imposed to compute the nozzle's choked mass

flow and the gas state at the outlet, for a given reservoir condition. After that, the state of the gas before the shock is computed imposing the previously computed mass flow and the total enthalpy and pressure conservation. The jet diameter before the shock was imposed as a function of the reservoir pressure, whose law was calibrated based on three US3D simulations performed on a relatively coarse mesh. By means of the same computations, it was also inferred a law for the velocity gradient ($\beta = \partial v / \partial y$). These values, together with the gas state, were used as input to the stagnation-line code to compute the stagnation-point pressure and heat flux.

The method was verified by comparing the US3D solution extracted along the stagnation line and the low-fidelity one, obtained with the above-mentioned procedure. The two different solutions are plotted in Figure 4.2. As one can see, the solution is well approximated. Nevertheless, the value of the stagnation pressure, and of the stagnation-point heat flux, turned out to be lower than the high-fidelity ones, probably due to some two-dimensional effects that were not accounted for. However, we remind that we make use of this procedure to compute a good approximation of the QoIs, being the accuracy guaranteed by the high-fidelity simulations. We finally remark that the whole procedure takes negligible time to converge, allowing us to cost-efficiently capture the input-output dependency.




Figure 4.2: Comparison between the high-fidelity solution extracted along the stagnation line (US3D, in solid blue line), and the low-fidelity one (in dashed red line). The solution is well approximated by the low-fidelity solver.

4.3 Mesh numerical error

When conservation laws are solved on a discretized representation of the physical domain, a numerical error is inherited by the solution. This error is a function of the grid density: it approaches zero as the cell dimension approaches the infinitesimal size. Several methodologies for estimating such a numerical error have been proposed in the literature. Historically, the first method was proposed by Richardson [204]: the numerical error, and, thus, the extrapolated solution, can be

characterized using the solution of two nested meshes. This method assumes that the order of convergence of the numerical scheme is known and respected. To get rid of this assumption, Roache [193] proposed to use three grids for evaluating also the observed order of convergence. More recently, Eça [145] developed a procedure that, unlike the first two methods, can be applied even outside the asymptotic range of spatial convergence. In this thesis, the numerical uncertainties associated with the used meshes were computed according to the procedure proposed by Eça [145]. Such a procedure is outlined in the following.

The numerical error of a given grid can be estimated as:

$$\epsilon_\phi \approx \delta_{\text{RE}} = \phi_i - \phi_0 = \alpha h_i^p, \quad (4.8)$$

where the quantity ϕ_i is a local or integral flow QoI at the grid refinement i , ϕ_0 represents the estimate of the exact solution, α a constant to be determined, h_i the typical cell size and p the observed order of grid convergence. In most applications, the exact solution of the QoI cannot be analytically computed. In these cases, it can be estimated as the asymptotic limit for the element size approaching the infinitesimal value. Hence, the QoI at the refinement i can be approximated as:

$$\phi_i = \phi_0 + \alpha h_i^p. \quad (4.9)$$

By using a minimum of four nested meshes, the values α , p , and ϕ_0 can be estimated through a least-squares methodology. A positive value of the order of convergence, p , implies a monotonic convergence, in contrast, a negative value a monotonic divergence.

The assumptions behind the use of Equations (4.8) and (4.9) are that the grids are I) in the asymptotic range, such that we can describe the error only through the highest-order term, and II) geometrically similar, such that we can represent the grid density solely by means of the quantity h_i . However, these two assumptions are not always met in practical applications. For instance, in hypersonic simulations, the employed limiters induce scatter in the data. In these cases, the use of Equation (4.9) may lead to poor fits and not reliable estimates, and a fixed order of convergence in the fit law is suggested, such as:

$$\phi_i = \phi_0 + \alpha h_i, \quad (4.10)$$

$$\phi_i = \phi_0 + \alpha h_i^2, \quad (4.11)$$

when the convergence behavior is monotonic, or

$$\phi_i = \phi_0 + \alpha_1 h_i + \alpha_2 h_i^2, \quad (4.12)$$

when it is not.

The procedure also allows for determining the uncertainty on the prediction associated with the mesh at the refinement i . Unlike the grid convergence index [193], such an evaluation also accounts for how well the data agree with the fit. The mesh numerical uncertainty reads:

$$U_\phi(\phi_i) = \begin{cases} F_s \epsilon_\phi(\phi_i) + \sigma + |\phi_i - \phi_{\text{fit}}|, & \text{if } \sigma < \Delta_\phi \\ 3 \frac{\sigma}{\Delta_\phi} (\epsilon_\phi(\phi_i) + \sigma + |\phi_i - \phi_{\text{fit}}|), & \text{otherwise} \end{cases} \quad (4.13)$$

where σ is the fit standard deviation, $\Delta_\phi = (\phi_i^{\text{max}} - \phi_i^{\text{min}})/(n_g - 1)$ is a data range parameter, n_g is the number of grids employed to perform the convergence study, and F_s is a safety factor equal to 1.25 if the error estimate is considered reliable, i.e. if $0.5 < p < 2.1$ and $\sigma < \Delta_\phi$, and to 3 otherwise.

4.4 US3D-Mutation⁺⁺ coupling

The coupling between the US3D solver and the Mutation⁺⁺ library was propaedeutic to perform the UQ studies as it allows us to modify the simulation model parameters through Mutation⁺⁺. In this section, we first describe the coupling, then we present the test cases carried out for verification purposes.

4.4.1 Description

The US3D solver was coupled to the Mutation⁺⁺ library by means of the code interface, which allows for expanding the code's capabilities without accessing its source code. Specifically, we wrote an interface to call the Mutation⁺⁺'s modules for I) viscosity and thermal conductivity, II) diffusion III) gas chemistry, and IV) GSI.

The interface works as follows. First, the start-up utility is invoked to create a Mutation⁺⁺ object, specifying the mixture, the mechanism, and the GSI files. At this point, Mutation⁺⁺ can be employed for the following routines:

The results of this section were included in the following publications:

- Baskaya, A., Capriati, M., Ninni, D., Bonelli, F., Pascazio, G., Turchi, A., Magin, T., Hickel, S., **Verification and Validation of Immersed Boundary Solvers for Hypersonic Flows with Gas-Surface Interactions**, AIAA AVIATION, 2022, Chicago, IL & Virtual.
- Capriati, M., Prata, K., Schwartzentruber, T., Candler, G., Magin, T. **Development of a nitridation gas-surface boundary condition for high-fidelity hypersonic simulations**, ECCOMAS, 2021, Virtual.

I) Computation of the viscosity and the thermal conductivity

Input of this function are the mass fractions, the density, and the temperature(s) of the gas. These quantities are rearranged to set the gas state in Mutation⁺⁺ through the partial densities and the temperature(s). After that, Mutation⁺⁺ is called to get the gas viscosity and the components of the thermal conductivity, which are then cast in an array according to the thermodynamic model employed in the simulation. These values represent the output of the function.

II) Computation of the diffusive fluxes

Input of this function is the number of the face where to compute the diffusive fluxes and the indices of the left and the right cells. First, the state at the interface is computed as average between the left and the right states. This is used to set the gas state in Mutation⁺⁺. Then, the molar fractions gradient at the interface is computed. It is approximated with finite difference, instead of using a multi-dimensional reconstruction approach, such as Green-Gauss or weighted least square, as we have access only to the left and the right variables. It is passed to Mutation⁺⁺ to obtain the diffusion velocities according to the Stefan-Maxwell formulation. Finally, the fluxes at the interface are computed as:

$$\text{Species } i \text{ diffusive flux} = V_i \rho_i A_{\text{interface}},$$

$$\text{Energy diffusive flux} = \sum_{i=1}^{n_s} (\text{Species } i \text{ diffusive flux}) \cdot h_i,$$

where V_i is the species i diffusion velocity, and $A_{\text{interface}}$ is the area of the interface. The Jacobian of the diffusive fluxes is analytically obtained from the SCEBD formulation as a function of the effective binary diffusion coefficients, which are also computed by means of Mutation⁺⁺. The diffusive fluxes and the relative Jacobian are the output of the function.

III) Computation of the gas chemistry source term

The input of this function is the number of the cell where to compute the chemical source term. US3D is called to get the values of the cell mass fractions, density, and temperatures, which are then used to set the gas state in Mutation⁺⁺. After that, Mutation⁺⁺ is called to get the chemical source terms and the analytical Jacobian, $\frac{\partial \dot{\omega}_i}{\partial \rho_j}$. The full Jacobian, $\frac{\partial \dot{\omega}_i}{\partial U_j}$, is obtained by adding the term $\frac{\partial \dot{\omega}_i}{\partial T} \frac{\partial T}{\partial U_j}$. The first term, $\frac{\partial \dot{\omega}_i}{\partial T}$, is the derivative of the source term with respect to the temperature and is computed numerically. The second term, $\frac{\partial T}{\partial U_j}$, represents the derivative of the temperature with respect to the conservative variables. It is computed analytically. The expression of the Jacobian can be found in [90]. The output of the function is the gas chemistry source term and the relative Jacobian.

IV) Gas surface interaction boundary condition

This function allows for solving both the SMB and the SEB BCs, as described in Section 2.3. A loop over all the faces belonging to the specific BC is performed. For each face, the distance between its center and the center of the cell next to it is computed. Such a distance is passed, together with the molar fractions of the cell, to Mutation⁺⁺ to compute the diffusive terms. In the only case of the SEB BC, also the temperature of the cell is passed to Mutation⁺⁺ to compute the conductive heat flux in the energy balance. A tentative surface state is passed to Mutation⁺⁺ to solve the GSI system. After that, the surface state is obtained. Finally, the solution is linearly extrapolated to the ghost cells, imposing zero pressure gradient in the normal direction.

Once the simulation is concluded, the finalizing utility is called to destroy the Mutation⁺⁺ object.

We remark that the US3D-Mutation⁺⁺ coupling results in a higher computational cost compared to US3D with its native modules. This increase is primarily due to the solution of more expensive, but more accurate, physical models, such as solving the Stefan-Maxwell system for computing diffusive velocities, as opposed to employing the SCEBD.

4.4.2 Verification test-cases

A series of test cases were performed to verify the coupling between Mutation⁺⁺ and US3D. The simulation results were compared against several reference sources depending on the specific test case. These include US3D with its native library, the stagnation-line code (which has been extensively used and validated coupled to Mutation⁺⁺), experimental data, and analytical formulations. A summary of the performed test cases can be found in Table 4.1. The first three cases verified single

Section	Test	Assessed Features
A.1	0D Reactor	Chemistry
A.2	1D Diffusion Problem	Transport
A.3	1D Catalytic Diffusion Problem	Diffusion and Catalysis
4.4.2	2D hypersonic flow over a cylinder	Chemistry, Transport
A.4	Axisymmetric flow over an ablator	Chemistry, Transport, GSI

Table 4.1: Summary of the verification test cases.

features of the coupling, i.e. chemistry, transport properties, and the catalytic BC. The last two test cases are more general: the first one is a hypersonic flow over a cylinder, and the second one is a subsonic nitrogen plasma flow over an ablative

sample. The former test case is presented in the following, while the others can be found in Appendix A.

Hypersonic air flow over an isothermal cylinder

Knight et al. [205] employed a hypersonic flow over a cylinder as test case to assess the capabilities of various CFD solvers, comparing the numerical solutions to the experimental data. We used the same test case to verify the coupling between Mutation⁺⁺ and US3D.

The mixture considered is a five-species air mixture: $S = [\text{N}_2, \text{O}_2, \text{NO}, \text{N}, \text{O}]$. Homogeneous FRC is modeled based on the reaction rates from Park et al. [79]. At the inlet, a supersonic inflow was prescribed, with free-stream values reported in Table 4.2. The inlet mass fractions are those at LTE. On the cylinder, an inert isothermal surface BC with $T_w = 300$ K was applied.

M_∞	u_∞ [m/s]	T_∞ [K]	p_∞ [Pa]	ρ_∞ [kg/m ³]
8.98	5956	901	476	1.547×10^{-3}

Table 4.2: Free-stream conditions for the 2D hypersonic flow over a probe.

Two US3D simulations were run. One made use of the US3D native thermo-physical library (below indicated as ‘US3D native’). The second employed the coupled Mutation⁺⁺ library (below indicated as ‘US3D+M⁺⁺’). For the native US3D computation, diffusion velocities were computed according to the SCEBD formulation, while transport properties using the Gupta-Yos mixture rule. For the US3D+M⁺⁺ computation, the Stefan-Maxwell formulation was employed to compute diffusive velocities, while Chapman-Enskog expansion for the transport properties. Both simulations used NASA-9 polynomials to calculate thermodynamic properties under the assumption of thermal equilibrium.

The Mach number contours predicted using the two libraries are shown in Figure 4.3. No appreciable differences are observed between the two solutions. From the flowfield, we extracted the temperature and mass fractions profile along the stagnation line. The two solutions are illustrated in Figure 4.4. As one can see, the post-shock thermo-chemical relaxation shows an excellent agreement. However, slight discrepancies are observed in the BL, where the chemical relaxation exhibits a different behavior. Upon different tests, it was noticed that such a difference emerges only in the case of extremely refined mesh, where the value of the temperature in the first cells drops below 500 K. In these cases, the oxygen recombination predicted by the native library is lower than the one predicted by Mutation⁺⁺. Finally, the pressure and the heat flux distributions along the cylinder were extracted. They are




Figure 4.3: Mach field: US3D+M⁺⁺ solution on the top, US3D native solution on the bottom. The two solutions well agree with each other.

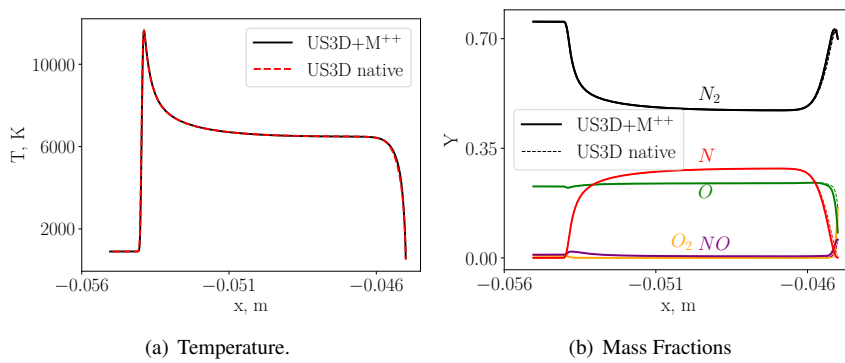


Figure 4.4: Comparison of the QoIs along the stagnation line: US3D+M⁺⁺ solution with solid line and US3D native solution with dashed line. The post-shock relaxation shows an excellent agreement, while slight discrepancies are observed in the BL.

plotted in Figure 4.5 along with the experimental data. The pressure distribution exhibits an excellent agreement between the two codes and with the experimental data. On the other hand, differences in the prediction of the heat flux are present, with the one obtained using the Mutation⁺⁺ library being more in agreement with the experimental data. Such a difference in the prediction is driven by the above-mentioned difference in the chemistry. In fact, if we compute the chemistry with the native library of US3D and the transport properties with Mutation⁺⁺ (referred to as ‘US3D chem’ in Figure 4.5(b)), we obtain an excellent agreement with the solution provided by US3D with its own native library.

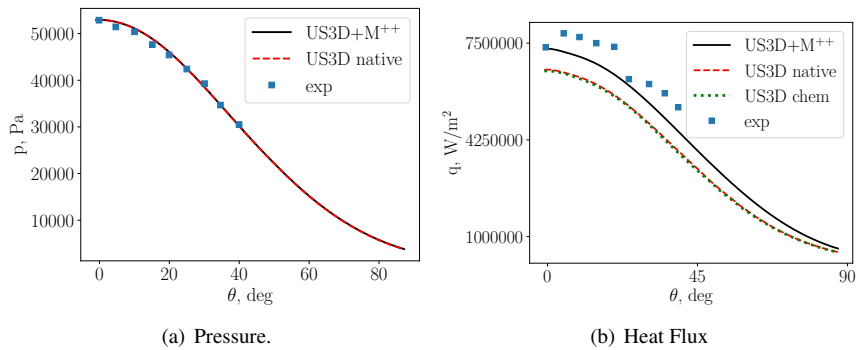


Figure 4.5: Comparison of the QoIs around the surface: US3D+M⁺⁺ solution with solid line and US3D native solution with dashed line. The experimental values are plotted with blue squares. The two pressure distributions well agree with each other and with the experimental data. The two heat flux distributions show some discrepancies. The US3D+M⁺⁺ better predict the experimental data.

4.5 Balancing numerical and epistemic uncertainties in an atmospheric entry flow

The numerical simulation of hypersonic flows is a challenging problem, especially when one aims to obtain reliable predictions of the thermal loads. In this context, appropriate meshes are essential to achieve accurate results [191, 206–212]. For example, the computational grid must be aligned to the shock, preferably with highly stretched cells in the shock region. This practice avoids the appearance of numerical errors (mostly in the form of non-physical vortices), which, advected downstream,

The result of this section are included in the article:

- Capriati, M., Cortesi, A., Magin, T., Congedo, P. **Stagnation point heat flux characterization under numerical error and boundary conditions uncertainty**. European Journal of Mechanics / B Fluids, 2022.

corrupts the correct evaluation of the QoIs in the stagnation region. When a limited number of simulations are needed, a high-quality computational grid can be obtained by monitoring the solution and manually adapting the grid [191, 213].

On the other hand, when performing UQ studies, one has to perform a large number of simulations. It is common practice to optimize the mesh on the nominal conditions and to use it to perform the required computations. While this approach is usually robust in the absence of shock waves [214, 215], problems can arise when dealing with hypersonic flows. In this case, each computation might be characterized by different BCs leading to different shock positions. For example, in Tryoen et al. [55], a fixed computational grid was used to perform all the simulations required for training a surrogate model, yielding difficulties in its construction. In such a scenario, each computation needs an aligned mesh. Manually adjusting the grids is not feasible in UQ studies and it is preferable to resort to automatized tools to align the grid and to avoid that numerical errors corrupt the surrogate model training.

Besides numerical errors associated with poor mesh-shock alignment, the numerical solution is always depend on the mesh refinement. In deterministic simulations, a grid-independent solution is typically sought. However, this practice might be not necessary in a UQ framework, where using a very efficient numerical representation would be advantageous to minimize the computational effort in training the surrogate model. One might employ a grid characterized by a numerical uncertainty on a QoI which, rather than approaching zero, is only lower than the propagated uncertainty on the same quantity. In fact, a grid characterized by a vanishing numerical uncertainty may be a waste of computational budget when the variability induced by the BCs uncertainties still dominate the prediction.

Following, we propose a methodology for balancing the numerical and the epistemic uncertainties on the prediction of the surface pressure and heat flux of a hypersonic vehicle. We systematically employed grid adaptation tools to automatically prevent mesh-shock misalignment, guaranteeing the robustness of the predictions without manual intervention. The benefits of adopting this strategy are investigated by comparing the grid-aligned results against those obtained on the nominal mesh. Furthermore, we performed a rigorous grid convergence study to evaluate the numerical uncertainty expected by using a given mesh. The uncertainties in the BC were also propagated to estimate the variability of the QoIs. The latter was then compared to the numerical uncertainty to choose the most efficient mesh.

4.5.1 EXPERT atmospheric entry

We selected the EXPERT vehicle [216, 217] entry flow as case study. The free-stream nominal conditions, and associated uncertainties, were taken from the study of Tryoen et al. [55], where difficulties were highlighted in the surrogate model construction. The values are reported on Table 4.3. The surface is assumed to be catalytic; the effective reaction efficiency uncertainty is also given in the same table.

Variable	Nominal value	Distribution	Minimum	Maximum
ρ_∞ , kg/m ³	2.88×10^{-4}	Uniform	2.30×10^{-4}	3.46×10^{-4}
u_∞ , m/s	4868.6	Uniform	3985.8	5842.3
γ	0.0015	Uniform	0.001	0.002

Table 4.3: Nominal conditions and associated uniform uncertainties for free-stream density, ρ_∞ , free-stream velocity, u_∞ , and recombination probability, γ .

We described the gas as an air mixture of five species, $S = [\text{N}_2, \text{O}_2, \text{NO}, \text{N}, \text{O}]$, whose homogeneous FRC is described by means of the Park rates [79]. Diffusive fluxes were computed according to the Stefan-Maxwell formulation, viscosity and translation thermal conductivity by means of Chapman-Enskog expansion, and the internal thermal conductivity with the Eucken correction. The catalytic surface is modeled by solving a surface mass/energy balance through the Mutation⁺⁺ library. Inviscid fluxes are approximated by the modified Steger-Warming scheme with a MUSCL-limited reconstruction. The solution converged to steady-state with a first-order DPLR time integration. All the simulations residual dropped by 8 orders of magnitude.

In the studied conditions (Mach number around 15), a strong bow shock develops in front of the vehicle, as plotted in Figure 4.6. The consequent increase in temperature, Figure 4.7, drives the dissociation of the molecules into atoms. These diffuse through the BL and partially recombine because of the BL's temperature drop and the catalytic activity of the vehicle surface.

4.5.2 Numerical uncertainty on deterministic simulations

We built four geometrically similar meshes to assess the spatial convergence and error of the QoIs. The starting nominal computational grid contains 39x41 nodes; the other three finer grids were obtained by progressively doubling the cells in both directions. The set of four nested grids is reported in Table 4.4. The normalized characteristic length, h_i/h_1 , the wall-normal distance of the first cell at the stagnation point, Δn , and the average time for 5.000 iterations to be performed on 16




Figure 4.6: Mach contour at nominal condition: strong bow shock develops in front of the vehicle.




Figure 4.7: Temperature (left) and mass fractions (right) values along the stagnation line at nominal conditions. The jump in temperature drives the dissociation of the molecules. The resulting atoms diffuse through the BL and partially recombine because of the surface's temperature drop and the catalytic activity.

cores are provided in the same table. The value of the sonic Reynolds number in the stagnation-point cell is also reported; it is defined as:

$$Re_c = \frac{\Delta n \rho c}{\mu_w}, \quad (4.14)$$

where c is the wall sonic velocity, and μ_w the viscosity. As a general rule of thumb, hypersonic heat fluxes are assumed to be well computed using a first normal spacing of 10^{-6} m; anyway it was shown that this value is affected, for example, by the

surface temperature [212]. A better criterion of convergence is the sonic Reynolds number, which should be kept below the unity [206, 207].

Mesh	Nodes	Cells	h_i/h_1	Δn , m	Re_c	Time, s
I	305x321	97280	1	$1.25 \cdot 10^{-6}$	0.2	≈ 2490
II	153x161	24320	2	$2.50 \cdot 10^{-6}$	0.4	≈ 675
III	77x81	6080	4	$5.00 \cdot 10^{-6}$	0.8	≈ 205
IV	39x41	1520	8	$1.00 \cdot 10^{-5}$	1.6	≈ 90

Table 4.4: Numerical grids used in the study: tag of the mesh, number of nodes, number of cells, normalized characteristic length, h_i/h_1 , wall-normal distance of the first cell at the stagnation point, Δn , sonic Reynolds number based on the first physical cell, and time required to perform 5.000 iterations on 16 cores.

The mesh III and the mesh IV are shown, respectively, in Figure 4.8(a) and 4.8(b): exploiting the axisymmetry of the flow, the 3D problem was reduced to a 2D configuration; further reduction of the computational cost was achieved by simulating only the half part of the 2D domain.

As above mentioned, it is essential to align the mesh to the shock. For this reason, we used US3D tailoring routine [11]: once the simulation is converged on a nominal grid, the routine computes the position of the shock and the numerical grid is aligned. The simulation is then converged on the new grid. As it can be appreciated in Figure 4.8(e), the tailoring tool allows for capturing a much less diffuse shock, improving the after-shock flow predictions. An example of tailored mesh is shown in Figure 4.8(c).

The angular distribution of pressure and heat flux are, respectively, plotted in Figure 4.9(a) and 4.9(b). One can see that most of the numerical error concentrates around the stagnation point, where the QoIs reach their maximum value. Their grid dependency was investigated both as stagnation-point value and as integrated value along the surface.

The numerical uncertainty for each grid in Table 4.4 was computed following the procedure proposed by Eça [145], presented in Section 4.3, using the nominal condition reported in Table 4.3. Equation (4.9) was used to compute the numerical error relative to the stagnation and integrated pressure (force), as these values fall in the monotonically convergent range. The normalized values ($100 \cdot \left| \frac{\phi_i - \phi_0}{\phi_0} \right|$) are shown respectively in Figure 4.10(a) and 4.10(c): the numerical error decreases as the mesh is refined, with a slope proportional to the observed order of convergence, whose values are given on Table 4.5. In this case, the grid adaptation tool improves the orders of convergence, and, thus, the rate of reduction of the numerical error. These values turned out to be lower than the theoretical value of 2, expected from




Figure 4.8: Top: Numerical grids used in the study. Bottom: zoom of the Mach contour on the mesh IV.

the linear reconstruction. Two main reasons compete in the corruption of the order:

- I) to prevent numerical oscillations from spreading, limiters are generally employed to decrease the order of the linear reconstruction when a discontinuity, such as a shock, is detected. In this case, a mix of first and second-order spatial accuracy dominates the flowfield. This behavior was also observed by Roy [218] who investigated the spatial convergence of the surface pressure in a Mach 8 flow over a blunt body.
- II) The GSI module of Mutation⁺⁺, used for solving the mass-energy balance BC, is first-order accurate, enforcing the first-second order mix.

Unlike these properties, stagnation-point heat flux exhibits a non-monotonic con-




Figure 4.9: Angular distribution of the QoIs for the four meshes in Table 4.4 relative to the nominal conditions reported in Table 4.3. The solution obtained using the adapted mesh is plotted with dashed line, the one with a non-adapted mesh with solid line. The numerical error concentrates around the stagnation point, where the QoIs reach their maximum value.

	p_w [Pa]	F [N]	Q [W]
NA	0.74	0.34	1.60
A	1.63	1.15	1.25

Table 4.5: Orders of convergence on non-adapted (NA) and adapted (A) grids for the pressure, p_w , the force, F , and the heat load, Q .

vergence and Equation (4.9) cannot be applied. We employed the polynomial fit in Equation (4.12). Its use is also justified by the reasons mentioned above: the error is a mix of first and second-order errors. Similarly to the other two properties, the fitting is improved when the adaptation tool is used, as one can see from Figure 4.10(b). It is interesting to note that, by contrast, the surface integrated value of the heat flux is in the monotonic convergence range, as one can see in Figure 4.10(d).

Equation (4.13) allows for evaluating the numerical uncertainty related to the use of a given grid, as shown in Figure 4.11. It can be noticed that the methodology is very conservative: the uncertainty bars of a coarser mesh contain the ones of a finer one. The only exception is the heat flux uncertainty bar of the coarsest mesh, which does not contain the other three due to the use of the polynomial fit. A second significant effect of adapting the mesh can also be noticed: the uncertainty associated with each property decreases as a consequence of the better fitting and order of convergence.

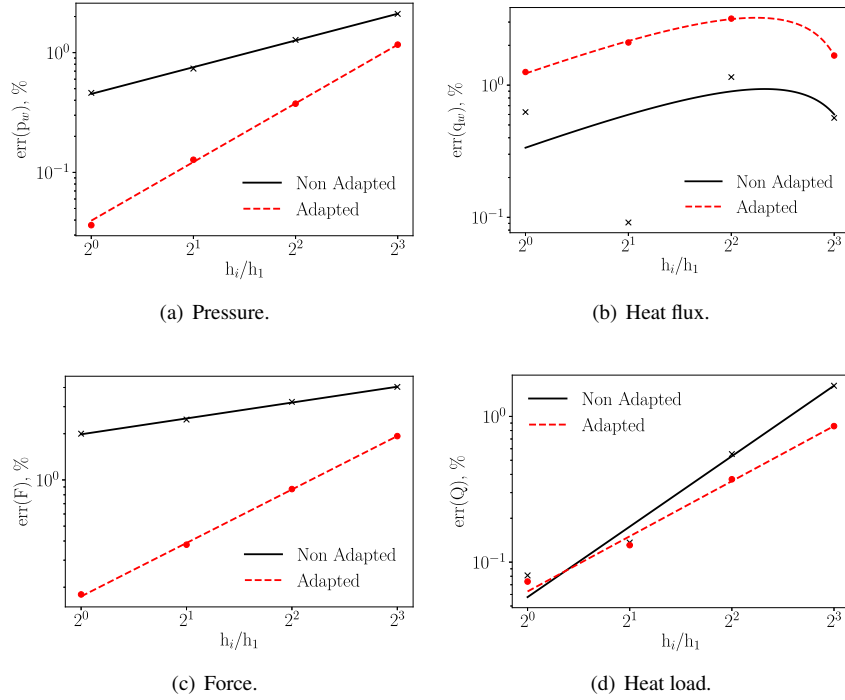


Figure 4.10: Relative error on the $QoIs$ as a function of the degree of refinement. The pressure, the force, and the heat load are in the monotonic convergence range. The observed order of convergence increases by employing grid adaption tools for the pressure and the force. The heat flux is outside the monotonic convergence range and the mesh adaption improves the fit.

4.5.3 Surrogate model construction and forward propagation

We estimated the variability of the stagnation-point pressure and heat flux by propagating the uncertainties on the free-stream density and velocity, and the recombination probability of the surface of the EXPERT vehicle. For simplicity, we neglected the uncertainty on the atmospheric chemistry model considered in Tryoen et al. [55]. Uniform uncertainties were chosen, namely of $\pm 20\%$ intervals around the nominal values for the density and the velocity, and a $[0.001; 0.002]$ interval for the recombination probability. The complete list of uncertainties is provided in Table 4.3.

We built a Kriging surrogate model by means of the UQLab software [189]. Specifically, we generated $N_t = 80$ points with a Sobol technique [159] to train the surrogate model, and $N_v = 20$ with a LHS strategy [158] for verification pur-

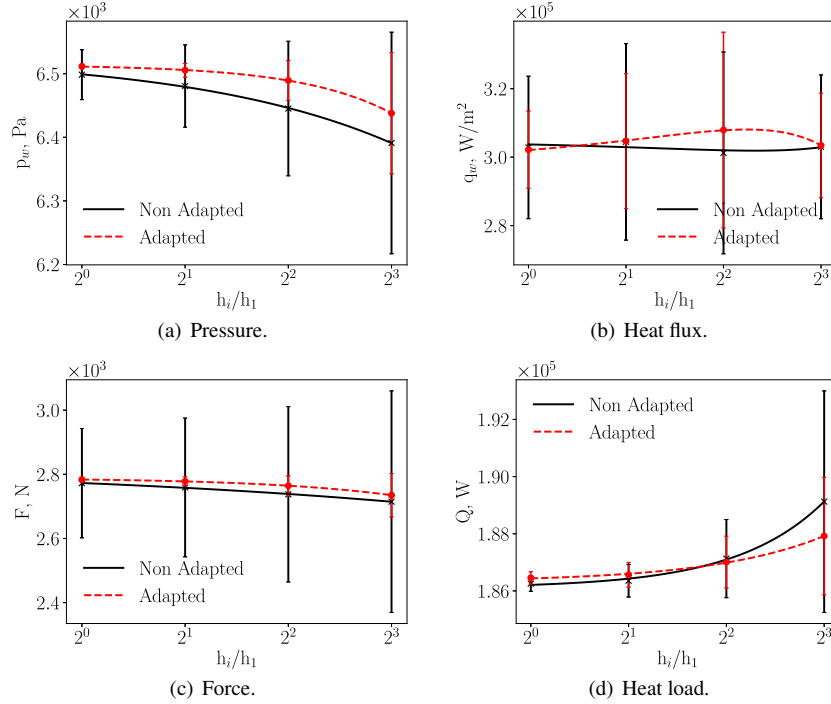


Figure 4.11: QoIs and relative numerical uncertainty as a function of the degree of refinement. The methodology is conservative: the uncertainty bars of a coarser grid consistently contain the finer grid ones; the only exception is the heat flux due to the used polynomial fit. Adaptation tools systematically reduce the numerical uncertainty.

poses. We decided to pay the price of performing a CFD simulation for each training/verification point on each mesh to assess the convergence of the statistical moments (mean and standard deviation). The stagnation-point pressure and heat flux were extracted from each solution and used to train the relative surrogate model for each grid.

The projections of the surrogate models obtained with the adapted mesh II are shown in Figure 4.12. By visual inspection, no non-physical predictions were observed. It is interesting to observe the surrogate model built using the not-adapted mesh II, shown in Figure 4.13(a). For the upper limit of free-stream density and velocity, an overshoot in heat flux is evident: it is a consequence of the poor alignment of the grid to the shock. In particular, in Figure 4.13(b), we can see how grid adaptation improves the solution, leading to physically sound values.

Non-physical behaviors were thus not observed in the surrogate models ob-

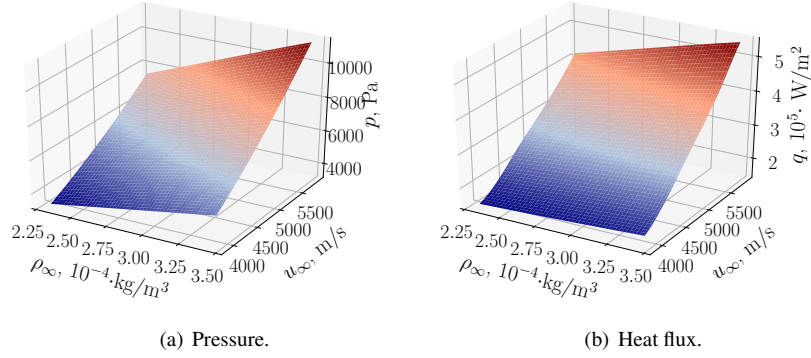


Figure 4.12: Projections of the surrogate models obtained for the adapted mesh II: no non-physical values are observed.

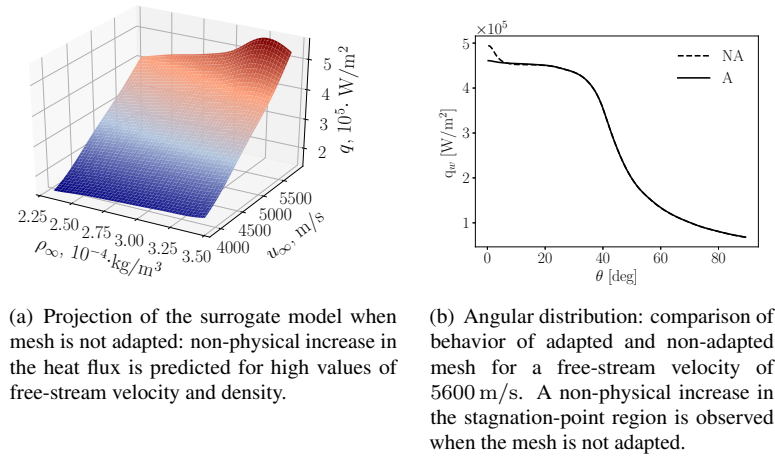


Figure 4.13: Heat flux obtained using mesh II.

tained using adapted grids.

Furthermore, the values returned by the surrogate models were compared with those obtained by the CFD on the verification points to assess the model quality. As the QQplots illustrate in Figure 4.14, the surrogate models performs well, as also confirmed by the verification errors, consistently below 10^{-5} .

Once verified, the surrogate models were used to compute mean (hereafter indicated with a $\hat{\cdot}$) and standard deviation ($\sigma(\cdot)$) of the QoIs, propagating the uncertainties in the BCs by means of a MC method. The mean and the standard deviation

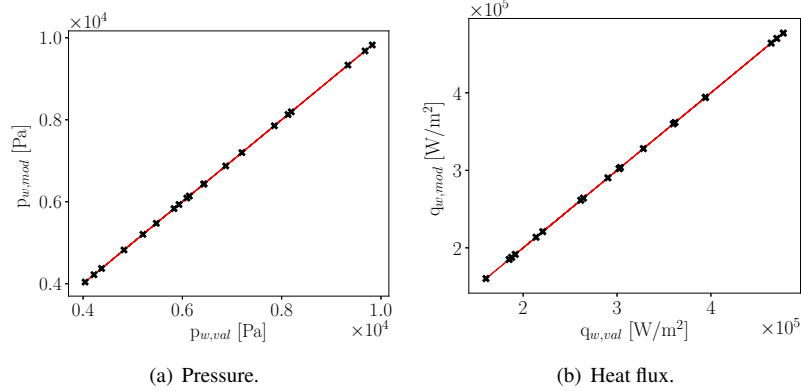


Figure 4.14: QQplot obtained for the adapted mesh II: the values returned by the surrogate model are plotted against the ones obtain with CFD computations. The surrogate models well predict the verification points.

obtained for the heat flux are shown in Figure 4.15(a) and 4.15(b), as a function of the degree of the refinement of the mesh employed to perform the simulations to train the surrogate model. As one can see, both mean and standard deviation follow the same polynomial convergence behavior of the heat flux computed at nominal conditions.




Figure 4.15: Statistical moments of the heat flux, obtained by propagating the uncertainties on the BCs, as a function of the degree of refinement of the meshes. They follow the same polynomial behavior observed in the nominal case.

At this point, it makes sense to compare the UQ-driven standard deviation and the numerical uncertainties of each mesh. We define an average numerical uncertainty, and a relative standard deviation, associated with each grid (g) and

properties (ϕ) as:

$$U_g(\phi) = \frac{\sum_{t=1}^{n_t} U_{g,t}(\phi)}{n_t}, \quad (4.15)$$

$$S_g(\phi) = \sqrt{\frac{\sum_{t=1}^{n_t} (U_g(\phi) - U_{g,t}(\phi))^2}{n_t - 1}}, \quad (4.16)$$

where $U_{g,t}$ is the numerical uncertainty associated to each simulation t , used to train the surrogate model, performed on the numerical grid g . Their values, in the case of non-adapted and adapted grids, divided by the UQ-driven standard deviation of each grid, $(U_g(\phi) \pm S_g(\phi))/\sigma(\phi, h_i/h_1) \cdot 100$, are shown in Figure 4.16: small ratios indicate that the numerical uncertainty is negligible compared to the one induced by the uncertainties in the BCs.

As one can see in Figure 4.16(a), the numerical uncertainty associated with the pressure decreases with the degree of the refinement and it is systematically lower when the mesh is adapted. Moreover, it can be seen that its magnitude is reduced more by adapting the grid than by refining it. Furthermore, in the chosen uncertainty space, even the coarsest non-adapted grid has a numerical uncertainty of at least one order of magnitude lower than the UQ-driven standard deviation, which makes accurate the use of mesh IV for surface pressure estimate.




Figure 4.16: Ratio between the average grid numerical uncertainty, $U_g(\phi)$, and the UQ-driven standard deviation, $\sigma(\phi, h_i/h_1)$, plotted with solid lines as a function of the degree of refinement. The standard deviation of the grid numerical uncertainty, $(S_g(\phi)/\sigma(\phi, h_i/h_1))$, divided by UQ-driven standard deviation is shown with opaque area. The numerical uncertainty is systematically lower when the mesh is adapted.

Regarding the heat flux, Figure 4.16(b), the trend is not monotonic as for the pressure, but the average numerical uncertainty, and the associated standard deviation, is systematically lower in the case of the adapted grid. Remark that the significant standard deviation for the non-adapted grid is biased by few simulations

where the carbuncle effect is observed. In this case, as one can see in Figure 4.17, the numerical uncertainty is driven high because of the bad scatter of the data. Adaptation, preventing the carbuncle effect from spreading, reduces the numerical uncertainty to acceptable values and is necessary for estimating the heat flux.




Figure 4.17: Numerical uncertainty as a function of the degree of refinement for a free-stream velocity of 5600 m/s. It is driven high when the carbuncle effect corrupts the solution on non-adapted grids.

From what observed above, the efficient mesh IV, together with grid adaption tools, can be employed to train the surrogate model in a very efficient way, as its numerical uncertainty on the prediction is at least one order of magnitude lower than the relative UQ-driven standard deviation.

4.6 Summary

In this chapter, we first described how to discretize the governing equation according to a FV formulation, highlighting the main numerical challenges. After that, we presented the numerical solvers employed in this work: I) the high-fidelity US3D solver, II) the VKI stagnation-line code that solves the Navier-Stokes equations in a quasi-1D fashion, and III) a code to compute the stagnation-point pressure and heat flux experienced by a probe exposed to a supersonic jet. Then, we presented a methodology to estimate the numerical error associated with the use of a given mesh. It accounts for the scatter of data and is robust for practical applications. Next, the coupling between the US3D solver and Mutation⁺⁺ library was described. It allows for using Mutation⁺⁺ for computing chemistry and transport properties, as well as to solve the GSI balances. Several test cases were run for verification purposes. They show excellent agreement with the reference solutions. Interestingly, we showed that the heat flux on a cylinder exposed to a hypersonic flow obtained by US3D agrees more with the experimental values when the chemistry is computed using Mutation⁺⁺.

Finally, we illustrated a systematic quantification of the numerical error and the uncertainty-induced variability for the simulation of hypersonic flows. Specifically, we showed that the use of grid alignment tools should always be adopted. In fact, they guarantee a better trend of the numerical error and improve the order of convergence. This behavior translates into having less numerical uncertainty associated with the estimate of the QoIs on a given mesh. Furthermore, they are essential in a UQ context, where it is not possible to manually adjust the mesh for each free-stream condition. In fact, non-physical peaks in the heat flux estimate were observed for some training points when the nominal mesh was used and the consequent surrogate model is not reliable. By contrast, grid adaptation tools prevent carbuncle effects from possibly corrupting the solution in specific training points. In this case, a regular behavior of the QoIs was observed with a resulting good quality of the associated surrogate model.

We also performed a grid convergence study to evaluate the numerical uncertainty expected by using a given mesh. It was compared to the UQ-driven standard deviation, allowing us to choose the most efficient mesh to perform the simulations required to train the surrogate model: the numerical uncertainty associated with the coarsest, and most efficient, grid turned out to be at least one order of magnitude less than the UQ-driven standard deviation. An efficient and robust surrogate model can be built employing this mesh together with alignment tools. Lastly, it was noticed that the magnitude of the numerical uncertainty was reduced more by adapting the grid than by refining it.

5

Holistic characterization of an under-expanded high-enthalpy jet under uncertainty

In this chapter, we present a methodology to efficiently characterize an under-expanded high-enthalpy jet obtained in the von Karman Institute's Plasmatron facility, for which no robust rebuilding procedure existed to date. Given the high computational cost of the high-fidelity simulations needed to describe the flow, we built an adaptive/multi-fidelity surrogate model to approximate the solver. A Bayesian inference method provided then rigorous estimates of the uncertainties of the rebuilt quantities. We show that the reservoir pressure and temperature, and the nitrogen catalytic recombination coefficient, can be accurately determined from the available measurements. Contrarily, the test conditions did not allow us to estimate the oxygen catalytic recombination coefficient. Finally, the characterized uncertainties were propagated through the numerical solver yielding an uncertainty-based high-fidelity representation of the hypersonic flow's structure variability.

The results of this chapter are included in the following publications:

- Capriati, M., Turchi, A., Congedo, P., Magin, T. **Holistic characterisation of an under-expanded high-enthalpy jet under uncertainty** (*under submission*).
- Capriati, M., Turchi, A., Congedo, P., Magin, T. **Heat flux characterization of an under-expanded/supersonic plasma jet over a catalytic probe**, EUCASS, 2022, Lille, France.

5.1 Overview

As we saw in Section 2.4.1, high-enthalpy subsonic flows, which duplicate the in-flight chemically reacting BL, can be obtained in the VKI Plasmatron facility. In this context, coupled numerical-experimental approaches have been developed and employed to deterministically extract the effective catalytic recombination coefficients [136] and the ablation reaction efficiencies [3, 219] of the TPMs. Recently, stochastic methodologies have been devised to improve the robustness of such a characterization [51, 52, 61–63]. In these studies, the subsonic nature of the flow allowed for reducing the dimension, and the cost, of the problem. In fact, it was possible to reconstruct the QoIs performing CFD simulations of the only stagnation line, making the UQ studies very efficient.

In Section 2.4.1, we presented a supersonic campaign that was recently conducted in the VKI Plasmatron facility to obtain high magnitudes of the heat flux over the TPMs. Supersonic flows involve various physical effects, including compressible structures and non-equilibrium phenomena. While the latter are generally confined in the BL in subsonic flows, they embrace the whole field in the supersonic testing. Because of these features, characterizing such flows is a challenging task, especially when one is interested in assessing the uncertainty of the QoIs. In fact, multi-dimensional high-fidelity computations are needed to simulate all the different features of the flow. For example, Diaz et al. [220] used high-fidelity computations to build a polynomial regression and evaluate the nozzle inflow conditions in an arc jet experiment as a function of the probe-measured heat flux and pressure. Brune et al. [53] built a surrogate model on high-fidelity computations to perform a sensitivity analysis on the impact of 47 uncertainty parameters on the heat flux and pressure predicted on a probe exposed to a supersonic flow, also in an arc jet experiment. In the latter case, 1500 high-fidelity CFD simulations were needed to train and verify the surrogate model. Such a large number of high-fidelity simulations might lead to a computationally expensive numerical campaign.

The rebuilding of supersonic flows requires a general and robust methodology, mixing accurate measurements and multi-physics simulations while accounting for several sources of uncertainty. Building on the development of stochastic approaches, we propose a multi-fidelity-based Bayesian methodology to rebuild the flow free-stream conditions and the catalytic efficiency of the materials in supersonic high-enthalpy experiments. Given the costly nature of the high-fidelity simulations needed to accurately describe the flow, attention was focused on optimizing the method. A surrogate model was built by means of an adaptive/multi-fidelity strategy to improve the efficiency while preserving the accuracy. The methodology allowed us to characterize the high-enthalpy under-expanded jet described in Section 2.4.1,

for which a standardized rebuilding procedure did not exist.

First, we present a detailed description of the flow based on a nominal high-fidelity simulation. Then, we outline the results of the characterization of the uncertainties on the nozzle's reservoir conditions, and on the catalytic efficiency of the probe, given the measurements at hand. Indeed, the heat flux at the wall is strongly dependent on the atomic recombination at the surface, which can be described using an effective catalytic efficiency of the material. Another relevant result concerns the assessment of the robustness of the high-fidelity prediction of the hypersonic flow structure, obtained by propagating the reservoir uncertainties.

5.2 Deterministic simulations

In this section, we present the results obtained under nominal conditions with the US3D solver. We simulated the flow from the exit of the plasma torch to the probe location. It was assumed to be laminar¹ and steady. Thus, we solved chemical reacting Navier-Stokes equations of an air mixture of five species, $S=[N_2, O_2, NO, N, O]$, in thermal equilibrium. The reaction rate coefficients were taken from Park et al. [79]. Thermodynamics properties were obtained using the NASA polynomials. Diffusive mass fluxes were computed using the SCEBD model, while viscosity and thermal conductivity according to the Gupta-Yos mixture rule supplied with the Eucken correction for the internal modes².


We employed the modified Steger-Warming scheme for the computation of the numerical fluxes with a MUSCL-limited reconstruction to obtain second-order accuracy and the full matrix point relaxation [221] for the time integration.

The physical domain was discretized as shown in Figure 5.1(a); from the left, in a clockwise direction, we can see I) the exit of the plasma torch/nozzle inlet, II) the sonic nozzle surface, III) the expansion chamber, and IV) the catalytic probe. Unlike the EXPERT application treated in the previous chapter, US3D cannot directly connect the probe surface to the inlet and its tailoring utility cannot be used. For this reason, attention was posed during the mesh generation to well capture the shock, as shown in Figure 5.1(b). The straight walls of the chamber were cut using a quarter of circumference to improve the alignment of the mesh to the shock. We remark that such a cut is not expected to affect the solution, as this area is not characterized by gradients in the flow quantities.


Four nested grids, listed in Table 5.1, were generated to assess the grid convergence by systematically doubling the number of the nodes in both the x and y

¹A RANS simulation was carried out to verify that the heat flux experienced by the probe is not sensitive to turbulence. It is an expected behavior given the high temperatures involved in the experiment.

²We computed the transport properties employing the native library of US3D. In fact, this simulation showed convergence issues when using the coupling with Mutation⁺⁺. This is probably due to the computation of the diffusive fluxes, where the face gradients were approximated with finite differences.



(a) Numerical domain, from the left, in the clockwise direction: I) exit of plasma torch/nozzle inlet, II) sonic nozzle surface, III) expansion chamber, and IV) probe.



(b) Zoom of the numerical domain in the shock region: grid aligned to the shock.

Figure 5.1: Details of Mesh III of Table 5.1.

directions.

Regarding the simulation BCs, the total temperature and pressure were imposed at the nozzle inlet. Here, the chemical composition was assumed to be at equilibrium. The nozzle surface was characterized by a non-reacting isothermal condition (whose temperature was kept constant at 1500 K for all the simulations). The chamber surface was characterized by an inlet condition for numerical stability

Index	Cells	Δn , m	h_i	t_{CPU} , min
I	172 224	5×10^{-7}	1	≈ 1600
II	43 056	1×10^{-6}	2	≈ 200
III	10 764	2×10^{-6}	4	≈ 30
IV	2691	4×10^{-6}	8	≈ 4

Table 5.1: Details of the meshes used: index, number of cells, length of the first cell at the stagnation point (Δn), normalized characteristic mesh (h_i), and time to converge (t_{CPU}).

imposing a small horizontal velocity (1 m/s), room temperature, and a chamber pressure of 5.5 hPa (also these conditions were kept constant for all the simulations). The probe was modeled as an isothermal/catalytic surface, solving a SMB through Mutation⁺⁺. A temperature of 350 K was imposed. Finally, a supersonic outlet condition was prescribed on the exit section; all the other surfaces were characterized by a symmetry BC.

The values of the total temperature and pressure at the nozzle entrance, and of the recombination probabilities on the catalytic probe were set according to the training points, as it will be discussed in Section 5.3. In the following, we will describe the results obtained under nominal conditions ($T_0 = 7500$ K, $p_0 = 16500$ Pa, $\gamma_N = 0.0736$, and $\gamma_O = 0.1170$). The nominal catalytic efficiencies were taken from Bellas [99].

5.2.1 Flow structure

Two numerical representations of the experiment, based on the above-mentioned nominal conditions, are given in Figure 5.2 to illustrate the flow features. The undisturbed flow is shown on the top, while the probe-disturbed one at the bottom.

The flow structure is rather complex, particularly:

1. a Prandtl-Meyer expansion fan develops at the exit of the nozzle to expand the flow to the chamber pressure;
2. a jet boundary separates the jet from the quiescent chamber gas;
3. a barrel shock deviates the streamlines to align them to the jet boundary;
4. a Mach disk, characteristic of highly under-expanded jet, is generated, when the flow is undisturbed, because of the oblique shock singular reflection;
5. a reflected shock develops at the triple point (red circle) where the barrel shock intersects the Mach disk;




Figure 5.2: Numerical flow structure obtained with a total temperature of 7500 K and a total pressure of 16500 Pa imposed at the entrance of the sonic nozzle. Density gradient contours. Undisturbed flow on the top, probe-disturbed on the bottom. Structure: (a) Prandtl-Mayer expansion fan, (b) Jet boundary, (c) Barrel shock, (d) Mach disk, (e) Reflected shock, and (f) Detached shock. The triple point with a red circle.

6. a detached shock appears because of the inference of the intrusive probe. Behind it, the flow expands because of the body's curvature. The detached shock interacts with the barrel one. The transmitted detached shock is reflected on the jet boundary, deviating it.

The temperature profile along the stagnation line was extracted from the two simulations; it is shown in Figure 5.3. As one can see in Figure 5.2 and 5.3, being the flow supersonic, the disturbance does not travel upstream and the structure of the flow in the undisturbed and disturbed case is nearly identical up to the edge of shock. Thus, we can safely characterize the nozzle inlet conditions by intrusively probing the flow.

The temperature and Mach contours are respectively shown at the top and the bottom of Figure 5.4. The flow first accelerates in the nozzle to the sonic condition, then it continues the expansion in the chamber, reaching a peak Mach number of around 4.5, just before the detached shock. It re-accelerates behind it. Contrary, the temperature drops during the expansion and sharply increases in the shock layer. Both the BL developing along the nozzle wall and the jet boundary are characterized




Figure 5.3: Distribution of temperature along the jet central line: profile of the undisturbed flow in solid black line, probe-disturbed one in dashed red lines. The flow is nearly identical up to the edge of shock.

by an extended temperature gradient, which explains the thick density gradient in these regions observed in Figure 5.2.

It is also worth analyzing the chemistry of the flow. Despite the temperature drop, the flow appears to be chemically frozen through the whole expansion, as shown in Figure 5.5, due to the pronounced velocity gradients. Strong diffusion is visible across the jet boundary, where the dissociated gas meets the quiescent one, mostly characterized by molecules at room temperature. The atomic nitrogen and oxygen mass fractions profiles, as well as the temperature one, were extracted along the central line and shown in Figure 5.6. As one can see, the atoms recombine in the BL developing in front of the probe, mostly driven by surface catalysis. The atomic nitrogen also partially recombines along the wall of the nozzle.

5.2.2 Numerical uncertainty

The same nominal conditions were imposed to characterize the numerical uncertainty associated with each of the four meshes reported in Table 5.1, based on the formulation presented in Section 4.3. The obtained order of convergence, asymptotic value, and numerical uncertainty of Mesh I are reported in Table 5.2 for the stagnation-point pressure and heat flux, as well as for the nozzle mass flow rate. The experimental uncertainty on these observables is also given in the same table.

The dependency of the results versus the grid refinement for the three observables is shown in Figure 5.7. The coarse solutions are characterized by high numerical uncertainty, which reduces refining the grid. Specifically, as one can see in Table 5.2, the Mesh I has a numerical uncertainty that is lower than the




Figure 5.4: Temperature (top) and Mach (bottom) contours. The flow expands reaching a Mach number of around 4.5 before the shock. The temperature drops during the expansion and increases at the shock layer.

Observable	p	ϕ_0	U_1	$2 \sigma_{\text{exp}}$
p , Pa	1.57	2436.90	38.57	50
q , W/m ²	1.57	4.48×10^6	7.5×10^4	4.48×10^5
m , kg/s	0.55	5.8×10^{-3}	1.1×10^{-4}	3×10^{-4}

Table 5.2: Results of the convergence study performed on the meshes in Table 5.1. For each observable: order of convergence, p , asymptotic value, ϕ_0 , numerical uncertainty of Mesh I, U_1 , and experimental uncertainty, σ_{exp} .

experimental one for all the observables, making it a very accurate representation of the problem. As we will discuss in the next section, one may use many simulations on the efficient, but less accurate, Mesh IV to train a surrogate model, retaining fewer simulations on Mesh I for accuracy purposes.




Figure 5.5: Mass fractions contours: atomic oxygen (top) and atomic nitrogen (bottom). The flow results chemically frozen through the whole expansion.




Figure 5.6: Temperature (left) and mass fractions (right) distribution along the jet center line, from the nozzle inlet to the probe stagnation point. The flow partially recombines in the BL, mostly driven by surface catalysis.

5.3 Uncertainty-based results

As stated at the beginning of the chapter, we are interested in characterizing the experimental flow conditions. The total pressure, and the related experimental




Figure 5.7: Dependency of the results versus the grid refinement for the three observables. Cross: CFD value; solid line: fit law according to Equation (4.9); bars: numerical uncertainty according to Equation (4.13).

uncertainties, can be easily measured at the wall upstream of the nozzle. On the other hand, the total temperature is not directly accessible and has to be rebuilt by means of auxiliary measurements. As we discussed in the previous section, the undisturbed and probe-disturbed flows are identical up to the edge of the shock layer. Thus, we can safely characterize the nozzle inlet conditions by intrusively probing the flow. For example, downstream of the nozzle the experimental setups allow for measuring the stagnation-point heat flux and pressure. The mass flow rate in the nozzle, being controlled by a rotameter, is also known within its uncertainty. Once the inlet conditions are characterized, one can compute, through CFD simulations, the free-stream conditions at any point at which a sample would be exposed to the plasma flow.

Bayesian inversion is a suitable framework to characterize the uncertainty of the QoIs, given the uncertainty on auxiliary observations. Based on the above description, inputs to the UQ problem are:

$$\mathbf{x} = [p_0, T_0, \gamma_N, \gamma_O]$$

where, besides the total pressure and temperature at the nozzle inlet, we have introduced the catalytic efficiencies of the probes for nitrogen (γ_N) and oxygen (γ_O) recombination. In fact, these coefficients strongly affect the heat flux predictions, and their values are not well-known. Interestingly, the proposed methodology does also allow us to characterize the catalytic response of the material.

With the above-mentioned entrance, a mathematical model is needed to numerically predict the experimental observations:

$$\mathbf{y} = [q^{\text{exp}}, p^{\text{exp}}, m^{\text{exp}}].$$

Regarding the uncertainties, a Gaussian uncertainty, reported in Table 5.3, was prescribed to these quantities that are measured. For those that are not experimentally accessible and need to be characterized, a non-informative uniform prior

uncertainty was chosen, as shown in Table 5.4. Note that we prescribed log-uniform distributions to the recombination efficiencies as their priors span several orders of magnitude.

Quantity	μ	2σ
p , Pa	2500.0	50.0
q , MW/m ²	4.48	0.448
m , kg/s	6.0×10^{-3}	3×10^{-4}
p_0 , Pa	16 500.0	500.0

Table 5.3: Informative Gaussian experimental uncertainty.

Quantity	min	max
T_0 , K	5600.0	8400.0
$\log_{10} \gamma_N$	-4	0.0
$\log_{10} \gamma_O$	-4	0.0

Table 5.4: Non-informative uniform uncertainty.

Following, we first introduce the construction of the surrogate model and the results of the forward propagation. Then, we present the results of the characterization of both the flow conditions and the catalytic properties of the probe obtained by performing a Bayesian inversion. The calibrated inflow conditions were finally propagated to characterize the variability of the flow structure. These results are presented next.

5.3.1 Surrogate model construction and forward propagation

A surrogate model, which expresses \mathbf{y} as a function of \mathbf{x} , was built to accelerate both the forwards analysis and the Bayesian inverse problem. In previous works from del Val [173] and Turchi et al. [51, 52], efficient low-fidelity CFD solvers were employed as a mathematical model to reconstruct a subsonic plasma flow. However, the previously-discussed compressible features of the supersonic case cannot be described using a low-fidelity model. Furthermore, it was pointed out in Section 4.2.3 that the low-fidelity model that we derived to describe the experiment predicts lower stagnation-point pressure and heat flux. For example, in the nominal condition, the low-fidelity model predicts $p^{\text{LF}} = 2115.9$ Pa. The error computed with respect to the high-fidelity asymptotic value in Table 5.2 is 321 Pa, much higher than the experimental uncertainty. For these reasons, high-fidelity simulations must

be employed to fully capture the complex physics and to accurately compute the experimental observations. This gain in accuracy comes at the expense of efficiency, which is not ideal when working in a UQ framework. It was restored by building the surrogate model in an adaptive/multi-fidelity fashion, as discussed in Section 3.4.

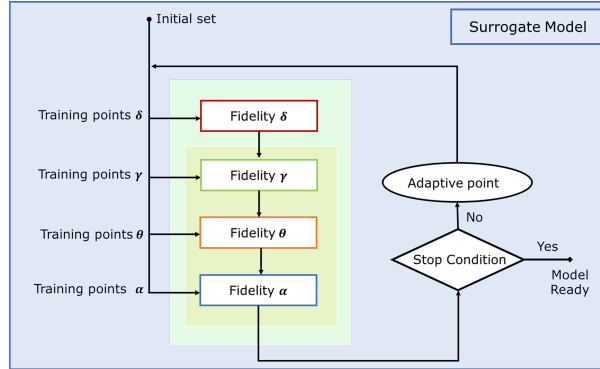


Figure 5.8: Sketch of the adaptive/multi-fidelity surrogate model construction.

The surrogate model is constructed as sketched in Figure 5.8. The lowest-fidelity model (δ) is computed according to the cheap representation of the problem described in Section 4.2.3. The other three fidelity, γ , θ , and α , rely on US3D computations on three differently refined grids, namely Mesh IV, III, and I, characterized by a different numerical uncertainty, as discussed in Section 5.2.2.

First, an ordinary Kriging was built by means of the δ low-fidelity training points. Then, a hierarchical Kriging was progressively trained for fidelity γ , θ , and α using the corresponding l -fidelity CFD training points and the $l - 1$ -fidelity surrogate representations as trend.

The initial number of training points was set to $N_\delta = 160$, $N_\gamma = 20$, $N_\theta = 20$, and $N_\alpha = 5$; another $N_v = 5$ independent points on the α fidelity were computed for verification purposes. The training set was then adaptively refined employing the Kriging standard deviation of the heat flux model, using the method outlined in Section 3.4.3. A normalized error was introduced as:

$$NRMSE = \sqrt{\frac{\sum_{i=0}^{N_v} (Y_{v,i} - \hat{Y}_i)^2}{N_v}} \frac{100}{\max Y_{v,i} - \min Y_{v,i}}, \quad (5.1)$$

to evaluate the accuracy of the surrogate model ($Y_{v,i}$ indicates the CFD estimate, while \hat{Y}_i the Kriging prediction). As one can see in Figure 5.9, the $NRMSE$ on the heat flux decreased from the 1.874% to 0.4936% sampling a total of 82 δ , 26 γ , and 24 θ training points. This set of points was further enriched with 3 α points to get a $NRMSE$ on the pressure and on the mass flow rate prediction respectively of the 0.43196% and of the 0.22846%.




Figure 5.9: Iterative $NRMSE$ for the three observables.

It is worth noting that most of these points were added on the border of the four-dimensional hypercube, as shown in Figure 5.10. This is expected by exploring techniques. In fact, the borders of the hypercube are characterized by the highest uncertainty. Anyway, we remark that the majority of the added points were obtained employing the cheap δ model, with a relatively null impact on the overall computational cost.




Figure 5.10: Added points: the majority of the points were sampled on the border of the hypercube.

The gain in accuracy obtained by employing lower-fidelity simulations in a hierarchical fashion is highlighted in Figure 5.11, where the model predictions on the heat flux are plotted against the CFD verification values. The multi-fidelity values (in green) almost lie on the 45-degree line, whereas the model predictions obtained using only high-fidelity (in red) points mispredict the response. Further increase in accuracy, due to the adaptive sampling strategy, is evident on the same plot: these predictions (in blue) perfectly lay on the 45-degree line.

We also remark that the initial number of δ points could have been increased with an almost null effect on the total CFD computational cost. Nevertheless, it was observed that while the $NRMSE$ associated with the lowest-fidelity model,




Figure 5.11: QQplot: CFD response VS Surrogate model response. Points from the adaptive/multi-fidelity surrogate model on the 45-degree line.

computed on 50 δ validation points, consistently decreased with the number of low-fidelity training points (shown in dashed line in Figure 5.12), the $NRMSE$ associated with the α fidelity does not necessarily follow the same trend (solid lines in Figure 5.12). In fact, the $NRMSE$ on the α model first decreases for all the quantities when the δ surrogate model is constructed on 160 points and used as trend to build the γ model (0 δ points means that the δ model is not included in the construction of the α model). Further refining the δ model, leaving unchanged the number of training points of the other levels, does not consistently improve the high-fidelity predictions. Furthermore, a reduced number of training points improves the computational cost associated with the surrogate model construction. For these reasons, the number of initial δ training points was set to 160, refining it where needed.




Figure 5.12: $NRMSE$ as a function of the number of low-fidelity points associated with the α surrogate model (solid line), and with the δ model (dashed line). For the quantities: heat flux (in black), pressure (in red), mass flow rate (in blue).

As we said previously, the tailoring tool of US3D cannot be used in such an

application. Hence, the surrogate model was visually inspected to make sure that spurious behaviors were not present. The projection of the heat flux on the total pressure and total temperature space is plotted in Figure 5.13. As one can see, the predictions are physically sound, and signs of carbuncle were not detected. We remind that the peak Mach number in this application is around 4.5, much lower than the 15 of the EXPERT case. This makes the solution less prone to the appearance of numerical errors across the shock.




Figure 5.13: Projection of the heat flux surrogate model on the T_0 and p_0 dimensions. No carbuncle effect observed.

The verified surrogate model was then employed to propagate the uncertainties on T_0 , γ_N and γ_O , in Table 5.4, and the one on p_0 in Table 5.3. A number of 10^5 LHS points were propagated through the surrogate model. The obtained probability densities of the observables are plotted in Figure 5.14 against the experimental ones. As one can see, the former are larger and contain the latter. An inverse problem can be then solved to reduce the uncertainties on the nozzle inlet conditions and on the catalytic efficiencies.

5.3.2 Characterization of the flow conditions and of the material catalytic response

The same surrogate models were employed to perform the Bayesian inversion. A total of 10 MCMC chains characterized by 10^5 iterations were generated. The first 20% of the points of each chain was neglected as Burn-In. The multivariate potential scale reduction factor, computed with Gelmen-Rubin diagnostic, is $R = 1.0035$, ensuring that the chain is well converged.

The posterior marginals relative to the four QoIs are plotted in Figure 5.15. Specifi-




Figure 5.14: Propagated uncertainties in red. Experimental uncertainties in black.

cally, the reservoir's total temperature probabilistic density present a well-defined peak at 7000 K, with a standard deviation of around 300 K, leading to a coefficient of variation ($CV = \sigma/|\mu| \cdot 100$) of the 4.28%. The reservoir's total pressure probability density is Gaussian-shaped, with a mean of around 16730 Pa, and a standard deviation of 140 Pa ($CV = 0.84\%$). Interestingly, we almost halved the experimental uncertainty by using complementary measurements. The probe's nitrogen catalytic efficiency turned out to be also well-characterized, with a probability density peak value of around 0.15, and support ranging from 0.01 to 1. These values are in agreement with those found by del Val et al. [62] in a study performed on a subsonic test in the Plasmatron, where copper's γ_N and γ_O were inferred using the same variable. On the other hand, it was not possible to characterize the efficiency of the oxygen recombination. In fact, the atomic oxygen fraction is much lower than the atomic nitrogen one, see Figure 5.6. Hence, also the associated fraction of heat released during the recombination process is lower compared to the fraction released by the nitrogen reaction. This makes it impossible to use the heat flux measurement for its estimation. To better characterize this quantity, one should perform an experiment at a lower temperature, in a range where oxygen is dissociated, but nitrogen is mostly in molecular form. The joint distribution of γ_N and γ_O is plotted in Figure 5.16. As one can see, the two QoIs are correlated and a characterization of the oxygen recombination coefficient would improve also the characterization of the nitrogen one.

5.3.3 Flow structure variability

The prior and posterior distributions on the total pressure and temperature at the entrance of the nozzle (p_0, T_0) were propagated to investigate the variability of the flow structure. To this end, an algorithm was written to automatically detect the compressible features in each simulation. It follows a series of steps:

1. from the solution field, the stagnation line and 24 vertical lines are extracted in the range $x = [0.16 - 0.27]$ m.

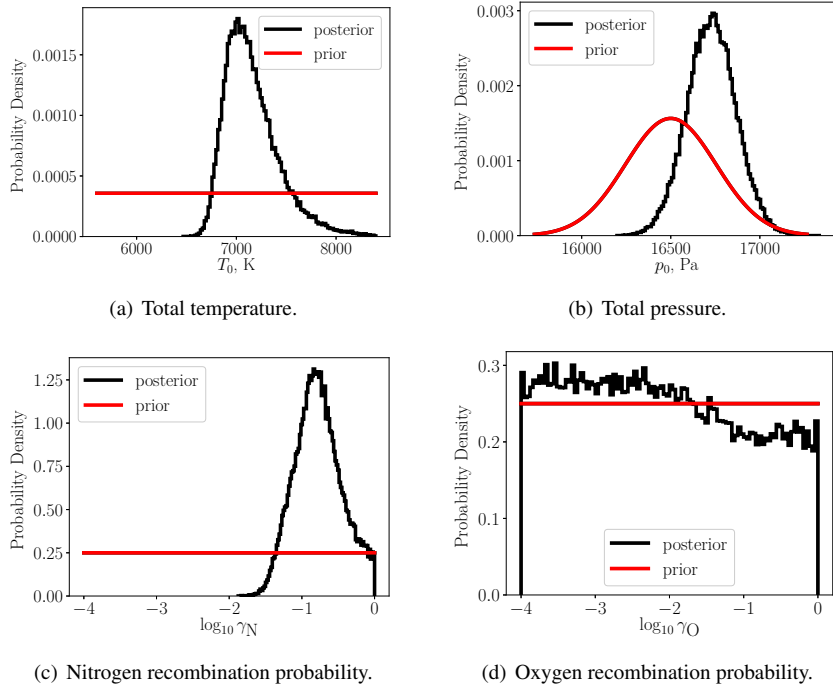


Figure 5.15: Prior and posterior marginal distributions for the four QoIs.




Figure 5.16: Joint distribution of γ_O and γ_N .

2. The points belonging to the jet boundary, the barrel shock, and the detached shock are identified by detecting the relative maxima in the density gradient lines. The reconstructed points are shown in Figure 5.17.
3. The reconstructed points are then fitted for two reasons: I) to reduce the noise in the rebuilt features, and II) to obtain the fit coefficients, which are more convenient to construct a surrogate model and to propagate the uncertainties. Two different fits are performed to reconstruct the detached shock: the first up to the point in which it intersects the barrel shock, and from here to $x = 0.27$ m, as its shape changes at the intersection point. The jet boundary, the barrel shock, and the second part of the detached shock were fitted with a third-order polynomial. The term $\sqrt{x - x_{\text{shock}}}$ was added to the polynomial used to fit the first part of the detached shock to enforce a vertical tangent at the shock location.
4. The stand-off distance is computed as the difference between the x coordinate of the relative maximum of the density gradient along the stagnation line, before the increase in the BL, and the coordinate of the stagnation point.
5. The coordinates of the point in which the detached shock intersects the barrel shock are computed as the point in which the fit of the detached shock intersects the one of the barrel shock.




Figure 5.17: Reconstructed points belonging to the jet boundary (plotted in yellow), the barrel shock (in green), and the detached shock (in red).

A surrogate model was built for the coefficients of the fits of each compressible feature, for the coordinates of the intersection point, and for the stand-off-distance. We performed other 15 high-fidelity simulations to improve the quality of the surrogate model. The simulation inlet conditions were randomly sampled from the reduced posterior distribution.

Both the prior and the posterior distributions were then propagated through the surrogate models. The propagated uncertainties on the flow structure are shown




Figure 5.18: Prior and posterior uncertainty on the flow structure: jet boundary (in orange), barrel shock (in green), and detached shock (in red). Prior: mean of the propagation plotted with dashed line, bounded by dash-dotted lines indicating the 95% interval of confidence. Posterior: the continuous line represents the mean, while the shadow areas represent the uncertainty within the 95% of confidence.

in Figure 5.18. As one can see, the barrel shock exhibits the largest variability with respect to the inlet conditions. This gets dramatically reduced when the total pressure and temperature are characterized. The same holds true also for the jet boundary and the detached shock, although the differences are much less pronounced. We can better appreciate the improvement in the predictions of these features by looking at Figure 5.19(a) and Figure 5.19(b), where the shock stand-off distance and the intersection point variability are respectively shown. The most likely distance of the shock from the stagnation point increases from 2.90 mm to 2.95 mm when we characterize the inlet conditions, while the associated uncertainty decreases from $\sigma = 0.06$ mm ($CV = 2.08\%$) to $\sigma = 0.038$ mm ($CV = 1.29\%$). Likewise, the uncertainty in the intersection location is greatly reduced. Overall, the characterization of the inlet conditions yield a much more robust prediction of the flow structure.

5.4 Summary

We proposed a multi-fidelity Bayesian-based methodology to rebuild the free-stream conditions and the probe catalytic efficiencies in supersonic high-enthalpy experiments. We applied it to the characterization of an experiment performed in the VKI Plasmatron facility, for which no standardized rebuilding procedure existed. Such a flow can be accurately described by means of high-fidelity simulations. Considering the large number of calculations needed for UQ studies and their associated computational cost, we built an adaptive/multi-fidelity surrogate model




Figure 5.19: Prior (in red) and posterior (in black) uncertainties in the flow features.

to mimic the high-fidelity CFD response. The initial set of training points was further enriched by an exploring infill strategy. Verification tests highlighted the gain in accuracy obtained by constructing the Kriging surrogate model leveraging lower-fidelity representations. The surrogate model was also visually inspected to verify the absence of carbuncle effects. Once verified, the surrogate model was employed to solve the inverse problem, allowing us to obtain rigorous estimates of the uncertainties in the rebuilt quantities.

The Bayesian analysis revealed pronounced peaks for the probability density of the total temperature and total pressure at the entrance of the sonic nozzle, yielding a robust characterization of the test conditions with a coefficient of variation below the 4.3%. An important outcome of the study was the determination of the nitridation catalytic efficiency of the calorimeter probe used to measure the heat flux. Nevertheless, given the relatively small amount of atomic oxygen versus atomic nitrogen present in the flow, it was not possible to estimate the oxygen recombination efficiency.

Finally, the reservoir uncertainties were propagated to assess the robustness of the prediction of the flow structure, identifying its variability. The barrel shock exhibited the greatest variability with respect to the total pressure and temperature, which gets largely reduced when these quantities are characterized. The same holds true for the stand-off-distance and for the point in which the detached shock intersects the barrel one.

We believe that the methodology developed is general and can be applied to characterize those experimental campaigns requiring high-fidelity computations.

6

Bayesian calibration of a finite-rate nitridation model from molecular beam and plasma wind tunnel experiments

The modeling of gas-surface interaction phenomena is crucial for predicting the heat flux and the mass loss experienced by hypersonic vehicles. Gas-surface interactions refer to the phenomena occurring between the reacting gas and the material. An important part of the modeling concerns the description of the surface chemical reactions. In this regards, we propose to calibrate, in a Bayesian sense, the rates of the elementary reactions between a nitrogen gas and a carbon surface. We rely both on molecular beam and plasma wind tunnel observations. The former provides detailed data on the chemical mechanisms, but are characterized by pressures non representative of atmospheric entries. By contrast, plasma wind tunnel experiments are conducted at representative pressures, but contains only macroscopic information. The parameters' posterior distributions are then propagated through the models representing the two experiments. The calibrated model turned out to be

The results of this chapter are included in the following publications:

- Capriati, M., del Val, A., Schwartzentruber, T., Minton, T., Congedo, P., Magin, T. **Bayesian calibration of a finite-rate nitridation model from molecular beam and plasma wind tunnel experiments**, EUCASS-CEAS, 2023, Lausanne, France.
- Capriati, M., Prata, K., Schwartzentruber, T., Candler, G., Magin, T. **Development of a nitridation gas-surface boundary condition for high-fidelity hypersonic simulations**, ECCOMAS, 2021, Virtual.

able to explain both experiments.

6.1 Overview

One key aspect of the GSI modeling is the description of the chemical reactions between the surface and the reacting gas. Among others, reactions involving atomic nitrogen are important in entry applications, for example in Earth and in Titan atmosphere. In fact, the nitridation reaction ($N + C_b \rightarrow CN$) ablates the TPMs, while the exothermic recombination reaction ($N + N \rightarrow N_2$) increases the heat flux experienced by the spacecraft.

Several experimental works addressed the characterization of the efficiencies of the two reactions. For instance, the nitridation efficiency was estimated, over a wide range of temperatures and pressures, by means of experiments in shock tube [222], in furnace-heated quartz flow tube coupled to microwave discharge [27], in ICP facilities [3, 107, 223–225], and using a molecular beam-surface scattering apparatus [142]. The studies [107, 142, 223] included also the determination of the nitrogen recombination efficiency. The values proposed in these works largely differ from each other, suggesting a poor understanding of both the nitridation and the recombination reactions.

It is important to remark that, in most of the above-mentioned studies, coupled numerical simulations supported the experiments in determining those QoIs that were not experimentally accessible, adding an extra layer of modeling uncertainty. For example, the concentration of the atomic nitrogen in ICP facilities, needed to compute the reaction efficiencies, was determined through numerical simulations in [3, 224, 225]. In such computations, the surface recombination was neglected, potentially biasing the determination of the gas concentrations. Furthermore, CFD simulations contain assumptions also on the gas-phase modeling, for instance on the rates of the homogeneous chemistry and on the transport properties. Consequently, the assumptions contained in the modeling of these processes directly affect the characterization of the surface reactions [67]. However, it is virtually impossible to characterize all these sources of uncertainties, along with the surface reaction probabilities, given the limited amount of observations that can be acquired in these facilities. Contrary to ICP, in the molecular beam-surface scattering experiments, the ablative sample is placed in a high-vacuum chamber and bombarded with a beam of molecules or atoms. Thus, few gas collisions are expected to happen and phenomena such as gas chemistry or transport are negligible. As a consequence, simulating these experiments requires no assumptions in the modeling of the gas. However, while these experiments manage to isolate GSI phenomenon, we cannot fully trust that they are representative of hypersonic flight conditions because of the very low pressures involved.

From a modeling standpoint, phenomenological Arrhenius laws for the nitridation reaction have been proposed by Suzuki et al. [224], and, later, by Helber et al. [3]. Such Arrhenius laws capture the only temperature dependence of the reaction and are expected to be valid at pressures close to the ones where they have been calibrated. Recently, Prata et al. [1] derived the ACA model based both on molecular beam and ICP experiments. To capture the different probabilities observed at different temperatures and pressures, both these dependencies were accounted for in the model. Furthermore, the model was calibrated to capture the main trends of the reaction efficiencies experimentally observed. Among these reactions, the ones involving atomic nitrogen were calibrated on the molecular beam experiments performed by Murray et al. [108], and the ICP ones from Lutz [107] and Helber et al. [3]. Very recently, the activation energies proposed in the ACA model have been compared with theoretical results by Nieman et al. [226]. In this study, the reaction energies were computed based on the periodic density functional theory, considering different reaction sites. Some of the theoretical energies well compared with the ACA model, but further refinement was suggested to others.

Few UQ works addressed a robust characterization of the nitridation reaction. Upadhyay et al. [60] employed a Bayesian method to infer the reaction probability accounting for the experimental and modeling uncertainties in explaining the data from Zhang et al. [27]. Four different models, differing in the underlying assumptions, were used for the calibration. The temperature dependence of the nitridation reaction was captured, for each model, calibrating a power law. A Bayesian framework was also used by del Val et al. [63] to infer the nitridation efficiency from the ICP data from Helber et al. [3]. First, a sensitivity analysis was performed, highlighting little influence of the BL edge conditions on the computation of the experimental observations, being the nitridation reaction efficiency the most sensitive parameter. After that, the nitridation probabilities were first individually inferred from each surface-temperature-dependent measurement, then an Arrhenius law was calibrated considering them jointly. No heterogeneous recombination was considered. In a subsequent work by the same authors [64], this aspect was tackled, along with other model-form uncertainties. Although the analysis did not highlighted evidences of the recombination reaction, it showed that its uncertainty broadens the posterior distribution of the nitridation probability.

All the above-mentioned inferences relied on experimental observations obtained in a single facility, involving similar operative pressures. Hence, a phenomenological model was adequate to describe the task at hand. To date, no stochastic calibration has been performed attempting to capture the reaction pressure dependence by informing an FRC model combining observations from very different experiments.

As we will see in Section 6.2, we used the nitridation reactions of the ACA model to numerically simulate the subsonic experimental campaign performed by Helber et al. [3], described in Section 2.4.1. It turned out that some of the observed features (e.g. the surface recession rates) were not fully in agreement with the experimental values. This issue underlined the necessity of including the experimental, as well as the parametric, uncertainties during the calibration of the model parameters. In this context, it was shown that Bayesian methods offer a robust framework for calibrating the phenomenological efficiencies of the surface reactions [27, 63, 64].

With respect to these works, we propose to calibrate the reactions of a more complex FRC model, i.e. the elementary reactions involving nitrogen of the ACA model [1]. The reaction rates, along with their uncertainties, are characterized by means of two different sets of observations. The first set comprises highly-informative molecular beam data. As we cannot fully trust that these data are representative of atmospheric entry conditions, we included in the study also the Plasmatron data, obtained at pressures of interest. First, the Plasmatron data were used to assess the predictive capability of the model calibrated only on molecular beam data. Then, we employed them jointly during the inference. Unlike the work from Prata et al. [1], we simulated the Plasmatron experiments with a CFD solver, where the ablative BC is embedded to properly describe the BL developing in front of the test sample. The molecular beam data were explained through the same analytical model proposed by Prata et al. [1]. The parameters' posterior distributions were obtained by constructing an MCMC. The posterior distributions were then propagated through both the CFD and the 0D model to obtain the posterior predictions. These are compared to the experimental counterparts.

6.2 Deterministic simulations of a Plasmatron experimental campaign

The subsonic Plasmatron experimental campaign described in Section 2.4.1 was numerically simulated using both the stagnation-line code and the US3D solver. The flow was assumed to be steady, laminar, and axisymmetric. The stagnation-line code computations were performed using the Roe scheme and an implicit local time step, while the US3D ones employed the modified Steger-Warming scheme and the DPLR time integration. Both the codes relied on Mutation⁺⁺ for the computation of the chemical rates (with Arrhenius coefficients taken from Olynick et al. [80]) and of the transport properties. Mutation⁺⁺ was also used to solve the GSI balances. In fact, the ablative response of the surface was modeled using the ACA model implemented in Mutation⁺⁺, see Section 2.3. For each of the test conditions (G4,

G5, G6, and G7), a SMB was solved imposing the experimental surface temperatures reported in Table 2.2. The inflow conditions were those at the edge of the BL rebuilt by Helber et al. [3]. These values are also provided in Table 2.2. US3D simulations required also the specification of the chamber pressure (1500 Pa) at the outlet section.

For the four test conditions, the stagnation-point mass blowing rate (\dot{m}) was extracted from the numerical solutions. These values are shown in Figure 6.1, along with the experimental counterparts, against the surface temperature. As one can see, the values predicted by the stagnation-line code and by US3D are in good agreement. However, both of them predict a mass blowing rate lower than what was experimentally observed. Furthermore, they lack in capturing the upward trend of \dot{m} with respect to the surface temperature.




Figure 6.1: Stagnation-point mass blowing rates obtained, for the four test conditions described in Section 2.4.1, with the stagnation-line code (black circle), and with the US3D solver (red cross). Experimental uncertainties indicated with black error bar. The two numerical solutions are in good agreement. However, both of them predict lower values of \dot{m} than the experiments and lack in capturing the upward trend of \dot{m} with respect to the surface temperature.

We remind that the ACA model was calibrated also on these experimental data. However, the model predictions were not directly compared to experimental mass blowing rates, but with the values of the nitridation probabilities derived by Helber et al. [3]. The latter were computed through a coupled numerical-experimental methodology, assuming no surface recombination of the nitrogen atoms. By contrast, the ACA model predicts a high recombination probability, see Figure 6.2, in agreement with what was observed in the experiments performed by Murray et al. [108] and Lutz [107]. Furthermore, the model used by Prata

et al. [1] to reproduce the Plasmatron experiments made strong assumptions on the modeling of the BL developing in front of the sample. In fact, the same 0D approach used to model the molecular beam experiment was adopted. Such a model does not account for the reduction of atomic nitrogen at the surface dependent on the reactions. As a result, although the equivalent nitridation probabilities are in fair agreement with the values proposed by Helber et al. [3], see Figure 6.2, the high recombination probability reduces the atomic nitrogen available at the surface for nitridation, hence the mass blowing rate ($\dot{m} \propto \gamma_{\text{CN}} n_{\text{N}}$).

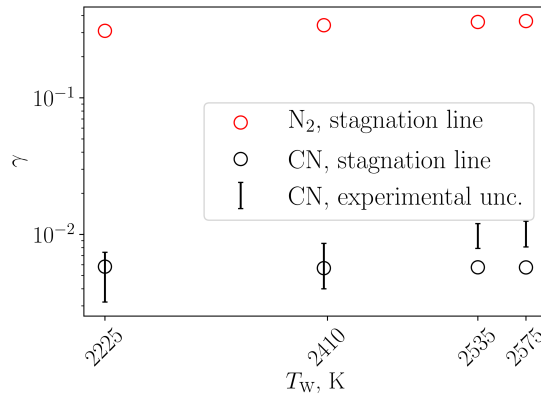


Figure 6.2: Probabilities of the nitridation (black) and recombination (red) reactions for the four test conditions described in Section 2.4.1, obtained with the stagnation-line code (circle). Uncertainties from the numerical-experimental rebuilding by Helber et al. plotted with error bars. High recombination probabilities are predicted. Nitridation probabilities in fair agreement with the numerical-experimental values.

We finally remark that the differences in the predictions between the stagnation-line code and the US3D solver are negligible compared to the experimental uncertainty, see Figure 6.1. Thus, it is possible to perform the UQ studies employing the efficient 1D model without any loss of accuracy.

6.3 Stochastic inverse problem definition

The results obtained in the previous section highlighted the necessity of I) directly exploiting pure experimental observations to calibrate the model (for example, the blowing mass rates), and II) employing an accurate model to reproduce the Plasmatron experiments. Furthermore, the ACA model was constructed deterministically, without accounting for experimental and parametric uncertainties. To improve the robustness of the predictions, we employed a Bayesian formalism to

compute the posterior distributions of the FRC reaction rates. Additionally, we directly used the experimental mass blowing rates to infer the reactions and CFD simulations to reconstruct the flow in the Plasmatron experiments. Following, we present the quantities that we intend to calibrate, along with their prior uncertainties, the observations used for the calibration, and the computational models.

Based on preliminary investigations, we chose to characterize four elementary reaction efficiencies (X_3, X_4, X_5, X_6), and the total active site density (B) of the ACA model; the subscript of the efficiencies denotes the index of the reaction in Table 2.1. Additionally, although it is common practice to consider uncertain only the pre-exponential coefficients of the reaction rates when dealing with homogeneous FRC [34, 55], we decided to infer also five activation energies (E_1, E_3, E_4, E_5, E_6) as different values can be found in literature, see Table 6.1. Hence, the vector of parameters that we intend to characterize is:

$$\boldsymbol{x} = [E_1, E_3, E_4, E_5, E_6, X_3, X_4, X_5, X_6, B].$$

The prior distributions on \boldsymbol{x} are reported in Table 6.1: non-informative uniform distributions over a wide range are chosen to encompass several values available in the literature. Remark that we prescribed log-uniform distributions to the elementary reaction efficiencies and the total active site density because their priors span several orders of magnitude.

Parameter	min	max	Prata	Nieman
E_1, K	0	3000	2500	0, 2436, and 4525*
E_3, K	0	15000	7000	10676
E_4, K	0	15000	2000	0
E_5, K	0	35000	21000	6150
E_6, K	15000	35000	20676	30286
$\log_{10}(X_3)$	-4	2	0.176	-
$\log_{10}(X_4)$	-4	0	-0.301	-
$\log_{10}(X_5)$	-4	0	-1	-
$\log_{10}(X_6)$	5	12	8	-
$\log_{10}(B)$	-7	-3	-5	-

Table 6.1: Non-informative uniform prior distributions on the parameters to calibrate, along with the values proposed by Prata et al. and Nieman et al.

*The three values correspond to three different adsorption sites.

The set of observations was divided into two groups. The first one comprises the efficiencies estimated by Murray et al. [108] in the molecular beam-surface scattering experiments:

$$\boldsymbol{y}^{\text{MB}} = [\gamma_{\text{CN}}, \gamma_{\text{N}_2}]_i \quad i = 1, \dots, 11.$$

We further divided this group into two sub-groups. The first one does not include the data obtained for $T_w < 1000$ K (hereafter referred to as ‘set A’). It is the same subset used by Prata et al. [1] to construct the ACA model. The second subset includes all the available points (hereafter referred to as ‘set B’). We assumed these observations to be independent and the model error to be negligible with respect to the experimental uncertainty, σ_i , on the measurement y_i (i^{th} component of \mathbf{y}). For the set A, the likelihood reads:

$$\pi^{\text{MB,A}}(\mathbf{y}^{\text{MB}}|\mathbf{x}) = \prod_{i=5}^{22} \mathcal{N}(y_i^{\text{MB}}|\mathcal{M}_i^{\text{MB}}(\mathbf{x}), (\sigma_i^{\text{MB}})^2), \quad (6.1)$$

while, for the set B, it is:

$$\pi^{\text{MB,B}}(\mathbf{y}^{\text{MB}}|\mathbf{x}) = \prod_{i=1}^{22} \mathcal{N}(y_i^{\text{MB}}|\mathcal{M}_i^{\text{MB}}(\mathbf{x}), (\sigma_i^{\text{MB}})^2). \quad (6.2)$$

The symbol $\mathcal{M}_i^{\text{MB}}(\mathbf{x})$ stands for a mathematical model which expresses y_i^{MB} as a function of \mathbf{x} . The 0D model described in Section 2.2.4 (*Air carbon ablation model* subsection) was employed to explain the molecular beam data. Along with the parameters \mathbf{x} , the surface temperature and the impinging nitrogen flux are inputs of the model. No uncertainty was considered on these quantities.

The second group comprises the blowing mass rates observed in the Plasmatron experiments:

$$\mathbf{y}^{\text{ICP}} = \dot{m}_i \quad i = 1, \dots, 4.$$

In this case, the likelihood reads:

$$\pi^{\text{ICP}}(\mathbf{y}^{\text{ICP}}|\mathbf{x}) = \prod_{i=1}^4 \mathcal{N}(y_i^{\text{ICP}}|\mathcal{M}_i^{\text{ICP}}(\mathbf{x}), (\sigma_i^{\text{ICP}})^2). \quad (6.3)$$

The symbol $\mathcal{M}_i^{\text{ICP}}(\mathbf{x})$ stands for a mathematical model which expresses y_i^{ICP} as a function of \mathbf{x} . It consists of a surrogate model trained on stagnation-line CFD computations. The settings of the simulations were discussed in Section 6.2. Along with the ACA model rate coefficients, the surface temperature and the conditions at the edge of the BL are inputs of the model.

A total of four different calibrations were performed. They are summarized in Table 6.2, along with the likelihoods used to solve the Bayesian inverse problem. In the first two tests, we used only molecular beam-surface scatter data to calibrate the model. The Plasmatron data were employed to assess the predictive capability of the model at high pressure. In the third test, we calibrated the model only on Plasmatron data to verify that such data can be explained through the ACA model. Finally, in the last test, the calibration was performed imposing both sets of observations.

Test	Molecular beam	Plasmatron	Likelihood
A	Calibration	Assessment	$\pi^{\text{MB,A}}$
B	Calibration	Assessment	$\pi^{\text{MB,B}}$
C	-	Calibration	π^{ICP}
D	Calibration	Calibration	$\pi^{\text{MB,B}} \cdot \pi^{\text{ICP}}$

Table 6.2: Matrix of the performed tests.

The quantities to calibrate turned out to be highly correlated in a high-dimensional space. For this reason, the posterior distributions were obtained by constructing an MCMC using the AIES algorithm [169]. The UQLAB software [189] was employed both to construct the MCMC chains and to train the surrogate model.

Surrogate model construction

A Kriging model was constructed to accelerate the evaluation of the mass blowing rate as a function of both the ACA model and the experimental conditions. First, we attempted to create a global surrogate model able to approximate the CFD response for the four experiments. Thus, on the top of the 10-dimensional \mathbf{x} input, the model was constructed also on the surface temperature, T_w , and on the temperature, T_e , the velocity, u_e and the pressure, p_e , at the edge of the BL. The resulting model:

$$\mathcal{M}_i^{\text{ICP}}(\mathbf{x}) \approx \mathcal{M}^{\text{K,ICP}}(\mathbf{x}, T_{w,i}, T_{e,i}, u_{e,i}, p_{e,i})$$

has a total input dimension of 14. The training space was chosen large enough to accommodate both the prior on \mathbf{x} and the four experimental conditions in Table 2.2, along with the relative uncertainties. Such a space was normalized to simplify the construction of the surrogate model. Its accuracy was evaluated by computing a normalized error based on 150 LHS validation points. Such an error reads:

$$NRMSE = \sqrt{\frac{1}{N_v} \sum_{i=0}^{N_v} \frac{(Y_{v,i} - \hat{Y}_i)^2}{Y_{v,i}^2}} 100, \quad (6.4)$$

where $\hat{Y}_i = \mathcal{M}^{\text{K,ICP}}(\mathbf{x}, T_{w,i}, T_{e,i}, u_{e,i}, p_{e,i})$ is the Kriging model evaluation at the validation point i , while $Y_{v,i}$ is the respective CFD response. The large input dimension, along with the large priors prescribed, deteriorates the rate of convergence of the model. In fact, as one can see in Figure 6.3, the error drops from around 7%, when 200 LHS training points are employed, to around the 2.3% with 6200 LHS points. Such a surrogate model was not very efficient to be used coupled to the sequential AIES algorithm.

For this reason, we employed it to perform a Sobol analysis, aiming at identifying the most sensitive variables and reducing the input dimension. The resulting




Figure 6.3: NRMSE of the global surrogate model built on 14-dimensions with dashed red line. NRMSE of the 6-dimensions G6 surrogate model with solid black line.

total Sobol indices are shown in Figure 6.4. The quantities at the edge of the BL show poor influence on the blowing mass rate. Similarly, the surface temperature exhibits a small influence. It is important to note that such an index was computed on the entire experimental range (≈ 600 K), while the experimental uncertainty on each test is only around 20 K. Hence, the influence of the uncertainty on the surface temperature for each single experiment can be considered negligible. A similar behavior was also observed by del Val et al. [63].




Figure 6.4: Total Sobol indices for the mass blowing rates. The parameters at the BL edge and the surface temperature show poor influence on the mass blowing rate.

Based on this analysis, we decided to build a surrogate model for each test, neglecting the uncertainties on the surface temperature and on the parameters at the edge of the BL. This reduced the input dimension to 10. Further reduction in dimension was achieved by performing a change of variables. In fact, the reaction

rates are in the form:

$$k_i \propto X_i \exp\left(-\frac{E_i}{T_w}\right), \quad (6.5)$$

allowing us to introduce the variable:

$$P_i = X_i \exp\left(-\frac{E_i}{T_w}\right), \quad (6.6)$$

as T_w is considered constant in each surrogate model. The resulting surrogate model:

$$\mathcal{M}_i^{\text{ICP}}(\mathbf{x}) \approx \mathcal{M}_i^{\text{K,ICP}}(P_1, P_2, P_3, P_4, E_1, B)$$

has only an input dimension of 6. The convergence of the model built on the G6 experiment was assessed by means of 150 LHS points. It is shown in Figure 6.3: the error drops from around 5% when 200 LHS points are used, to around 1.2% with 5200 LHS points. For efficiency reasons, we decided to use the surrogate model built on 3000 points with a $NRMSE = 1.6\%$.

6.4 Results of the calibration

In this section, we describe the results obtained by calibrating the model employing the four sets of observations presented in Table 6.2. For all the conditions, 100 MCMC chains were constructed with an AIES algorithm. For each of them, we sampled enough points to have, after the burn-in process, a low multivariate potential scale reduction factor from the Gelmen-Rubin diagnostic ($R_A = 1.1$, $R_B = 1.1$, $R_C = 2.5$, $R_D = 1.2$). For the case C, which shows a slightly high value of R , we assessed the chain convergence by visually inspecting the posterior distributions. In the following, we first present the marginal distributions of the calibrated parameters, then the posterior predictions of the model.

6.4.1 Posterior marginal distributions

The marginal distributions of the calibrated parameters for the two nitridation reactions are plotted in Figures 6.5(a), 6.5(b), 6.6(a), and 6.6(b), while in Figures 6.5(c), and 6.6(c) the two joint distributions relating the reaction efficiency and the activation energy of the selected reactions are given. The values reported by Prata et al. [1] and by Nieman et al. [226] are also shown in the same figures. When the model is calibrated on the molecular beam data, both the nitridation reactions turned out to be well characterized, especially when using set B, as it contains information also at low temperatures. On the other hand, when the model is calibrated using only the Plasmatron data (set C), poor characterization of the parameters is observed. However, when the Plasmatron data are combined with the molecular beam ones (set D) an improved characterization is achieved. This is especially true with regard

to reaction 3, whose parameters present more pronounced peaks. It makes sense as reaction 3 was included in the model to capture the ablative behavior at high pressure.

When the model is calibrated by means of set A, the posterior distribution includes both the values proposed by Prata et al. and by Nieman et al., which makes them consistent with the data employed in the calibration. However, when low-temperature experimental efficiencies are included (set B), all the distributions shift towards lower values. As we will discuss later, this allows the model to predict the non-decreasing trend of the nitridation efficiency observed at low temperatures in the molecular beam experiments, see Figures 6.10 and 6.11. In this case, the distribution of the calibrated activation energies departs from the values proposed by Prata et al. and by Nieman et al., suggesting some incompatibility of the latter in explaining the molecular beam nitridation efficiencies at low temperatures through the proposed model. The same holds true when the model is calibrated with set D.

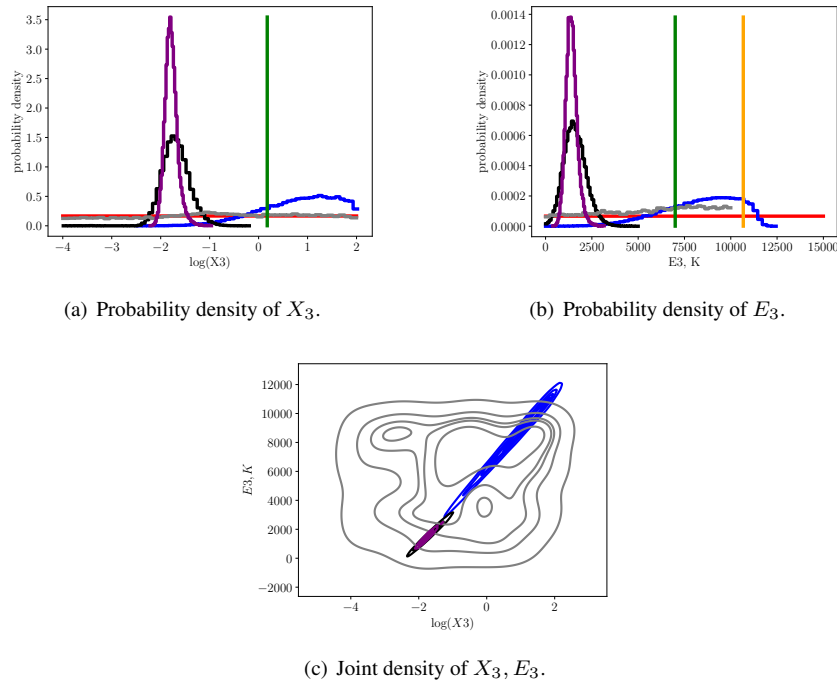


Figure 6.5: Probability densities of the efficiency and the activation energy of reaction 3. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al.

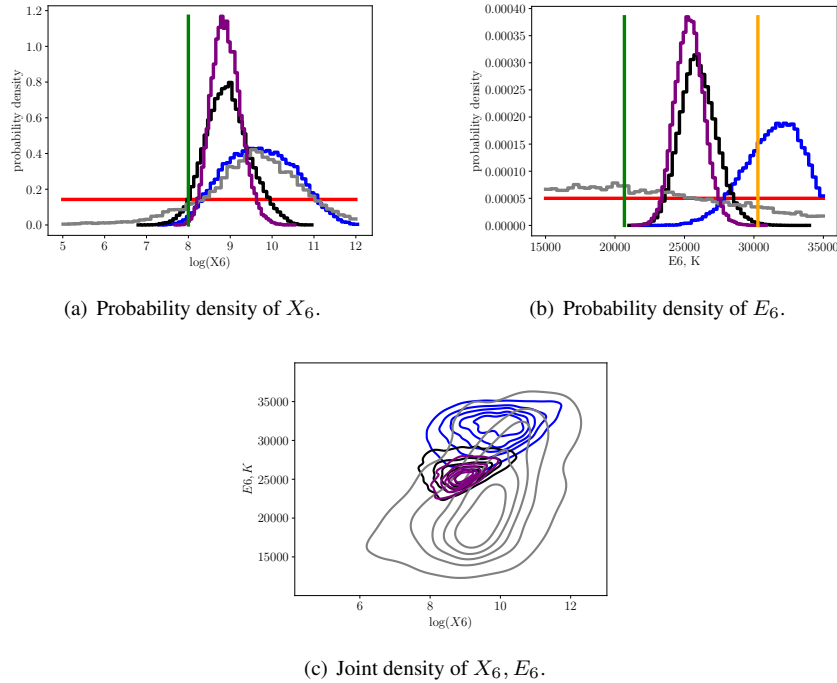


Figure 6.6: Probability densities of the efficiency and the activation energy of reaction 6. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al.

The marginal distributions of the calibrated parameters for the two recombination reactions are shown in Figures 6.7(a), 6.7(b), 6.8(a), and 6.8(b), while in Figures 6.7(c), and 6.8(c) the two joint distributions for the reaction efficiency and the activation energy are given. The LH mechanism (reaction 5, $N_s + N_s \rightarrow N_2 + 2s$) appears to be the dominant recombination mechanism, in agreement with what was observed by Murray et al. [142]. In fact, the ER reaction (reaction 4, $N_s + N \rightarrow N_2 + s$) appears to be poorly characterized with the molecular beam experiments considered. The distribution for the activation energy of the LH mechanism has a peak very close to the value proposed by Prata et al. when set A is used. Such peak shifts toward a lower value employing set B, closer to the value computed by Nieman et al. Also in this case, the parameters show a poor characterization when using set C. When the set D is employed, the reaction 4 is predicted to be unlikely. In fact, the joint distribution moves toward low values of the efficiency and high values of the activation energy.

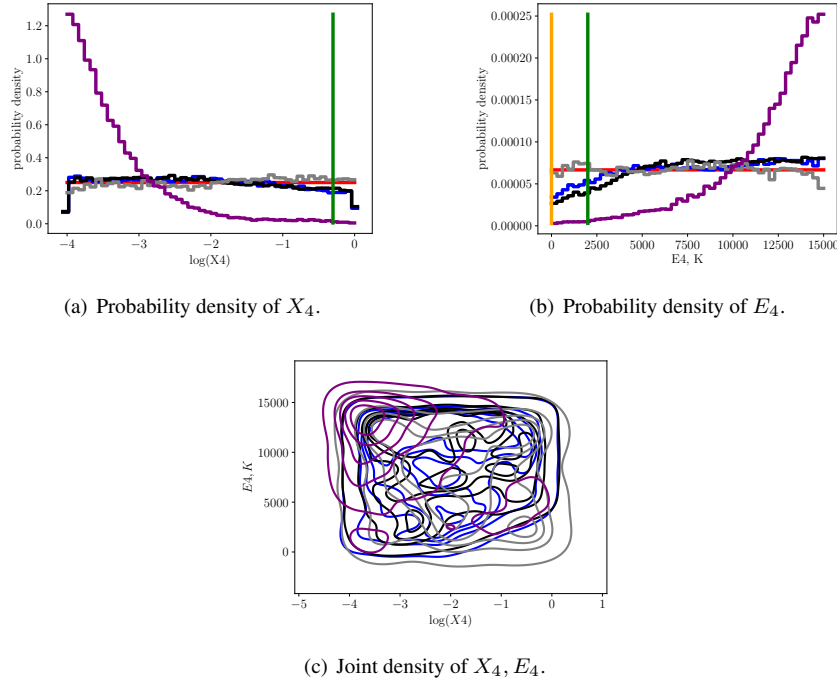


Figure 6.7: Probability densities of the efficiency and the activation energy of reaction 4. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al.

The activation energy of the absorption reaction, see Figure 6.9(a), is also well characterized, with a distribution tending towards small values, close to the lowest value proposed by Nieman et al. Interestingly, the more the model is refined the more this value approaches zero. Such a low value of activation energy allows the model to predict the very high values of recombination efficiencies experimentally observed in the molecular beam experiments.

Finally, the total active site density, see Figure 6.9(b), has a broad distribution with a peak close to the value employed by Prata when employing the set A. It moves towards small values using sets B and D, while is poorly characterized with set C, although tending to high values.

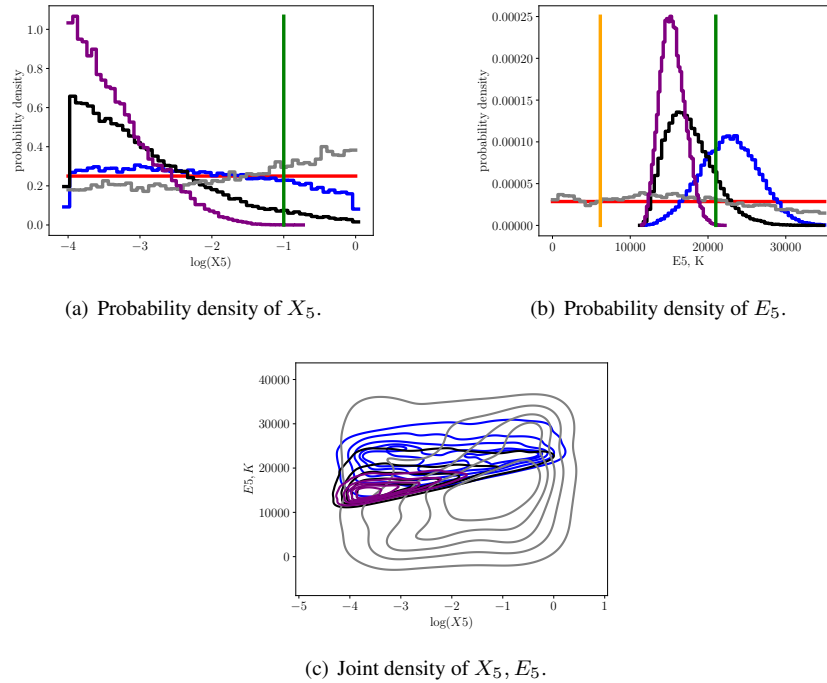


Figure 6.8: Probability densities of the efficiency and the activation energy of reaction 5. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values obtained by Prata et al., and **orange lines** the values by Nieman et al.

6.4.2 Model posterior predictions

From molecular beam data

The posterior distributions obtained by calibrating the model using set A and set B were propagated through the 0D model to obtain the efficiencies posterior predictions both at low and high pressure.

The model posterior predictions from set A are plotted in Figures 6.10(a) and 6.10(b). The molecular beam experimental uncertainties [142], and the marginals on the nitridation reaction obtained by del Val et al. [63] are shown in the figures. The latter were obtained by directly calibrating these efficiencies employing the same Plasmatron experiments. The predictions of the calibrated model are in perfect agreement with the molecular beam points used for the calibration. Large model uncertainties are encountered in the predictions at high pressure as no high-pressure calibration point was used. This suggests the necessity for higher-pressure experi-




Figure 6.9: Probability densities of the activation energy of reaction 1 and total active site density. **Blue histograms** refer to the parameters posterior using set A, **black histograms** using set B, **gray histograms** using set C, **purple histograms** using set D, **red lines** indicate the prior distributions, **green lines** the values by Prata et al., and **orange lines** the values by Nieman et al. (the three different lines corresponding to three distinct adsorption sites).

mental points to reduce this uncertainty. The same can be achieved, as we will see next, by using experimental points characterized by a surface coverage¹ similar to the one at high pressure.

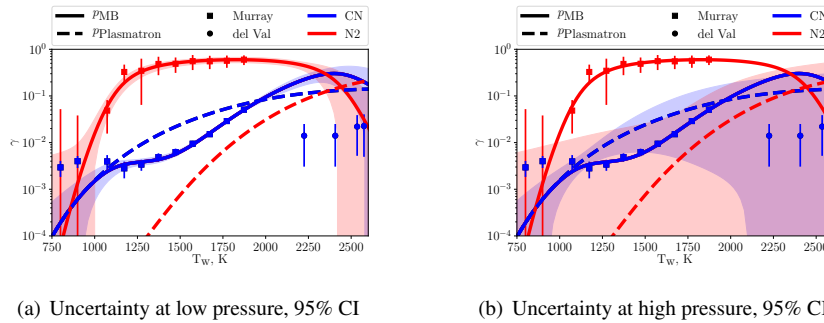


Figure 6.10: Predictions of the model calibrated with set A. The red and blue color indicates, respectively, the recombination and the nitridation efficiencies. The solid lines represent the mean of the prediction at molecular beam pressure (0.024 Pa), while the dashed lines at the Plasmatron pressure (1600 Pa). The shadow areas correspond to the prediction uncertainty (95% Confidence Interval (CI)), at low pressure on the left, and, on the right, at high pressure. The error bars with square as mean represent the molecular beam from Murray et al., while the error bars with circle as mean are the marginals obtained by del Val et al.

¹The surface coverage is defined as the ratio between the number of adsorbed atoms/molecules on a surface and the total number of active sites.

In fact, the uncertainty on the nitridation reaction at high pressure, shown in Figure 6.11(b), is drastically reduced when the model is calibrated with set B. Furthermore, the nitridation efficiencies at such high pressure exhibit excellent agreement with the marginals from del Val et al., suggesting that the two experiments can be accurately described through the use of the same calibrated model.

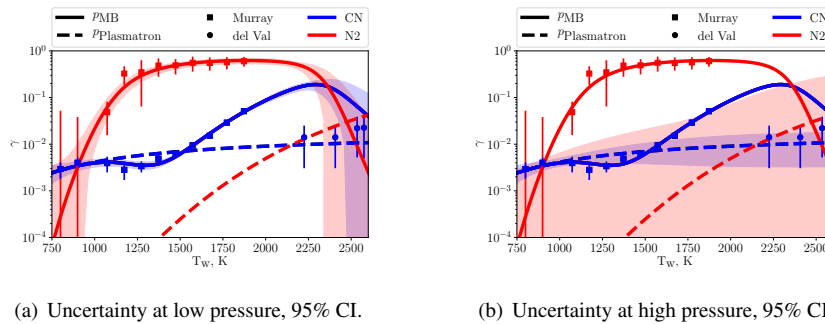


Figure 6.11: Predictions of the model calibrated with set B. The red and blue color indicates, respectively, the recombination and the nitridation efficiencies. The solid lines represent the mean of the prediction at molecular beam pressure (0.024 Pa), while the dashed lines at the Plasmatron pressure (1600 Pa). The shadow areas correspond to the prediction uncertainty (95% Confidence Interval (CI)), at low pressure on the left, and, on the right, at high pressure. The error bars with square as mean represent the molecular beam from Murray et al., while the error bars with circle as mean are the marginals obtained by del Val et al.

This improved characterization can be explained by looking at Figure 6.12, where the surface coverage at the two different pressures is plotted as a function of the surface temperature. As one can see, the value of the surface coverage at low-pressure/low-temperature is close to the almost flat value predicted at high pressure. Furthermore, the experimental uncertainty on the nitridation reaction efficiency at low temperatures is small. Thus, the low-temperature points contain information on the model behavior at high surface coverage.

On the other hand, the recombination efficiency at high pressure is still characterized by large variation as the relative experimental uncertainty at low-pressure/low-temperatures is high. The inclusion in the calibration of more accurate measurements at low temperatures would certainly help in decreasing the model uncertainty at high pressure.

Regarding the main trends, a sharp increase in the recombination efficiency is predicted at low pressure, in agreement with the assumptions of Murray et al. [142]. After a plateau, it decreases around 2250 K because the surface coverage sharply decreases at this temperature, as shown in Figure 6.12. The nitridation efficiency follows also the experimental trend: first, it increases, mostly driven by reaction 3 of

Table 2.1. After a brief downward trend due to reaction 3, the nitridation efficiency shows a steeper upward slope for higher surface temperatures, driven by reaction 6. For high temperatures, nitridation reaction efficiencies show a downward trend due to the decreasing surface coverage. At higher pressure, an almost flat trend in the nitridation probability is predicted, due to the nearly constant surface coverage.




Figure 6.12: Model predictions of the surface coverage using the set B for the calibration. Solid line refers to the mean predictions at molecular beam pressure and the dashed line at Plasmatron pressure. The uncertainty corresponding to the 95% CI is shown, for both the pressure, with shadow area.

The parameters' posteriors were also propagated through the surrogate models to obtain the uncertainties on the stagnation-point mass blowing rate for the cases G4, G5, G6, and G7. Their values, along with the experimental uncertainties, are shown in Figures 6.13(a) and 6.13(b). It can be seen that the propagated

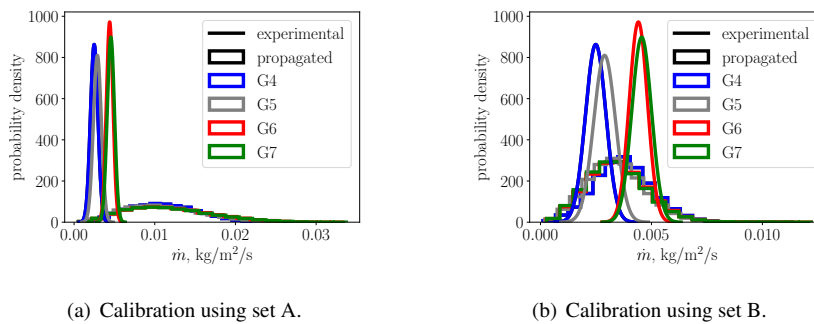


Figure 6.13: Mass blowing rate distributions. The continuous lines refer to the experimental uncertainty while the histograms show the propagated distributions. Blue, gray, red, and green colors indicate, respectively, the G4, G5, G6, and G7 case.

uncertainties are large when the model is calibrated using set A. They dramatically

decrease when the calibration is performed by means of the set B. Furthermore, such distributions share much of the support with their experimental counterparts, which is interesting as no Plasmatron data was used in the inference. However, the calibrated model appears to be less sensitive to the surface temperature (which increases from G4 to G7) than what was experimentally observed.

From Plasmatron data

Although Plasmatron data poorly inform the parameters, such an analysis allowed us to assess whether the model is able to correctly capture the experimental data. The posterior distribution of the model parameters calibrated with set C was propagated through the four surrogate models to obtain the posterior predictions of the stagnation-point mass blowing rate for the cases G4, G5, G6, and G7. These are shown in Figure 6.14, along with the experimental uncertainties. Overall, the propagated uncertainties are in good agreement with their experimental counterparts, and the upward trend with respect to the surface temperature (which increases from G4 to G7) is respected. Only for the G5 case, the mean prediction is slightly higher than the experimental one. This analysis verified that the model is capable of well explaining the Plasmatron experimental data.




Figure 6.14: Mass blowing rate distributions obtained propagating the parameters posteriors calibrated with set C. The continuous lines refer to the experimental uncertainty while the histograms show the propagated distributions. Blue, gray, red, and green colors indicate, respectively, the G4, G5, G6, and G7 case.

From molecular beam and Plasmatron data

The parameters posteriors obtained calibrating the model with set D were propagated through the 0D model. The posterior predictions are shown in Figure 6.15. While the main behavior of the model is similar to the one obtained when the calibration was performed with set B, the uncertainty at low pressures is slightly reduced at low and high temperatures. The uncertainties in both the nitridation and

the recombination reactions at high pressure exhibit a larger reduction, as a natural consequence of the inclusion of the high-pressure Plasmatron data.




Figure 6.15: Predictions of the model calibrated with set D. The red and blue color indicates, respectively, the recombination and the nitridation efficiencies. The solid lines represent the mean of the prediction at molecular beam pressure (0.024 Pa), while the dashed lines at the Plasmatron pressure (1600 Pa). The shadow areas correspond to the prediction uncertainty (95% Confidence Interval (CI)), at low pressure on the left, and, on the right, at high pressure. The error bars with square as mean represent the molecular beam from Murray et al., while the error bars with circle as mean are the marginals obtained by del Val et al.

The parameters' posterior was also propagated through the four surrogate models representing the case G4, G5, G6, and G7. The results are shown in Figure 6.15. As one can see, the posterior predictions of the stagnation-point mass blowing rate are contained in the relative experimental uncertainties. However, the predicted upward trend with respect to the surface temperature (which increases from G4 to G7) is negligible. Since the model cannot predict such a behavior jointly to the molecular beam data, the posterior predictions are shrunk in a reduced part of the experimental uncertainty, where all four measurements agree with each other.

Joint distributions of the reaction probabilities

Finally, 200 points were sampled from the posterior distributions relative to the calibrations performed using set B, C, and D. They were propagated directly through the CFD model, for the conditions G4, G5, G6, and G7. The numerical results were post-processed, using the Equation 2.81 and 2.82, to obtain the equivalent probability of the nitridation and the recombination reactions. The posterior joint distributions (γ_{CN} , γ_{N2}) for the three different calibrations and for the four different cases are shown in Figure 6.17. On the same plots, the joint distributions obtained by del Val et al. [64] are shown. These were obtained by performing the calibration directly on these efficiencies.

First of all, we can notice the excellent agreement between the joint distributions




Figure 6.16: Mass blowing rate distributions obtained propagating the parameters posteriors calibrated with set D. The continuous lines refer to the experimental uncertainty while the histograms show the propagated distributions. Blue, gray, red, and green colors indicate, respectively, the G4, G5, G6, and G7 case.

from del Val et al. and the ones corresponding to the calibration using the set C. It makes sense because, although the two calibrated models are different, they attempt to predict the same behavior on the same calibration points. The joint distributions from the calibration on set B share part of the support with the distribution from del Val et al. It is an expected result, from what was discussed about the posterior predictions on the mass blowing rates, that turned out to be almost in agreement with the experimental values. Lastly, the joint distributions from the calibration on set D are included in the joint distribution from del Val et al.

Interestingly, as a result of calibrating the model using jointly the Plasmatron and the molecular beam data, we were able to reduce the posterior uncertainties on the two efficiencies with respect to the results by del Val et al. The largest reduction was observed with respect to the recombination reaction, whose value is predicted to not exceed 0.1.

6.5 Summary

In this chapter, we presented the results of the calibration of an FRC model for carbon ablation. Specifically, we considered the reactions involving atomic nitrogen of the state-of-the-art ACA model.

First, we performed deterministic numerical simulations of a Plasmatron experiment. The ablative behavior of the sample was described by means of the ACA model, which turned out to predict lower stagnation-point mass blowing rates than what was experimentally observed.

For this reason, we calibrated the model using a Bayesian framework. Based on preliminary investigations, we decided to characterize the pre-exponential factors and the activation energies of selected reactions, and the total active site density.




Figure 6.17: Joint distribution of the reaction probabilities.

Four different calibrations were performed. The first two employed the molecular beam data for the inference and the Plasmatron data to assess the model predictive capability at high pressure. The third inference relied only on the Plasmatron data, while the last on all the observations jointly. Regarding the numerical models, we used the 0D reactor proposed by Prata et al. [1] to explain the molecular beam data, while a surrogate model, built on CFD simulations equipped with the specific ablative BC, was employed to describe the Plasmatron experiment. Since the deterministic simulations highlighted negligible numerical error associated with the dimension reduction contained in the stagnation-line code, we decided to train the surrogate model based on stagnation-line simulations. To further increase the efficiency of the method, the surrogate model was built for each experimental point individually, after having performed an appropriate dimension reduction of the input space.

Overall, molecular beam data were found to be highly informative for calibrating the FRC model. Contrary, Plasmatron data appeared to be inadequate to infer such a complex model. However, the inclusion of these data in the calibration process was important to provide information at pressures typical of the targeted

applications. Specifically, both the nitridation reactions turned out to be well characterized with the data at hand, especially when exploiting together information from molecular beam and Plasmatron experiments. Regarding the recombination reactions, the LH mechanism appeared to be the dominant mechanism, especially when the Plasmatron data were combined with the molecular beam ones. In fact, the analysis suggested the ER mechanism to be unlikely. Also the activation energy of the adsorption process was found to be well characterized. Its distribution moved towards the lowest value suggested by Nieman et al. [226] as the observation data set is enriched. Finally, the probability of the total number of active sites also tended towards low values, especially when including low-temperature points from the molecular beam experiments.

The parameters posteriors obtained for each set of calibration were then propagated through both the 0D model and the surrogates models representing the four conditions in the Plasmatron experiment. Interestingly, when the model is calibrated using all the molecular beam points, the propagated mass blowing rates share much of the support with their experimental counterparts. This improved calibration was explained by the fact that high-pressure and low-pressure/low-temperature conditions are characterized by a very similar surface coverage. Thus, the molecular beam points at low pressure and low temperature are highly informative for the high-pressure behavior. When the model was calibrated only on Plasmatron data, the propagated uncertainty on the mass blowing rates compared almost perfectly with the experimental uncertainties. This suggested that the FRC can well explain the Plasmatron data. When the model was calibrated using both molecular beam and Plasmatron data a good agreement was found with all the measurements used for the calibration. However, the upward trend of the mass blowing rate with respect to the surface temperature appeared to be much less pronounced than what was experimentally observed.

Finally, we compared the joint distributions of the post-processed phenomenological efficiencies with the values obtained by del Val et al. [64]. Excellent agreement was found when the model was calibrated only on Plasmatron data. When using all the measurements, the joint distributions are a subset of the ones obtained by del Val et al. [64]. Interestingly, the use of a higher-fidelity ablation model and the consequent possibility to calibrate it on different experiments, allowed us to reduce the posterior uncertainties obtained by del Val et al. at high pressures.

The calibrated model, and the relative uncertainties, presented in this chapter can be used in the future for predicting the ablative behavior of a carbon surface exposed to a nitrogen flow, for example in Titan atmospheric entry applications. We remind that in the last analysis, we used all the data to infer the model, leaving no experimental point for validation. Future perspectives include the assessment of the calibrated model including new Plasmatron data obtained at different pressures.

Another interesting direction would be to calibrate the whole ACA model, including those reactions involving atomic oxygen.

7

Conclusions

This thesis focused on the prediction and calibration of atmospheric entry flows. Their description is rather complex. In fact, they might include strong shocks, high-temperature effects, homogeneous and heterogeneous finite-rate chemistry, etc. Hence, accurate predictions can be achieved by employing high-fidelity CFD solvers. The physical models embedded within, and their parameters, are generally affected by severe uncertainties. For this reason, a rigorous quantification of the uncertainties is necessary. Thus, we proposed to develop tools and methodologies to obtain accurate predictions of atmospheric entry conditions using high-fidelity simulations and state-of-the-art experimental data.

Following, we summarize the main contributions of this thesis and we provide some potential perspectives.

The **first** objective was **to produce high-fidelity reference solutions**. More precisely, we targeted both in-flight and on-ground applications, ranging from hypersonic to subsonic regimes, where the flow could have been assumed as steady, laminar, and axisymmetric. In this regard, the following accomplishments were achieved:

- A1: Accurate numerical representations were obtained by **employing the US3D solver**, which is a three-dimensional FV code, specifically designed for hypersonic applications. Thus, it is expected to describe complex flow features, such as shock structures, in a robust manner. Additionally, it is equipped with a tailoring routine to automatically align the mesh to the shock, improving

the accuracy of those solutions characterized by strong shocks.

We **coupled US3D with the Mutation⁺⁺ library**, which provides state-of-the-art closures for homogeneous chemistry and transport properties. Several test cases have been presented for verifying the coupling. Particularly, it was shown that such a coupling led to an improved agreement between the numerical and experimental heat flux predicted on the surface of a cylinder exposed to a hypersonic flow. However, we also noticed applications, such as the one described in Chapter 5, where the use of the coupling with Mutation⁺⁺ led to convergence issues. This is probably due to the computation of the diffusive fluxes. In fact, the molar fraction gradient at the interface, needed to compute the diffusive fluxes, was approximated using finite-difference, instead of more accurate Green-Gauss or weighted least square approaches. In all the applications targeted in this thesis, we described catalytic or ablative materials exposed to reacting flows. For this reason, **we introduced in US3D specific BCs to enable the solution of surface mass and energy balances** through Mutation⁺⁺. Furthermore, **we included in the library a state-of-the-art FRC model**, which captures both the temperature and pressure dependency of the ablative reactions. Also in this case, we presented several verification test cases to assess the goodness of the implementation of the BCs in US3D and of the FRC model in Mutation⁺⁺. The latter was also validated against experimental data.

We consistently used the US3D solutions to assess the accuracy of those predictions obtained by means of lower-fidelity approximations and decide whether they were sufficiently accurate to perform UQ studies.

The **second objective was to investigate the benefits and the challenges associated with the systematic use of a high-fidelity solver in a UQ framework**. Based on the applications targeted, we reached the following accomplishments:

- A2 In Chapter 4, we presented a methodology **to balance the numerical error associated with the use of a given mesh and the problem-related uncertainties to produce an optimal representation for a given computational budget**. We applied it to characterize the surface pressure and heat flux experienced by the EXPERT vehicle in a trajectory point of its atmospheric entry. The flow was described by means of two-dimensional axisymmetric high-fidelity computations coupled to a SEB catalytic BC.

It was shown that a **relatively coarse mesh can be employed to accurately estimate all the QoIs**. In fact, the associated numerical uncertainty was at least one order of magnitude lower than the variability induced by the uncertainties in the BCs. It was also highlighted the importance of **automatically aligning the mesh to the shock** at each training point to prevent that numerical errors, arising across it, corrupt the estimation of the stagnation-point

heat flux, resulting in an inaccurate surrogate model.

In this case, resorting to two-dimensional axisymmetric solutions allowed for **computing off-stagnation-point values of the heat flux and of the pressure**. They can be potentially employed in an inverse problem by comparing them to the measurements provided by off-stagnation-point sensors, as in Cortesi et al. [57].

We remark that in this study we decided to run a simulation for each training point and each mesh and that we constructed each surrogate model independently. A much more efficient methodology would have been to fuse the information coming from the different mesh levels to reduce the number of computations required on the most refined grids.

A3: As shown in Chapter 5, **we devised an adaptive/multi-fidelity strategy** to reduce the computational cost associated with the construction of a surrogate model from high-fidelity simulations. We employed it to **describe an under-expanded high-enthalpy jet** obtained in the Plasmatron facility.

We first constructed a **low-fidelity model** to compute the pressure and heat flux at the stagnation point of a probe exposed to such a jet. However, it **predicted significantly lower values** compared to the high-fidelity US3D reference solution. Hence, **US3D high-fidelity computations were necessary** to accurately perform the UQ analysis. Furthermore, the mesh required to correctly compute all the QoIs was too expensive to be directly used to build a surrogate model. We leveraged the low-fidelity representation and the solutions from coarser meshes to limit the number of high-fidelity computations to those required for accuracy reasons. Particularly, we have shown that **fusing this information by means of a hierarchical Kriging strategy dramatically improves the quality of the surrogate model** with respect to a Kriging trained only on high-fidelity points, with little increase in the computational burden. Further improvement in the model accuracy was achieved by employing an **adaptive sampling strategy**.

We have also highlighted that the high-fidelity solution permits to resolve the **compressible features** of the under-expanded jet (jet boundary, barrel shock, and detached shock).

The **third** objective was to **use the experimental data to reduce the uncertainty on the prediction of targeted QoIs**. In this regard:

A4: Also in Chapter 5, **we characterized the conditions at the inlet of the sonic nozzle, and the catalytic efficiencies of the probes**, of the under-expanded high-enthalpy jet by means of the measurements of the stagnation-point pressure and heat flux, and of the mass flow in the nozzle. A Bayesian framework, coupled with the previously-mentioned adaptive/multi-fidelity surrogate model, was used to solve the inverse problem.

We have shown that we obtained **pronounced peaks** for the probability density of the **total pressure and temperature** at the inlet of the nozzle, and for the **efficiency of the nitrogen recombination** reaction. However, given the limited amount of atomic oxygen in the flow, we were not able to determine its recombination reaction.

Another important contribution concerned the **assessment of the variability of the main features of the under-expanded jet** with respect to the posterior uncertainties on the reservoir conditions, yielding a robust characterization of the flow.

A5: In Chapter 6, we **inferred the rates of the elementary reactions of a detailed ablation model from Plasmatron and molecular beam experiments**. In this case, the low-fidelity stagnation-line code was found to be accurate to simulate the Plasmatron experiments. Additionally, a 0D approach was sufficient to describe the molecular beam experiments. For this reason, such **efficient computations** were used **in the UQ analysis**.

On the other hand, a **high-fidelity ablation model** was needed to jointly **describe both the experiments**. It is characterized by a high number of input parameters. We built a surrogate model to approximate the stagnation-line computations. Given the high dimension of the input space, its construction required a large number of training points. To reduce this number, and have a more efficient surrogate model, we resorted to a **sensitivity analysis to reduce the input dimension**.

The surrogate model and the 0D solver were used to perform a Bayesian analysis based both on Plasmatron and molecular beam data. Overall, we found that combining these information led to an **accurate characterization of most of the model parameters**. Furthermore, the so-calibrated model was **capable of describing both the experimental observations**. However, we highlighted that the model predictions **at high pressure were less sensitive to the surface temperature** than what was experimentally observed.

Perspectives and future developments

In this section, we outline some perspectives and future developments that naturally emerge from the achievements and the limitations of the tools and methodologies developed in this thesis.

The multi-fidelity strategy developed in Chapter 5 might be employed to characterize flows requiring high-fidelity simulations. Furthermore, the lessons learned from the performed analysis might drive future characterizations in including additional flow features for calibration and validation purposes. In this regard:

- P1: A current effort is being performed to characterize an experimental campaign conducted in the Plasmatron facility in which a supersonic jet was obtained by means of a convergent-divergent nozzle. The setup of the experimental campaign allowed for the measurement of the stagnation-point heat flux and pressure, the mass flow rate, the pressure along the nozzle, and of the shock-stand-off distance. These measurements will be used in the inference/validation process.
- P2: The totality of the past studies concerning the characterization of the efficiencies of the ablative reactions from Plasmatron data were performed employing solely stagnation-point measurements, justifying the use of efficient low-fidelity stagnation-line computations. Very recently, Fagnani et al. [132] developed an experimental methodology for obtaining the surface mapping of the recession rate and of the temperature. These data can be exploited, along with the methodologies developed in this thesis, to further enrich the inverse problem.

Future works might also be devoted to improving the methodology. In fact, we focused the attention on applications where the flow could have been assumed as *steady*, *laminar*, and *axisymmetric*. Obviously, we were far from covering the whole spectrum of regimes that might be encountered in hypersonic flows. A broader and more general future perspective includes the characterization of more complex flows, for instance, turbulent and unsteady, for which *higher resolution* is required, introducing new challenges. For example:

- P3: One can characterize the flow in Longshot experiments. Although such a flow does not present several high-temperature physical features encountered in the applications treated in this manuscript, such as gas chemistry and GSI, it is expected to be turbulent and in thermal non-equilibrium. Furthermore, the reservoir conditions sharply decrease during the flow expansion and such an unsteady feature should be accounted for. All in all, a robust characterization should be performed accounting for turbulent, thermal non-equilibrium, and unsteady effects. The major challenges foreseen in this application regard the correct modeling of the physics. In a previous deterministic investigation by Geratz et al. [227], we analyzed the sensitivity of the results to real gas and RANS models. None of the models employed was able to correctly reproduce the experimentally observed Mach number at the end of the nozzle. Such a discrepancy might be driven by an over-simplification contained in the RANS model or by unsteady effects, i.e. the decay of the reservoir pressure. This highlighted the necessity of a higher fidelity in the CFD computations coupled to a UQ strategy. Concerning this point, an improved methodology should be employed for the construction of the surrogate model, which should be designed to also capture the time dependency of the response.

Regarding the characterization of the FRC ablation model treated in Chapter 6, we saw that all the experimental data were used to calibrate the model, leaving no information for its validation.

P4: Recently, a new experimental campaign was conducted in the Plasmatron facility to assess the pressure dependence of the nitridation reaction [228]. These new experimental observations can be analyzed and employed to validate the previously calibrated FRC model.

We also saw in Chapter 6 that the calibrated FRC model was not able to capture the upward trend of the stagnation-point mass blowing rate with respect to the surface temperature. This might be due to the fact that we are modeling two distinct carbon surfaces with the same model, potentially introducing a model inadequacy.

P5: The Bayesian analysis can be enriched by explicitly accounting for the model inadequacy in the calibration of the FRC model. In this regard, Leoni [229] has shown that a '*Full Maximum A Posteriori*' approximation is able to accurately estimate the model discrepancy. Future works might be oriented at including such a methodology in the calibration procedure.

Finally, we focused only on the calibration of the reactions of the ACA model involving atomic nitrogen. However, the full model consists also of reactions involving atomic oxygen.

P6: Future works can be devoted to the characterization of the full model. Although it might seem a trivial extension of the current work, we expect that new challenges will arise. In fact, the total dimension of the parameters to calibrate will more than double, as well as the experimental points for the Bayesian analysis. The construction of the resulting MCMC is foreseen to evolve extremely slowly, and more efficient tools might be necessary to solve the inverse problem.



US3D-Mutation⁺⁺ Coupling Verification Test Cases

In Section 4.4, we described the coupling that we have performed between the US3D solver and the Mutation⁺⁺ library. A series of verification test cases, summarized in Table 4.1, were performed to assess the goodness of the coupling. The test '*2D hypersonic flow over a probe*' was presented in Section 4.4. In this section, we revise the results associated with those other test cases that were not shown in Section 4.4.

A.1 0D Reactor

This first test case verified that the chemical source term was correctly computed. To this end, we considered an air mixture composed of five species, $S = [\text{N}_2, \text{O}_2, \text{NO}, \text{N}, \text{O}]$, in an adiabatic box. Starting from the chemical non-equilibrium initialization provided in Table A.1, the system was left free to time-march towards the equilibrium condition according to the chemical rates by Park [5]. The solution provided by the US3D solver with its native library, referred to as 'US3D native', was compared to the one obtained with US3D coupled to Mutation⁺⁺, referred to as 'US3D+M⁺⁺'. The code-to-code comparison, shown in Figure A.1, shows excellent agreement, verifying the implementation.

ρ , kg/m ³	T , K	u , m/s	$Y(\text{N}_2)$	$Y(\text{O}_2)$
0.01	7000	0.0	0.767	0.233

Table A.1: Initial conditions for the 0D reactor case.




Figure A.1: Evolution in time of the mass fractions in an air mixture composed of five species in the 0D reactor case. The code-to-code comparison shows excellent agreement.

A.2 1D Diffusion Problem

The second test case verified the coupling for the transport properties. The simulation consisted of a 1D tube with isothermal end walls at different temperatures. The boundary and initial conditions are provided in Table A.2. The tube is 3 mm long. Also in this case, the solution obtained with the coupling (viscosity and thermal conductivity were computed with the Chapman-Enskog expansion and the Eucken correction, while diffusive velocities by means of the Stefan-Maxwell formulation) is compared to the one produced by US3D with its native library (viscosity and thermal conductivity were computed using the Gupta-Yos mixture rule and the Eucken correction, while diffusive velocities using the SCEBD formulation). We remark that, although the transport models employed in the two simulations are different, they are expected to provide results in reasonable agreement at the temperatures considered ($T_{\max} = 4800$ K).

ρ , kg/m ³	T , K	T_{left} , K	T_{right} , K	u , m/s	$Y(\text{N}_2)$	$Y(\text{O}_2)$
0.02	1000	800	4800	0.0	0.767	0.233

Table A.2: Boundary and initial conditions for the 1D diffusion case.

In this test case, the temperature gradient leads to chemical reactions, which, in turn, drive the mass diffusion. Temperature and mass fractions distributions along the tube are shown in Figure A.2. As one can see, also in this case the agreement between the two codes is excellent, verifying the coupling for transport properties.




Figure A.2: Comparison of the temperature (left) and mass fractions (right) along the tube for the 1D diffusion case. The solutions of the two codes perfectly agree with each other.

A.3 1D Catalytic Diffusion Problem

This test case verified the implementation of the catalytic BC for a simple $[N_2, N]$ binary mixture along a 1D 0.2 m long tube. The simulation BCs are given in Table A.3. One side of the tube is characterized by a reservoir condition, while the other by a catalytic isothermal BC. Nitrogen recombination reactions were imposed through a phenomenological approach. Four different probabilities, reported in Table A.3, were used.

p , Pa	T , K	T_{wall} , K	u , m/s	$Y(N_2)$	$Y(N)$	γ_N
100	3000	3000	0.0	0.0	1.0	[0.001, 0.01, 0.1, 1.0]

Table A.3: BCs for the 1D catalytic diffusion case.

The results obtained by US3D coupled to Mutation⁺⁺ were compared to the analytical solution proposed by Bariselli et al. [230]. They are shown in Figure A.3. Also in this case, the agreement is excellent.

A.4 Nitrogen flow over an ablative surface

This last test case was performed to verify the GSI coupling with respect to the ablative reactions. To this end, the G5 subsonic Plasmatron experiment, detailed




Figure A.3: N_2 mass fractions for different recombination coefficients γ for the 1D catalytic diffusion problem. The code solution perfectly agrees with its analytical counterpart.

in Section 2.4.1, was simulated both with US3D coupled with Mutation⁺⁺ (below indicated as ‘US3D+M⁺⁺’) and with the stagnation-line code (below indicated as ‘stagline’).

The simulation consists of an ionized nitrogen flow over an isothermal reactive surface. Specifically, we modeled the nitridation surface reaction ($N + C_b \rightarrow CN$) using a phenomenologic approach, whose reaction probability was taken from Helber et al. [3]:

$$\gamma_{CN} = 7.91 \cdot 10^{-2} \exp\left(-\frac{5663}{T_w}\right). \quad (\text{A.1})$$

The test free stream conditions were imposed at the inlet and the experimental chamber pressure at the outlet. On the sample surface, a SMB was solved. A symmetry BC was applied to all the other faces. The experimental conditions are reported in Section 2.4.1. The results from US3D were extracted along the stagnation line and compared to the ones returned by the stagnation-line code. As one can see in Figure A.4 the two solutions agree well both for the temperature and the mass fractions profile in the BL. The mass blowing rate at the stagnation point is 0.003 410 kg/m²/s for the US3D computation and 0.003 414 kg/m²/s for the stagnation-line code, resulting in an almost null error.

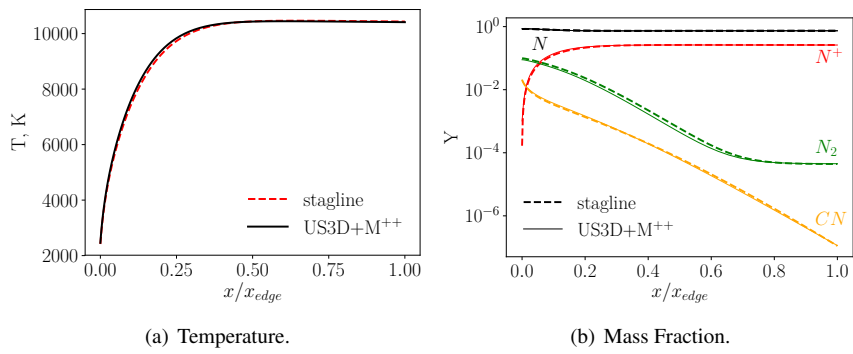


Figure A.4: Comparison of the QoIs along the stagnation line obtained with the stagnation-line code (dashed line) and the US3D solver (solid line). The two solutions agree well both for the temperature and the mass fractions profile in the BL.

B

Transport matrices

In Chapter 2 we described the modeling of the transport fluxes in the Navier-Stokes equations. Following we provide the expressions of the transport matrices. Their components are:

$$G_{ij}^V = -\frac{X_i X_j}{D_{ij}}(1 + \phi_{ij}), \quad i \neq j, \quad (\text{B.1})$$

$$G_{ii}^V = \sum_{\substack{j \neq i \\ j \in \mathcal{H}}} \frac{X_i X_j}{D_{ij}}(1 + \phi_{ij}) + \left(\frac{T_e}{T}\right)^2 \frac{X_i X_e}{D_{ie}}(1 + \phi_{ie}), \quad (\text{B.2})$$

$$G_{ij}^\mu = G_{ji}^\mu = \frac{X_i X_j}{n D_{ij}} \frac{1}{m_i + m_j} \left(\frac{6}{5} \frac{Q_{ij}^{(2,2)}}{Q_{ij}^{(1,1)}} - 2 \right), \quad i \neq j, \quad (\text{B.3})$$

$$G_{ii}^\mu = \sum_{\substack{j \neq i \\ j \in \mathcal{H}}} \frac{X_i X_j}{n D_{ij}} \frac{1}{m_i + m_j} \left(\frac{6}{5} \frac{m_j}{m_i} \frac{Q_{ij}^{(2,2)}}{Q_{ij}^{(1,1)}} + 2 \right) + \frac{X_i^2}{\mu_i}, \quad (\text{B.4})$$

$$G_{ij}^{\Lambda^T} = G_{ji}^{\Lambda^T} = \frac{1}{25k_B} \frac{X_i X_j}{n D_{ij}} \frac{m_i m_j}{(m_i + m_j)^2} \cdot \left(16 \frac{Q_{ij}^{(2,2)}}{Q_{ij}^{(1,1)}} + 12 \frac{5Q_{ij}^{(1,2)} - 4Q_{ij}^{(1,3)}}{Q_{ij}^{(1,1)}} - 55 \right), \quad i \neq j, \quad (\text{B.5})$$

$$\begin{aligned}
G_{ii}^{\lambda\Gamma} = & \frac{1}{25k_{\text{B}}} \sum_{\substack{j \neq i \\ j \in \mathcal{H}}} \left(30 \frac{m_i^2}{m_j^2} + 25 \frac{m_j}{m_i} - 12 \frac{m_j}{m_i} \frac{5Q_{ij}^{(1,2)} - 4Q_{ij}^{(1,3)}}{Q_{ij}^{(1,1)}} + 16 \frac{m_i}{m_j} \frac{Q_{ij}^{(2,2)}}{Q_{ij}^{(1,1)}} \right) \\
& \cdot \frac{X_i X_j}{nD_{ij}} \frac{m_i m_j}{(m_i + m_j)^2} + \frac{4}{15k_{\text{B}}} \frac{X_i^2 m_i}{\mu_i}.
\end{aligned} \tag{B.6}$$

Symbols ϕ_{ij} and ϕ_{ie} are correction functions, whose description can be found in [89, 90], and \mathcal{H} stands for the set of index of the heavy species.

References

- [1] K. Prata, T. Schwartzentruber, and T. Minton. *Air–Carbon Ablation Model for Hypersonic Flight from Molecular-Beam Data*. *AIAA Journal*, 60(2):627–640, 2022.
- [2] A. Fagnani. *Development of measurement techniques and study of the aerothermal response of space debris materials to atmospheric entry plasmas*. PhD thesis, Vrije Universiteit Brussel, 2023.
- [3] B. Helber, A. Turchi, and T. Magin. *Determination of active nitridation reaction efficiency of graphite in inductively coupled plasma flows*. *Carbon*, 125:582–594, 2017.
- [4] J. Anderson. *Hypersonic and High Temperature Gas Dynamics*. McGraw Hill, New York, 1989.
- [5] C. Park. *Nonequilibrium Hypersonic Aerothermodynamics*. Springer Berlin Heidelberg, 1990.
- [6] O. Uyanna and H. Najafi. *Thermal protection systems for space vehicles: A review on technology development, current challenges and future prospects*. *Acta Astronautica*, 176:341–356, 2020.
- [7] O. Chazot and F. Panerai. *High-Enthalpy Facilities and Plasma Wind Tunnels for Aerothermodynamics Ground Testing*, pages 329–342. 2015.
- [8] B. Richards and K. Enkenhus. *Hypersonic testing in the VKI longshot free-piston tunnel*. *AIAA Journal*, 8(6):1020–1025, 1970.
- [9] B. Bottin, O. Chazot, M. Carbonaro, V. Haegen, and S. Paris. *The VKI Plasmatron Characteristics and Performance*. page 26, 04 2000.
- [10] A. Kolesnikov. *Conditions of simulation of stagnation point heat transfer from a high-enthalpy flow*. *Fluid Dynamics*, 28:131–137, 1993.
- [11] G. Candler, H. Johnson, I. Nompelis, V. Gidzak, P. Subbareddy, and M. Barnhardt. *Development of the US3D Code for Advanced Compressible and Reacting Flow Simulations*. In 53rd AIAA Aerospace Sciences Meeting. American Institute of Aeronautics and Astronautics, 2015.

- [12] D. Hash, J. Olejniczak, M. Wright, D. Prabhu, M. Pulsonetti, B. Hollis, P. Gnoffo, M. Barnhardt, I. Nompelis, and G. Candler. *FIRE II Calculations for Hypersonic Nonequilibrium Aerothermodynamics Code Verification: DPLR, LAURA, and US3D*. In 45th AIAA Aerospace Sciences Meeting and Exhibit. American Institute of Aeronautics and Astronautics, 2007.
- [13] O. Schroeder, J. Brock, E. Stern, and G. Candler. *A coupled ablation approach using Icarus and US3D*. In AIAA Scitech 2021 Forum. American Institute of Aeronautics and Astronautics, 2021.
- [14] L. Trevino and G. Candler. *Numerical Simulation of Regular Surface Patterns on Sublimating Ablative Materials*. In 53rd AIAA Aerospace Sciences Meeting. American Institute of Aeronautics and Astronautics, 2015.
- [15] M. Barnhardt, G. Candler, and M. MacLean. *CFD Analysis of CUBRC Base Flow Experiments*. In 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition. American Institute of Aeronautics and Astronautics, 2010.
- [16] J. Brock, P. Subbareddy, and G. Candler. *Detached-Eddy Simulations of Hypersonic Capsule Wake Flow*. AIAA Journal, 53(1):70–80, 2015.
- [17] M. MacLean, E. Mundy, T. Wadhams, M. Holden, M. Barnhardt, and G. Candler. *Experimental and Numerical Study of Laminar and Turbulent Base Flow on a Spherical Capsule*. In 47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition. American Institute of Aeronautics and Astronautics, 2009.
- [18] E. Stern, V. Gidzak, and G. Candler. *Estimation of Dynamic Stability Coefficients for Aerodynamic Decelerators Using CFD*. In 30th AIAA Applied Aerodynamics Conference. American Institute of Aeronautics and Astronautics, 2012.
- [19] T. Schwartzentruber. *Nonequilibrium Gas-Surface Interactions at High Temperature*. Technical report, Air Force Research Laboratory, 2020.
- [20] F. Milos and Y. Chen. *Comprehensive model for multicomponent ablation thermochemistry*. In 35th Aerospace Sciences Meeting and Exhibit. American Institute of Aeronautics and Astronautics, 06 1997.
- [21] C. Park. *Effects of atomic oxygen on graphite ablation*. AIAA Journal, 14(11):1640–1642, 1976.
- [22] S. Zhuktov and T. Abe. *Viscous Shock-Layer Simulation of Airflow past Ablating Blunt Body with Carbon Surface*. Journal of Thermophysics and Heat Transfer, 13(1):50–59, 01 1999.

- [23] C. Alba, R. Greendyke, and J. Marschall. *Development of a Nonequilibrium Finite-Rate Ablation Model for Radiating Earth Reentry Flows*. *Journal of Spacecraft and Rockets*, 53(1):98–120, 2016.
- [24] S. Poovathingal, T. Schwartzentruber, V. Murray, T. Minton, and G. Candler. *Finite-Rate Oxidation Model for Carbon Surfaces from Molecular Beam Experiments*. *AIAA Journal*, 55(5):1644–1658, 2017.
- [25] C. Alba, R. Greendyke, S. Lewis, R. Morgan, and T. McIntyre. *Numerical Modeling of Earth Reentry Flow with Surface Ablation*. *Journal of Spacecraft and Rockets*, 53(1):84–97, 2016.
- [26] G. Candler. *Nonequilibrium Processes in Hypervelocity Flows: An Analysis of Carbon Ablation Models*. In 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition. American Institute of Aeronautics and Astronautics, 01 2012.
- [27] L. Zhang, D. Pejaković, J. Marschall, M. Dougherty, and D. Fletcher. *Laboratory Investigation of the Active Nitridation of Graphite by Atomic Nitrogen*. *Journal of Thermophysics and Heat Transfer*, 26(1):10–21, 2012.
- [28] M. MacLean, J. Marschall, and D. Driver. *Finite-Rate Surface Chemistry Model, II: Coupling to Viscous Navier-Stokes Code*. In 42nd AIAA Thermophysics Conference. American Institute of Aeronautics and Astronautics, 2011.
- [29] J. Sacks, W. Welch, T. Mitchell, and H. Wynn. *Design and Analysis of Computer Experiments*. *Statistical Science*, 4(4):409–423, 1989.
- [30] B. Peherstorfer, K. Willcox, and M. Gunzburger. *Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization*. *SIAM Review*, 60(3):550–591, 2018.
- [31] B. Sudret, S. Marelli, and J. Wiart. *Surrogate models for uncertainty quantification: An overview*. In 2017 11th European Conference on Antennas and Propagation (EUCAP), pages 793–797, 03 2017.
- [32] M. Wright, D. Bose, and Y. Chen. *Probabilistic Modeling of Aerothermal and Thermal Protection Material Response Uncertainties*. *AIAA Journal*, 45(2):399–410, 2007.
- [33] G. Palmer, J. Olejniczak, and M. Wright. *Uncertainty Analysis of Aeroheating Predictions for Titan Entries*. In 44th AIAA Aerospace Sciences Meeting and Exhibit. American Institute of Aeronautics and Astronautics, 01 2006.

- [34] D. Bose, M. Wright, and T. Gökçen. *Uncertainty and Sensitivity Analysis of Thermochemical Modeling for Titan Atmospheric Entry*. 37th AIAA Thermophysics Conference, 06 2004.
- [35] S. Copeland, M. Mahzari, I. Cozmuta, and J. Alonso. *A Statistics-Based Material Property Analysis to Support Ablation Simulation UQ Efforts*. In 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA. American Institute of Aeronautics and Astronautics, 2012.
- [36] Y. Chen, T. Squire, B. Laub, and M. Wright. *Monte Carlo Analysis for Spacecraft Thermal Protection System Design*. In 9th AIAA/ASME Joint Thermophysics and Heat Transfer Conference. American Institute of Aeronautics and Astronautics, 2006.
- [37] S. Ghaffari, T. Magin, and G. Iaccarino. *Uncertainty Quantification of Radiative Heat Flux Modeling for Titan Atmospheric Entry*. In 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition. American Institute of Aeronautics and Astronautics, 2010.
- [38] S. Hosder and B. Bettis. *Uncertainty and Sensitivity Analysis for Reentry Flows with Inherent and Model-Form Uncertainties*. *Journal of Spacecraft and Rockets*, 49(2):193–206, 2012.
- [39] X. Wang, C. Yan, S. Ju, Y. Zheng, and J. Yu. *Uncertainty analysis of laminar and turbulent aeroheating predictions for Mars entry*. *International Journal of Heat and Mass Transfer*, 112:533–543, 2017.
- [40] N. Villedieu, F. Panerai, O. Chazot, and T. Magin. *Uncertainty Quantification For Gas-Surface Interaction In Plasmatron Facility*. In L. Ouwehand, editor, 7th European Symposium on Aerothermodynamics, volume 692 of *ESA Special Publication*, page 149, 2011.
- [41] B. Lockwood and D. Mavriplis. *Gradient-based methods for uncertainty quantification in hypersonic flows*. *Computers and Fluids*, 85:27–38, 2013. International Workshop on Future of CFD and Aerospace Sciences.
- [42] A. Alexeenko, A. Weaver, R. Greendyke, and J. Camberos. *Flowfield Uncertainty Analysis for Hypersonic CFD Simulations*. In 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition. American Institute of Aeronautics and Astronautics, 2010.

- [43] T. West, S. Hosder, and C. Johnston. *Multistep Uncertainty Quantification Approach Applied to Hypersonic Reentry Flows*. Journal of Spacecraft and Rockets, 51(1):296–310, 2014.
- [44] T. West and S. Hosder. *Uncertainty Quantification of Hypersonic Reentry Flows with Sparse Sampling and Stochastic Expansions*. Journal of Spacecraft and Rockets, 52(1):120–133, 2015.
- [45] A. Brune, S. Hosder, and K. Edquist. *Uncertainty Analysis of Fluid-Structure Interaction of a Deformable Hypersonic Inflatable Aerodynamic Decelerator*. Journal of Spacecraft and Rockets, 53(4):654–668, 2016.
- [46] M. Santos, S. Hosder, and T. West. *Aerothermal Uncertainty Quantification of Deployable Entry Technologies Using Multi-Fidelity Modeling*. In ASCEND 2021. American Institute of Aeronautics and Astronautics, 2021.
- [47] M. Santos, S. Hosder, and T. West. "Multifidelity Turbulent Heating Prediction of Hypersonic Inflatable Aerodynamic Decelerators with Surface Scalping". Journal of Spacecraft and Rockets, 58(5):1325–1338, 09 2021.
- [48] K. Quinlan, J. Movva, E. Stein, and A. Kupresanin. *Leveraging Multi-Fidelity Aerodynamic Databasing to Efficiently Represent a Hypersonic Design Space*. In ASCEND 2021. American Institute of Aeronautics and Astronautics, 2021.
- [49] C. Sorensen, P. Valentini, and T. Schwartzentruber. *Uncertainty Analysis of Reaction Rates in a Finite-Rate Surface-Catalysis Model*. Journal of Thermophysics and Heat Transfer, 26(3):407–416, 2012.
- [50] F. Sanson, P. Congedo, N. Villedieu, O. Chazot, T. Magin, and F. Panerai. *Quantification of uncertainty on the catalytic property of reusable thermal protection materials from high enthalpy experiments*. Experimental Thermal and Fluid Science (EXP THERM FLUID SCI), 82, 12 2016.
- [51] A. Turchi, P. Congedo, B. Helber, and T. Magin. *Thermochemical ablation modeling forward uncertainty analysis—Part II: application to plasma wind-tunnel testing*. International Journal of Thermal Sciences, 118:510–517, 2017.
- [52] A. Turchi, P. Congedo, and T. Magin. *Thermochemical ablation modeling forward uncertainty analysis—Part I: Numerical methods and effect of model parameters*. International Journal of Thermal Sciences, 118:497–509, 2017.
- [53] A. Brune, T. West, and L. White. *Calibration Probe Uncertainty and Validation for the Hypersonic Material Environmental Test System*. Journal of Thermophysics and Heat Transfer, 34(2):404–420, 04 2020.

- [54] O. Georgii and H. Volker. *Sensitivity Analysis Study of Expanding Hypersonic Flows of Nitrogen and Oxygen*. In 9th European conference for aeronautics and space sciences. EUCASS association, 2022.
- [55] J. Tryoen, P. M. Congedo, R. Abgrall, N. Villedieu, and T.E. Magin. *Bayesian-based method with metamodels for rebuilding freestream conditions in atmospheric entry flows*. AIAA Journal, **52**(10):2190 – 2197, 2014.
- [56] A. Cortesi, P. Constantine, T. Magin, and P. Congedo. *Forward and backward uncertainty quantification with active subspaces: Application to hypersonic flows around a cylinder*. Journal of Computational Physics, 407:109079, 2020.
- [57] A. Cortesi. *Simulations numériques prédictives pour la reconstruction des conditions en amont dans les écoulements de rentrée atmosphérique*. PhD thesis, Université de Bordeaux, 2018.
- [58] J. Ray, S. Kieweg, D. Dinzl, B. Carnes, V. Weirs, B. Freno, M. Howard, T. Smith, I. Nompelis, and G. Candler. *Estimation of Inflow Uncertainties in Laminar Hypersonic Double-Cone Experiments*. AIAA Journal, **58**(10):4461–4474, 2020.
- [59] K. Chowdhary, C. Hoang, K. Lee, J. Ray, V. Weirs, and B. Carnes. *Calibrating hypersonic turbulence flow models with the HIFiRE-1 experiment using data-driven machine-learned models*. Computer Methods in Applied Mechanics and Engineering, 401:115396, 2022.
- [60] R. Upadhyay, K. Miki, O. Ezekoye, and J. Marschall. *Uncertainty quantification of a graphite nitridation experiment using a Bayesian approach*. Experimental Thermal and Fluid Science, **35**(8):1588–1599, 2011.
- [61] F. Sanson, F. Panerai, T. Magin, and P. Congedo. *Robust reconstruction of the catalytic properties of thermal protection materials from sparse high-enthalpy facility experimental data*. Experimental Thermal and Fluid Science, **96**, 03 2018.
- [62] A. del Val, O. Le Maître, T. Magin, O. Chazot, and P. Congedo. *A surrogate-based optimal likelihood function for the Bayesian calibration of catalytic recombination in atmospheric entry protection materials*. Applied Mathematical Modelling, 101:791–810, 2022.
- [63] A. del Val, O. Le Maître, P. Congedo, and T. Magin. *Stochastic calibration of a carbon nitridation model from plasma wind tunnel experiments using a Bayesian formulation*. Carbon, 200:199–214, 11 2022.

- [64] A. del Val, T. Magin, and P. Congedo. *Quantification of model-form uncertainties affecting the calibration of a carbon nitridation model by means of Bayesian Model Averaging*. International Journal of Heat and Mass Transfer, 213:124271, 2023.
- [65] K. Miki, M. Panesi, E. Prudencio, A. Maurente, S. Cheung, J. Jagodzinski, D. Goldstein, S. Prudhomme, K. Schulz, C. Simmons, J. Strand, and P. Varghese. *On the (In)Validation of a Thermochemical Model with EAST Shock Tube Radiation Measurements*. In 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition. American Institute of Aeronautics and Astronautics, 2010.
- [66] K. Miki, M. Panesi, E. Prudencio, and S. Prudhomme. *Estimation of the nitrogen ionization reaction rate using electric arc shock tube data and Bayesian model analysis*. Physics of Plasmas, 19(2):023507, 02 2012.
- [67] A. del Val and O. Chazot. *Stochastic determination of thermal reaction rate coefficients for air plasmas*. The Journal of Chemical Physics, 159(6):064105, 08 2023.
- [68] J. Scoggins, V. Leroy, G. Bellas-Chatzigeorgis, B. Dias, and T. Magin. *Mutation + + : MUlticomponent Thermodynamic And Transport properties for IONized gases in C++*. SoftwareX, 12:100575, 2020.
- [69] P. Gnoffo. *Planetary-Entry Gas Dynamics*. Annual Review of Fluid Mechanics, 31(1):459–494, 1999.
- [70] E. Josyula and J. Burt. *Review of Rarefied Gas Effects in Hypersonic Applications*. Technical Report 1ADA582766, AIR FORCE RESEARCH LAB WRIGHT-PATTERSON AFB OH, 2011.
- [71] G. Bird. *Molecular gas dynamics and the direct simulation of gas flows*. Number 42 in Oxford engineering science series. Clarendon Press ; Oxford University Press, 1994.
- [72] L. Mieussens. *Discrete-Velocity Models and Numerical Schemes for the Boltzmann-BGK Equation in Plane and Axisymmetric Geometries*. Journal of Computational Physics, 162(2):429–466, 2000.
- [73] C. Baranger, Y. Dauvois, G. Marois, J. Mathé, J. Mathiaud, and L. Mieussens. *A BGK model for high temperature rarefied gas flows*. European Journal of Mechanics - B/Fluids, 80:1–12, 2020.
- [74] R. Gupta, C. Scott, and J. Moss. *Slip-boundary equations for multicomponent nonequilibrium airflow*. Technical Report 19860005060, NASA, 1985.

- [75] M. Capitelli, I. Armenise, D. Bruno, M. Cacciatore, R. Celiberto, G. Colonna, O. De Pascale, P. Diomede, F. Esposito, C. Gorse, K. Hassouni, A. Laricchiuta, S. Longo, D. Pagano, D. Pietanza, and M. Rutigliano. *Non-equilibrium plasma kinetics: a state-to-state approach*. Plasma Sources Science and Technology, 16(1):S30–S44, 02 2007.
- [76] F. Milos and D. Rasky. *Review of numerical procedures for computational surface thermochemistry*. Journal of Thermophysics and Heat Transfer, 8(1):24–34, 1994.
- [77] P. Gnoffo, R. Gupta, and J. Shinn. *Conservation equations and physical models for hypersonic air flows in thermal and chemical nonequilibrium*. Technical Report 19890006744, NASA, 1989.
- [78] V. Giovangigli. *Multicomponent Flow Modeling*. Birkhäuser Boston, MA, 1999.
- [79] C. Park, R. Jaffe, and H. Partridge. *Chemical-Kinetic Parameters of Hyperbolic Earth Entry*. Journal of Thermophysics and Heat Transfer, 15(1):76–90, 2001.
- [80] D. Olynick, Y. Chen, and M. Tauber. *Aerothermodynamics of the Stardust Sample Return Capsul*. Journal of Spacecraft and Rockets, 36:442–462, 1998.
- [81] T. Gokcen. *N₂-CH₄-Ar Chemical Kinetic Model for Simulations of Atmospheric Entry to Titan*. Journal of Thermophysics and Heat Transfer, 21, 06 2004.
- [82] W. Vincenti and C. Kruger. *Introduction to physical gas dynamics*. Krieger Publishing Company, 1965.
- [83] M. Panesi. *Physical Models for Nonequilibrium Plasma Flow Simulations at High Speed Re-Entry Conditions*. PhD thesis, Università degli Studi di Pisa, 2009.
- [84] M. Chase. *NIST-JANAF Thermochemical Tables, 4th Edition*. American Institute of Physics, -1, 08 1998.
- [85] L. Gurvich, I. Veits, and C. Alcock. *Thermodynamics properties of individual substances. Volume 1 - Elements O, H/D, T/, F, Cl, Br, I, He, Ne, Ar, Kr, Xe, Rn, S, N, P, and their compounds. Part 1 - Methods and computation. Part 2 - Tables (4th revised and enlarged edition)*. 01 1989.
- [86] S. Gordon and J. McBride. *Thermodynamic Data to 20,000 K For Monatomic Gases*. Technical Report 1999-208523, NASA, 1999.

- [87] B. McBride, M. Zehe, and S. Gordon. *NASA Glenn coefficients for calculating thermodynamic properties of individual species*. Technical report, NASA, 10 2002.
- [88] H. Ferziger, J. Kaper and E. Gross. *Mathematical Theory of Transport Processes in Gases*. North-Holland Publishing Company, 1972.
- [89] T. Magin and G. Degrez. *Transport algorithms for partially ionized and unmagnetized plasmas*. *Journal of Computational Physics*, 198(2):424–449, 2004.
- [90] J. Scoggins. *Development of numerical methods and study of coupled flow, radiation, and ablation phenomena for atmospheric entry*. PhD thesis, Université Paris-Saclay, 2017.
- [91] J. Ramshaw and C. Chang. *Ambipolar diffusion in two-temperature multicomponent plasmas*. *Plasma Chemistry and Plasma Processing*, 13(3):489–498, 09 1993.
- [92] J. Ramshaw. *Self-Consistent Effective Binary Interaction Approximation For Strongly Coupled Multifluid Dynamics*. *Journal of Non-Equilibrium Thermodynamics*, 23(2), 1998.
- [93] T. Magin and G. Degrez. *Transport properties of partially ionized and unmagnetized plasmas*. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 70:046412, Oct 2004.
- [94] C. Wilke. *A Viscosity Equation for Gas Mixtures*. *The Journal of Chemical Physics*, 18(4):517–519, 04 1950.
- [95] B. Armaly and K. Sutton. *Viscosity of multicomponent partially ionized gas mixtures*. In 15th Thermophysics Conference. American Institute of Aeronautics and Astronautics, 07 1980.
- [96] R. Gupta, J. Yos, R. Thompson, and K. Lee. *A Review of Reaction Rates and Thermodynamic and Transport Properties for an 11-Species Air Model for Chemical and Thermal Nonequilibrium Calculations to 30000K*. Technical report, NASA, 1990.
- [97] P. Schrooyen, A. Turchi, K. Hillewaert, P. Chatelain, and T. Magin. *Two-way coupled simulations of stagnation-point ablation with transient material response*. *International Journal of Thermal Sciences*, 134:639–652, 2018.
- [98] B. Dias, A. Zibitsker, J. Meurisse, and N. Mansour. *Towards a flow-material unified solver for heatshield modeling*. In AIAA SCITECH 2023 Forum. American Institute of Aeronautics and Astronautics, 03 2023.

- [99] G. Bellas-Chatzigeorgis. *Development of advanced gas-surface interaction models for chemically reacting flows for re-entry conditions*. PhD thesis, Politecnico di Milano, 2018.
- [100] A. Martin, H. Zhang, and K. Tagavi. *An introduction to the derivation of surface balance equations without the excruciating pain*. *International Journal of Heat and Mass Transfer*, 115:992–999, 2017.
- [101] A. Viladegut, Ü. Düzel, and O. Chazot. *Diffusion effects on the determination of surface catalysis in Inductively Coupled Plasma facility*. *Chemical Physics*, 485-486:88–97, 2017.
- [102] D. Bianchi, F. Nasuti, and E. Martelli. *Coupled Analysis of Flow and Surface Ablation in Carbon-Carbon Rocket Nozzles*. *Journal of Spacecraft and Rockets*, 46(3):492–500, 2009.
- [103] R. Goulard. *On Catalytic Recombination Rates in Hypersonic Stagnation Heat Transfer*. *Journal of Jet Propulsion*, 28(11):737–745, 1958.
- [104] P. Barbante. *Accurate and Efficient Modelling of High Temperature Nonequilibrium Air Flows*. PhD thesis, Université Libre de Bruxelles, 2001.
- [105] M. Barbato, S. Reggiani, C. Bruno, and J. Muylaert. *Model for Heterogeneous Catalysis on Metal Surfaces with Applications to Hypersonic Flows*. *Journal of Thermophysics and Heat Transfer*, 14(3):412–420, 2000.
- [106] J. Marschall and M. MacLean. "Finite-Rate Surface Chemistry Model, I: Formulation and Reaction System Examples". 06 2011.
- [107] A. Lutz. *Experimental investigation and analysis of high-enthalpy nitrogen flow over graphite*. PhD thesis, University of Vermont, 2015.
- [108] V. Murray and T. Minton. *Gas-surface interactions of atomic nitrogen with vitreous carbon*. *Carbon*, 150:85–92, 2019.
- [109] G. Bellas-Chatzigeorgis, A. Turchi, A. Viladegut, O. Chazot, P. Barbante, and T. Magin. *Development of catalytic and ablative gas-surface interaction models for the simulation of reacting gas mixtures*. In 23rd AIAA Computational Fluid Dynamics Conference. American Institute of Aeronautics and Astronautics, 2017.
- [110] K. Hillewaert. *Development of the discontinuous Galerkin method for large scale/high-resolution CFD and acoustics in industrial geometries*. PhD thesis, Université catholique de Louvain, 2013.

- [111] J. Lachaud and N. Mansour. *Porous-Material Analysis Toolbox Based on OpenFOAM and Applications*. Journal of Thermophysics and Heat Transfer, 28(2):191–202, 2014.
- [112] G. May, K. Devesse, A. Rangarajan, and T. Magin. *A Hybridized Discontinuous Galerkin Solver for High-Speed Compressible Flow*. Aerospace, 8(11):322, 2021.
- [113] A. Lani, T. Quintino, D. Kimpe, H. Deconinck, S. Vandewalle, and S. Poedts. *The COOLFluid Framework: Design Solutions for High Performance Object Oriented Scientific Computing Software*. In V. Sunderam, G. van Albada, P. Sloot, and J. Dongarra, editors, Computational Science – ICCS 2005, pages 279–286, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [114] D. Ninni. *Development of a multi-GPU solver for atmospheric entry flows with gas-surface interactions*. PhD thesis, Politecnico di Bari, 2022.
- [115] S. Hickel, C. Egerer, and J. Larsson. *Subgrid-scale modeling for implicit large eddy simulation of compressible flows and shock-turbulence interaction*. Physics of Fluids, 26(10):106101, 10 2014.
- [116] W. Maier, J. Needels, J. Alonso, F. Morgado, C. Garbacz, M. Fossati, O. Tumuklu, and K. Hanquist. *Development of Physical and Numerical Nonequilibrium Modeling Capabilities within the SU2-NEMO Code*. In AIAA AVIATION 2023 Forum. American Institute of Aeronautics and Astronautics, 06 2023.
- [117] A. Margaritis, C. Scherding, O. Marxen, P. Schmid, and T. Sayadi. *Development of a high-fidelity computational tool for chemically reacting hypersonic flow simulations*, 2022.
- [118] C. Scherding, G. Rigas, D. Sipp, P. Schmid, and T. Sayadi. *Data-driven framework for input/output lookup tables reduction: Application to hypersonic flows in chemical nonequilibrium*. Phys. Rev. Fluids, 8:023201, 2023.
- [119] B. Dias. *Thermal ablation and radiation modeling of meteor phenomena*. PhD thesis, Université catholique de Louvain, 2020.
- [120] M. Capriati, G. Bellas-Chatzigeorgis, A. Turchi, B. Helber, and T. Magin. *Thermal non-equilibrium modeling for ablative gas-surface interaction*. In HiSST: 2nd International Conference on High-Speed Vehicle Science and Technology, 2022.
- [121] M. Capriati, K. Prata, T. Schwartzentruber, G. Candler, and T. Magin. *Development of a Nitridation Gas-Surface Boundary Condition for High-Fidelity*

- Hypersonic Simulations*. In 14th WCCM-ECCOMAS Congress. CIMNE, 2021.
- [122] A. Gordeev, A. Kolesnikov, and V. Sakharov. *Flow and heat transfer in underexpanded nonequilibrium jets of an induction plasmatron*. Fluid Dynamics - FLUID DYN, 46:623–633, 08 2011.
- [123] A. Balter-Peterson, F. Nichols, B. Mifsud, and W. Love. *Arc jet testing in NASA Ames Research Center thermophysics facilities*. In AIAA 4th International Aerospace Planes Conference. American Institute of Aeronautics and Astronautics, 12 1992.
- [124] Ü. Düzel, O. Schroeder, H. Zhang, and A. Martin. *Numerical Simulation of an Arc Jet Test Section*. Journal of Thermophysics and Heat Transfer, 34(2):393–403, 04 2020.
- [125] C. Purpura, F. De Filippis, P. Barrera, and D. Mandanici. *Experimental characterisation of the CIRA plasma wind tunnel SCIROCCO test section*. Acta Astronautica, 62(6):410–421, 2008.
- [126] V. Murray. *Gas-surface interactions with sp² carbon in extreme environments*. PhD thesis, Montana state university, 2018.
- [127] P. Barbante and O. Chazot. *Flight Extrapolation of Plasma Wind Tunnel Stagnation Region Flowfield*. Journal of Thermophysics and Heat Transfer, 20(3):493–499, 2006.
- [128] A. Turchi, J. Matesanz Saiz, T. Magin, and O. Chazot. *Duplication of hypersonic stagnation-region aerothermochemistry and gas-surface interaction in high-enthalpy ground testing*. Experiments in Fluids, 62, 2021.
- [129] O. Chazot. *Experimental Studies on Hypersonic Stagnation Point Chemical Environment*. Technical report, von Karman Institute for Fluid Dynamics, 2006.
- [130] D. Vanden Abeele. *An efficient computational model for inductively coupled air plasma flows under thermal and chemical non-equilibrium*. PhD thesis, KU Leuven, 2000.
- [131] T. Magin. *A model for inductive plasma wind tunnels*. PhD thesis, Université Libre de Bruxelles, 2004.
- [132] A. Fagnani, B. Helber, A. Hubin, and O. Chazot. *3D infrared temperature maps measurements of ablative materials during plasma wind tunnel experiments*. Measurement Science and Technology, 34(7):075401, 7 2023.

- [133] A. Fagnani, D. Le Quang Huy, B. Helber, S. Demange, A. Turchi, O. Chazot, and A. Hubin. *Investigation of a Free-Stream Air Plasma Flow by Optical Emission Spectroscopy and Comparison to Magnetohydrodynamics Simulations*. In AIAA Scitech 2020 Forum. American Institute of Aeronautics and Astronautics, 01 2020.
- [134] A. Cipullo, B. Helber, F. Panerai, L. Zeni, and O. Chazot. *Investigation of Freestream Plasma Flow Produced by Inductively Coupled Plasma Wind Tunnel*. *Journal of Thermophysics and Heat Transfer*, 28(3):381–393, 2014.
- [135] A. Fagnani, B. Dias, P. Schrooyen, B. Helber, T. Magin, and O. Chazot. *Investigation of Quartz Ablation in the VKI Plasmatron Facility: Comparison Between Experimental and Numerical Results*. In AIAA AVIATION 2021 FORUM. American Institute of Aeronautics and Astronautics, 08 2021.
- [136] A. Viladegut and O. Chazot. *Empirical Modeling of Copper Catalysis for Enthalpy Determination in Plasma Facilities*. *Journal of Thermophysics and Heat Transfer*, 34(1):26–36, 2020.
- [137] B. Helber, O. Chazot, A. Hubin, and T. Magin. *Microstructure and gas-surface interaction studies of a low-density carbon-bonded carbon fiber composite in atmospheric entry plasmas*. *Composites Part A: Applied Science and Manufacturing*, 72:96–107, 2015.
- [138] F. Panerai and O. Chazot. *Characterization of gas/surface interactions for ceramic matrix composites in high enthalpy, low pressure air flow*. *Materials Chemistry and Physics*, 134(2):597–607, 2012.
- [139] D. Guariglia, B. Helber, and O. Chazot. *Very High Heat-Flux Measurements in Plasmatron with Subsonic and Supersonic Plasma Flow*. In 8th European Symposium on Aerothermodynamics for Space Vehicles. European Space Agency, 2015.
- [140] B. Helber, A. Turchi, O. Chazot, T. Magin, and A. Hubin. *Gas/Surface Interaction Study of Low-Density Ablators in Sub- and Supersonic Plasmas*. In 11th AIAA/ASME Joint Thermophysics and Heat Transfer Conference. American Institute of Aeronautics and Astronautics, 06 2014.
- [141] D. Le Quang, Y. Babou, O. Chazot, and P. Andre. *Investigation of Supersonic Air Plasma Jet Produced in the VKI Plasmatron Facility*. ESA Special Publication, 714:31, 2012.
- [142] V. Murray, P. Recio, A. Caracciolo, C. Miossec, N. Balucani, P. Casavecchia, and T. Minton. *Oxidation and nitridation of vitreous carbon at high temperatures*. *Carbon*, 167:388–402, 10 2020.

- [143] V. Murray, B. Marshall, P. Woodburn, and T. Minton. *Inelastic and Reactive Scattering Dynamics of Hyperthermal O and O₂ on Hot Vitreous Carbon Surfaces*. The Journal of Physical Chemistry C, 119(26):14780–14796, 07 2015.
- [144] A. Kleyn. *Molecular beams and chemical dynamics at surfaces*. Chemical Society Reviews, 32:87–95, 2003.
- [145] L. Eça and M. Hoekstra. *A procedure for the estimation of the numerical uncertainty of CFD calculations based on grid refinement studies*. Journal of Computational Physics, 262:104–130, 2014.
- [146] J. Slotnick, A. Khodadoust, J. Alonso, D. Darmofal, W. Gropp, E. Lurie, and D. Mavriplis. *CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences*. 2014.
- [147] C. Roy and W. Oberkampf. *A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing*. Computer Methods in Applied Mechanics and Engineering, 200(25):2131–2144, 2011.
- [148] C. Roy and M. Balch. *A holistic approach to uncertainty quantification with application to supersonic nozzle thrust*. 2:363–381, 01 2012.
- [149] M. Kennedy and A. O’Hagan. *Bayesian Calibration of Computer Models*. Journal of the Royal Statistical Society Series B: Statistical Methodology, 63(3):425–464, 09 2001.
- [150] G. Matheron. *Principles of geostatistics*. Economic Geology, 58:1246–1266, 1963.
- [151] D. Xiu and G. Karniadakis. *“The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations”*. SIAM Journal on Scientific Computing, 24(2):619–644, 01 2002.
- [152] D. Lucor, D. Xiu, C.-H. Su, and G. E. Karniadakis. *Predictability and uncertainty in CFD*. International Journal for Numerical Methods in Fluids, 43(5):483–505, 2003.
- [153] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [154] T. Sullivan. *Introduction to Uncertainty Quantification*, volume 63 of *Texts in Applied Mathematics*. Springer International Publishing, 2015.
- [155] O. Le Maître and O. Knio. *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*. Scientific Computation. Springer Netherlands, 2010.

- [156] R. Caflisch. *Monte Carlo and quasi-Monte Carlo methods*. Acta Numerica, 7:1–49, 1998.
- [157] N. Metropolis and S. Ulam. *The Monte Carlo Method*. Journal of the American Statistical Association, 44(247):335–341, 09 1949.
- [158] M. McKay, R. Beckman, and W. Conover. *A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code*. Technometrics, 21(2):239–245, 1979.
- [159] I. Sobol'. *On the distribution of points in a cube and the approximate evaluation of integrals*. USSR Computational Mathematics and Mathematical Physics, 7(4):86–112, 1967.
- [160] K. Sutton and R. Graves. *A general stagnation-point convective heating equation for arbitrary gas mixtures*. Technical Report 19720003329, NASA, 1971.
- [161] J. Wakefield. *Bayesian and Frequentist Regression Methods*. Springer, 2013.
- [162] P. Wagner, J. Nagel, S. Marelli, and B. Sudret. *UQLab user manual – Bayesian inversion for model calibration and validation*. Technical report, Chair of Risk, Safety and Uncertainty Quantification, ETH Zurich, Switzerland, 2022. Report UQLab-V2.0-113.
- [163] C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer New York, 2004.
- [164] N. Metropolis, A. Rosenbluth, MN. Rosenbluth, A. Teller, and E. Teller. *“Equation of State Calculations by Fast Computing Machines”*. The Journal of Chemical Physics, 21(6):1087–1092, 06 1953.
- [165] W. Hastings. *“Monte Carlo sampling methods using Markov chains and their applications”*. Biometrika, 57(1):97–109, 04 1970.
- [166] A. Gelman, W. R. Gilks, and G. O. Roberts. *Weak convergence and optimal scaling of random walk Metropolis algorithms*. The Annals of Applied Probability, 7(1):110 – 120, 1997.
- [167] H. Haario, E. Saksman, and J. Tamminen. *An Adaptive Metropolis Algorithm*. Bernoulli, 7(2):223–242, 2001.
- [168] A. Gelman, G. O. Roberts, and W. R. Gilks. *Efficient Metropolis jumping rules*. In J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, editors, Bayesian Statistics, pages 599–608. Oxford University Press, Oxford, 1996.

- [169] J. Goodman and J. Weare. *Ensemble samplers with affine invariance*. Communications in Applied Mathematics and Computational Science, 5(1):65–80, 01 2010.
- [170] H. Rosenbrock. *An Automatic Method for Finding the Greatest or Least Value of a Function*. The Computer Journal, 3(3):175–184, 01 1960.
- [171] A. Gelman and D. Rubin. *Inference from Iterative Simulation Using Multiple Sequences*. Statistical Science, 7(4):457–472, 1992.
- [172] S. Brooks and A. Gelman. *General Methods for Monitoring Convergence of Iterative Simulations*. Journal of Computational and Graphical Statistics, 7(4):434–455, 1998.
- [173] A. del Val. *Bayesian calibration and assessment of gas-surface interaction models and experiments in atmospheric entry plasmas*. PhD thesis, École polytechnique, 2021.
- [174] D. Krige. *A statistical approach to some basic mine valuation problems on the Witwatersrand*. Journal of the Southern African Institute of Mining and Metallurgy, 52(6):119–139, 1951.
- [175] Z. Han, C. Xu, L. Zhang, Y. Zhang, K. Zhang, and W. Song. *Efficient aerodynamic shape optimization using variable-fidelity surrogate models and multilevel computational grids*. Chinese Journal of Aeronautics, 33(1):31–47, 01 2020.
- [176] I. Abdallah, C. Lataniotis, and B. Sudret. *Parametric hierarchical kriging for multi-fidelity aero-servo-elastic simulators — Application to extreme loads on wind turbines*. Probabilistic Engineering Mechanics, 55:67–77, 2019.
- [177] H. Babae, P. Perdikaris, C. Chrysostomidis, and G. E. Karniadakis. *Multi-fidelity modelling of mixed convection based on experimental correlations and numerical simulations*. Journal of Fluid Mechanics, 809:895–917, 12 2016.
- [178] A. Sadagopan, D. Huang, U. Duzel, L. Martin, and K. Hanquist. *”Assessment of High-Temperature Effects on Hypersonic Aerothermoelastic Analysis using Multi-Fidelity Multi-Variate Surrogates”*. In AIAA Scitech 2021 Forum. American Institute of Aeronautics and Astronautics, 01 2021.
- [179] M. Kennedy and A. O’Hagan. *Predicting the Output from a Complex Computer Code When Fast Approximations Are Available*. Biometrika, 87(1):1–13, 2000.

- [180] R. Howarth. *Mining Geostatistics*. Mineralogical Magazine, 43(328):563–564, 1979.
- [181] A. Forrester, A. Sobester, and A. Keane. *Multi-fidelity optimization via surrogate modelling*. Proc. R. Soc. A, 463:3251–3269, 12 2007.
- [182] Z. Han and S. Görtz. *Hierarchical Kriging Model for Variable-Fidelity Surrogate Modeling*. AIAA Journal, 50(9):1885–1896, 2012.
- [183] A. Forrester, A. Sobester, and A. Keane. *Engineering Design via Surrogate Modelling: A Practical Guide*. Wiley, 2008.
- [184] D. Jones, M. Schonlau, and W. Welch. *Efficient Global Optimization of Expensive Black-Box Functions*. Journal of Global Optimization, 13:455–492, 12 1998.
- [185] V. Picheny, T. Wagner, and D. Ginsbourger. *A benchmark of kriging-based infill criteria for noisy optimization*. Structural and Multidisciplinary Optimization, 48, 09 2013.
- [186] C. Rasmussen and C. Williams. *Gaussian processes for machine learning*. The MIT Press, 2006.
- [187] C. Lataniotis, D. Wicaksono, S. Marelli, and B. Sudret. *UQLab user manual – Kriging (Gaussian process modeling)*. Technical report, Chair of Risk, Safety and Uncertainty Quantification, ETH Zurich, Switzerland, 2021. Report UQLab-V1.4-105.
- [188] Y. Zhang, Z. Han, and k. Zhang. *Variable-fidelity expected improvement method for efficient global optimization of expensive functions*. Structural and Multidisciplinary Optimization, 58, 10 2018.
- [189] S. Marelli and B. Sudret. *UQLab: A Framework for Uncertainty Quantification in Matlab*. In *Vulnerability, Uncertainty, and Risk*, pages 2554–2563. American Society of Civil Engineers, 2014.
- [190] G. Candler, P. Subbareddy, and I. Nompelis. *CFD Methods for Hypersonic Flows and Aerothermodynamics*, pages 203–237. 02 2015.
- [191] G. Candler, D. Mavriplis, and L. Trevino. *Current Status and Future Prospects for the Numerical Simulation of Hypersonic Flows*. In *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*. American Institute of Aeronautics and Astronautics, 2009.
- [192] G. Candler and I. Nompelis. *Computational Fluid Dynamics for Atmospheric Entry*. Technical Report ADA568031, NATO, 2009.

- [193] P. Roache. *Perspective: A Method for Uniform Reporting of Grid Refinement Studies*. Journal of Fluids Engineering, 116(3):405–413, 1994.
- [194] A. Munafò and T. Magin. *Modeling of stagnation-line nonequilibrium flows by means of quantum based collisional models*. Physics of Fluids, 26(9):097102, 2014.
- [195] J. Steger and R. Warming. *Flux vector splitting of the inviscid gasdynamic equations with application to finite-difference methods*. Journal of Computational Physics, 40(2):263–293, 1981.
- [196] B. van Leer. *Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov’s method*. Journal of Computational Physics, 32(1):101–136, 1979.
- [197] M. Wright, G. Candler, and D. Bose. *Data-Parallel Line Relaxation Method for the Navier-Stokes Equations*. AIAA Journal, 36(9):1603–1609, 1998.
- [198] P. Subbareddy and G. Candler. *A fully discrete, kinetic energy consistent finite-volume scheme for compressible flows*. Journal of Computational Physics, 228(5):1347–1364, 2009.
- [199] S. Kim. *A Multi-Dimensional Linear Reconstruction Scheme for Arbitrary Unstructured Mesh*. In 16th AIAA Computational Fluid Dynamics Conference. American Institute of Aeronautics and Astronautics, 06 2003.
- [200] A. Klomfass and S. Müller. *Calculation of stagnation streamline quantities in hypersonic blunt body flows*. Shock Waves, 7(1):13–23, 02 1997.
- [201] A. Munafo. *Multi-Scale models and computational methods for aerothermodynamics*. PhD thesis, Ecole Centrale Paris, 05 2014.
- [202] P. Roe. *Approximate Riemann solvers, parameter vectors, and difference schemes*. Journal of Computational Physics, 43(2):357–372, 1981.
- [203] M. Liou. *A sequel to AUSM, Part II: AUSM+-up for all speeds*. Journal of Computational Physics, 214(1):137–170, 2006.
- [204] L. Richardson and R. Glazebrook. *IX. The approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stresses in a masonry dam*. Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character, 210(459):307–357, 1911.
- [205] D. Knight, J. Longo, D. Drikakis, D. Gaitonde, A. Lani, I. Nompelis, B. Reimann, and L. Walpot. *Assessment of CFD capability for prediction of*

- hypersonic shock interactions*. Progress in Aerospace Sciences, 48-49:8–26, 2012.
- [206] P. Papadopoulos, E. Venkatapathy, D. Prabhu, M. Loomis, and D. Olynick. *Current grid-generation strategies and future requirements in hypersonic vehicle design, analysis and testing*. Applied Mathematical Modelling, 23(9):705–735, 1999.
- [207] G. Candler, M. Barnhardt, T. Drayna, I. Nompelis, D. Peterson, and P. Subbareddy. *Unstructured Grid Approaches for Accurate Aeroheating Simulations*. In 18th AIAA Computational Fluid Dynamics Conference. American Institute of Aeronautics and Astronautics, 2007.
- [208] S. Henderson and J. Menart. *Grid Study on Blunt Bodies with the Carbuncle Phenomenon*. In 39th AIAA Thermophysics Conference. American Institute of Aeronautics and Astronautics, 2007.
- [209] A. Bonfiglioli, M. Grottadaurea, R. Paciorri, and F. Sabetta. *An unstructured, three-dimensional, shock-fitting solver for hypersonic flows*. Computers & Fluids, 73(Supplement C):162 – 174, 2013.
- [210] M. Onofri, R. Paciorri, D. Cardillo, M. Grottadaurea, and A. Bonfiglioli. *Numerical Simulations of flows past IXV re-entry vehicle at CRAS*. In 3rd International ARA days, 2011.
- [211] Z. Gao, H. Xue, Z. Zhang, H. Liu, and C. Lee. *A hybrid numerical scheme for aeroheating computation of hypersonic reentry vehicles*. International Journal of Heat and Mass Transfer, 116:432–444, 2018.
- [212] S. Luo, J. Liu, and K. Li. *Grid convergence and influence of wall temperature in the calculation of thermochemical non-equilibrium heat flux*. Journal of Physics D: Applied Physics, 53(28):285502, 2020.
- [213] D. Saunders, S. Yoon, and M. Wright. *An Approach to Shock Envelope Grid Tailoring and Its Effect on Reentry Vehicle Solutions*. In 45th AIAA Aerospace Sciences Meeting and Exhibit. American Institute of Aeronautics and Astronautics, 2007.
- [214] S. Kawai and K. Shimoyama. *Kriging-model-based uncertainty quantification in computational fluid dynamics*. In 32nd AIAA Applied Aerodynamics Conference. American Institute of Aeronautics and Astronautics, 2014.
- [215] S. Salehi, M. Raisee, M. Cervantes, and A. Nourbakhsh. *Efficient uncertainty quantification of stochastic CFD problems using sparse polynomial chaos and compressed sensing*. Computers & Fluids, 154:296–321, 2017.

- [216] A. Barrio, M. Sudars, J. Gavira, R. Aulizio, F. Ratti, F. Massobrio, L. Walpot, G. Passarelli, J. Thoemel, and A. Thirkettle. *EXPERT - The ESA experimental re-entry test-bed. Trajectory and mission design*. 08 2011.
- [217] L. Walpot, H. Ottens, J. Muylaert, O. Bayle, P. Urmston, U. Thomas, G. Sacoccia, M. Caporicci, and C. Stavrinidis. *Heat Flux and Static Stability Predictions of the Expert Vehicle*. In D. Danesy, editor, Fifth European Symposium on Aerothermodynamics for Space Vehicles, volume 563 of *ESA Special Publication*, page 139, 01 2005.
- [218] C. Roy. *Grid Convergence Error Analysis for Mixed-Order Numerical Schemes*. *AIAA Journal*, 41(4):595–604, 2003.
- [219] B. Helber, O. Chazot, T. Magin, and A. Hubin. *Space and Time-Resolved Emission Spectroscopy of Carbon Phenolic Ablation in Air and Nitrogen Plasmas*. In 44th AIAA Thermophysics Conference. American Institute of Aeronautics and Astronautics, 06 2013.
- [220] P. Ventura Diaz, A. Parente, J. Meurisse, S. Yoon, and N. Mansour. *High-Fidelity Numerical Analysis of Arc-Jet Aerothermal Environments*. In AIAA Scitech 2020 Forum. American Institute of Aeronautics and Astronautics, 01 2020.
- [221] M. Wright, G. Candler, and M. Prampolini. *Data-parallel lower-upper relaxation method for the Navier-Stokes equations*. *AIAA Journal*, 34(7):1371–1377, 1996.
- [222] C. Park and D. Bogdanoff. *Shock-Tube Measurement of Nitridation Coefficient of Solid Carbon*. *Journal of Thermophysics and Heat Transfer*, 20(3):487–492, 2006.
- [223] G. Vignoles, A. Turchi, D. Bianchi, P. Blaineau, X. Lamboley, D. Le Quang Huy, C. Levet, O. Caty, and O. Chazot. *Ablative and catalytic behavior of carbon-based porous thermal protection materials in nitrogen plasmas*. *Carbon*, 134:376–390, 2018.
- [224] T. Suzuki, K. Fujita, and T. Sakai. *Graphite Nitridation in Lower Surface Temperature Regime*. *Journal of Thermophysics and Heat Transfer*, 24(1):212–215, 2010.
- [225] T. Suzuki, K. Fujita, K. Ando, and T. Sakai. *Experimental Study of Graphite Ablation in Nitrogen Flow*. *Journal of Thermophysics and Heat Transfer*, 22(3):382–389, 2008.

-
- [226] R. Nieman, M. Sands, Y. Wang, T. Minton, E. Mussoni, J. Engerer, and H. Guo. *Informing air-carbon ablation modeling with theoretical calculations of atomic oxygen and nitrogen interacting with carbon surfaces*. Physical Chemistry Chemical Physics, 2023.
- [227] M. Geratz, M. Capriati, G. Grossir, and T. Magin. *Influence of Physical Models on the Numerical Modeling of Hypersonic Nozzle Flow Expansion*. In 9th European conference for aeronautics and space sciences. EUCASS association, 2022.
- [228] T. Magin., M. Capriati, G. Kale, G. Bellas-Chatzigeorgis, A. Turchi, B. Helber, and O. Chazot. *PLASMUT - Nonequilibrium Gas-Surface Interactions at High Temperature, VKI Plasmatron and Mutation++ Library*. Technical report, Air Force Research Laboratory, 2021.
- [229] N. Leoni. *Bayesian inference of model error for the calibration of two-phase CFD codes*. PhD thesis, École polytechnique, 2022.
- [230] F. Bariselli, E. Torres, and T. Magin. *State-specific catalytic recombination boundary condition for DSMC methods in aerospace applications*. AIP Conference Proceedings, 1786(1):190009, 2016.

Titre: Inférence bayésienne multi-fidélité des conditions d'écoulement hypersonique et des paramètres de modèles de chimie hétérogène

Mots clés: Écoulements hypersoniques, Entrée atmosphériques, Quantification de l'incertitude, Reconstruction des conditions de l'écoulement amont

Résumé: La caractérisation de l'interaction entre un vaisseau spatial entrant dans une atmosphère planétaire et le gaz environnant est une tâche difficile, nécessitant des expériences précises et des simulations haute-fidélité. Les prédictions numériques dépendent fortement des incertitudes de modélisation et expérimentales qui s'accumulent pendant le processus d'inférence. Les méthodes de quantification de l'incertitude offrent une structure puissante pour tenir compte de plusieurs sources d'incertitude. Les prédictions haute-fidélité pour les applications aérospatiales combinent des modèles physiques suffisamment complets pour prendre en compte des caractéristiques d'écoulement complexes, et une quantification rigoureuse de l'incertitude. En raison du grand nombre de simulations requises pour mener des études de quantification de l'incertitude, des représentations efficaces à basse fidélité sont attrayantes. Cependant, les simplifications contenues dans ces modèles peuvent conduire à une précision réduite, détériorant potentiellement le résultat d'un problème d'inférence. Dans cette thèse, notre objectif est de développer des outils et des méthodologies pour effectuer des prédictions précises à l'aide de solveurs haute-fidélité et de données expérimentales de pointe. Cette approche implique l'utilisation du solveur CFD US3D et d'une structure globale UQ pour résoudre des problèmes d'inférence, ce qui permet d'inclure l'erreur de maillage et d'utiliser des stratégies multi-fidélité. La première contribution de cette thèse concerne la production de solutions haute-fidélité avec le solveur US3D pour chaque phénomène d'intérêt. Ce dernier a été couplé à Mutation++, que nous avons étendue pour incorporer un modèle innovant d'ablation. La deuxième contribution concerne une étude sur l'influence de l'erreur de maillage sur la convergence des simulations haute-fidélité sous incertitude. Nous avons construit un modèle de substitution efficace en équilibrant les erreurs

numériques associées au maillage et les incertitudes liées au problème. Nous avons appliqué cette méthodologie à la propagation des incertitudes du modèle pour caractériser la pression et le flux de chaleur subis par un véhicule de rentrée. Des résultats précis ont été obtenus avec un maillage grossier automatiquement aligné avec le choc pour chaque point d'entraînement. La troisième contribution concerne le développement d'une formulation multi-fidélité pour alléger le coût de calcul associé à la construction du modèle de substitution du solveur haute-fidélité et à son utilisation dans un problème d'inférence. Nous avons défini une méthodologie pour caractériser un jet hypersonique obtenu dans le moyen d'essai Plasmatron, pour laquelle aucune procédure standard de reconstruction des conditions en amont n'existait à ce jour. L'analyse nous a permis de caractériser les conditions d'écoulement à l'entrée de la tuyère et le coefficient de recombinaison catalytique de l'azote de la sonde utilisée pour mesurer le flux de chaleur et la pression. Les incertitudes caractérisées ont ensuite été propagées à travers le solveur numérique, fournissant une représentation haute-fidélité basée sur l'incertitude de la variabilité de la structure de l'écoulement supersonique. Dans la dernière application, nous avons élaboré une méthodologie pour l'étalonnage et l'évaluation d'un modèle chimique d'interaction gaz-surface à taux de réaction finis pour l'ablation. Plus précisément, nous avons déduit les taux de coefficients des réactions élémentaires se produisant entre une surface de carbone et un gaz d'azote à partir d'expériences de faisceau moléculaire-surface de diffusion, ainsi que Plasmatron. L'analyse a montré que les deux ensembles de données expérimentales sont compatibles avec le même modèle étalonné. En conclusion, nous avons proposé des outils stochastiques puissants pour déduire des conditions d'écoulement libre hypersonique et de paramètres de modèles chimiques hétérogènes.

Title: Multi-fidelity Bayesian inference of hypersonic flow free-stream conditions and heterogeneous chemistry model parameters

Keyword: Hypersonic flows, Atmospheric entry, Uncertainty Quantification, Freestream rebuilding

Abstract: The characterization of the interaction between a spacecraft entering a planetary atmosphere at hypersonic speeds and its surrounding gas is a challenging task, requiring accurate experiments and high-fidelity simulations. Numerical predictions strongly depend on modeling and experimental uncertainties, as well as numerical errors, which accumulate during the inference process. In this context, uncertainty quantification methods offer a powerful framework to account for several sources of uncertainty. High-fidelity predictions for aerospace applications combine: i) physical models sufficiently complete to account for complex flow features, leading to expensive and hard-to-perform numerical simulations, ii) rigorous uncertainty quantification permitting the use of experimental data to improve the prediction of quantities of interest. Because of the large number of simulations requested to conduct uncertainty quantification studies, efficient low-fidelity representations are appealing to reduce the computational effort. However, the simplifications contained in the low-fidelity models can lead to reduced accuracy, potentially deteriorating the outcome of an inference problem. In this thesis, our objective is to develop tools and methodologies to perform accurate predictions using high-fidelity solvers and state-of-the-art experimental data. This approach involves the use of the US3D CFD solver and an overall UQ framework to solve inference problems, which permits to include mesh error and to employ multi-fidelity strategies. The first contribution of this thesis concerns the production of high-fidelity solutions for each phenomenon of interest with the US3D solver. We coupled US3D to the open-source Mutation++ physicochemical library, which we expanded to incorporate a state-of-the-art ablation model. The second contribution concerns a study about the influence of the mesh error on the convergence of high-fidelity simulations under uncertainty. We constructed an efficient surrogate model by balancing the grid's numerical errors and the problem-related uncertainties. We applied this methodology to the propagation of model uncertainties to characterize the pressure and heat flux experienced by a re-entry vehicle.

Accurate results were obtained with a coarse mesh automatically aligned to the shock for each training point. The third contribution concerns the development of a multi-fidelity formulation to alleviate the computational cost associated with the construction of the surrogate model for the high-fidelity solver and its use in an inference problem. In particular, we defined a methodology to characterize an under-expanded high-enthalpy jet obtained in the von Karman Institute Plasmatron facility, for which no standardized rebuilding procedure for the free-stream conditions existed to date. The description of such a flow required axisymmetric simulations. The analysis allowed us to characterize the flow conditions at the entrance of the nozzle and the nitrogen catalytic recombination coefficient of the probe used to measure the heat flux and pressure at the stagnation point. The characterized uncertainties were then propagated through the numerical solver yielding an uncertainty-based high-fidelity representation of the supersonic flow structure variability. In the last application, we devised a methodology for the calibration and assessment of a finite-rate chemistry gas-surface interaction model for ablation. Specifically, we inferred the rates coefficients of the elementary reactions occurring between a carbon surface and a nitrogen gas from both molecular beam-surface scattering and Plasmatron experiments. The analysis highlighted that both experimental data sets are compatible with the same calibrated model. In conclusion, we proposed powerful stochastic tools, encompassing one or more fidelity levels, to infer hypersonic flow free-stream conditions and heterogeneous chemical model parameters.