



**von KARMAN INSTITUTE
FOR FLUID DYNAMICS**

Uncertainty Quantification of Aerothermal Flow-Material Simulations of Low-Density Ablative Thermal Protection Systems

Thesis submitted by Joffrey Coheur

in partial fulfillment of the requirements for the degree of Doctor of Philosophy
(PhD) in Engineering Sciences (ULiège & UCLouvain)

December 2021

Supervisors: Professor Maarten Arnst (Université de Liège)
Aerospace & Mechanical Engineering
Computational and Stochastic Modeling
and Professor Philippe Chatelain (Université catholique de Louvain)
Institute of Mechanics, Materials and Civil Engineering
Thermodynamics and Fluid Mechanics
and Professor Thierry Magin (von Karman Institute for Fluid Dynamics)
Aeronautics and Aerospace

Thesis jury :

Maarten Arnst (Université de Liège)
Philippe Chatelain (Université catholique de Louvain)
Pietro Congedo (Inria Saclay Île-de-France)
Jean Lachaud (Université de Bordeaux)
Thierry Magin (von Karman Institute for Fluid Dynamics)
Miltiadis Papalexandris (Université catholique de Louvain)
Vincent Terrapon (Université de Liège)

Thesis jury:

The present dissertation has been evaluated by the members of the Jury (sorted by alphabetical order):

Prof. Maarten Arnst	(Supervisor)	Université de Liège;
Prof. Philippe Chatelain	(Supervisor)	Université catholique de Louvain;
Prof. Pietro Congedo		Inria Saclay Île-de-France;
Prof. Jean Lachaud		Université de Bordeaux;
Prof. Thierry Magin	(Supervisor)	von Karman Institute for Fluid Dynamics;
Prof. Miltiadis Papalexandris	(Secretary)	Université catholique de Louvain;
Prof. Vincent Terrapon	(Chair)	Université de Liège.

Funding:

This research was financially supported by the Fund for Research Training in Industry and Agriculture (FRIA) 1E05418F provided by the Belgian Fund for Scientific Research (F.R.S.-FNRS). The scientific stay at NASA Ames was supported by the Fédération Wallonie-Bruxelles (FWB) and by the Wallonie-Bruxelles International (WBI) fellowship for excellence (WBI.World).

Computational resources have been provided by the Consortium des équipements de Calcul Intensif (CECI), funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S-FNRS) under grant no. 2.5020.11.

Uncertainty Quantification of Aerothermal Flow-Material Simulations of Low-Density Ablative Thermal Protection Systems

Coheur Joffrey

Abstract

Essential to space missions involving an atmospheric entry, the thermal protection system (TPS) shields the spacecraft and its payload from the severe aerothermal loads. Low-density carbon/phenolic composite materials have gained renewed interest to serve as ablative thermal protection materials (TPMs). These materials can accommodate the high heating rates and heat loads encountered during the atmospheric entry, at hypersonic velocities, by absorbing part of the incoming heat through physico-chemical transformations. One of the main endothermic processes is the pyrolysis of the resin compound, whereby volatile products are released, leaving a carbonaceous residue on the fibers.

Whereas new experimental data have already been published to characterize the decomposition of these low-density carbon/phenolic materials, they are yet to be exploited for the inversion of physico-chemical models. In addition, the issue of uncertainty quantification, required to assess the reliability of the numerical model and the physico-chemical models, is yet to be addressed. Therefore, the overarching objective of this thesis is to contribute to the development of an uncertainty-quantified numerical modeling of the ablation of new porous composite materials and to the analysis of the impact of uncertainty on TPS design. To that aim, we first address the development and the uncertainty characterization of physico-chemical models for resin pyrolysis on the basis of new experimental data relevant to the pyrolytic decomposition of the phenolic resin used in carbon/phenolic composite TPMs. Then, we analyze the impact of the uncertainty in the physico-chemical models on the numerical modeling of ablation of TPS by means of non-intrusive stochastic methods.

The central contribution of this thesis is to infer from these new experimental data an uncertainty-quantified pyrolysis model. We adopt a Bayesian probabilistic approach to account for uncertainties in the model identification. We use an approximate likelihood function involving a weighted distance between the model predictions and the time-dependent experimental data. To sample from the posterior, we use gradient-informed Markov chain Monte Carlo methods with an adaptive selection of the numerical parameters. We develop a versatile code for performing uncertainty characterization using Bayesian inference tools on engineering problems in which the proposed methods are implemented. To select the decomposition mechanisms to be represented in the pyrolysis model, we proceed by progressively increasing the complexity of the pyrolysis model until a satisfactory fit to the data is ultimately obtained. To improve the computational time, we derive a fast semi-analytical solution for the resin pyrolysis using multi-component parallel reactions both for the case of a constant temperature and the case of a linear heating rate. The pyrolysis model thus obtained involves six reactions and has 48 parameters.

A second contribution is the assessment of the impact of uncertainties on the material response of an ablating TPS relevant to in-flight performance prediction. Using the six-reaction pyrolysis model, we demonstrate its use in a numerical simulation of heat shield surface recession in a Martian entry. We provide probabilistic projections of the recession of the surface, the production of gaseous species at the surface, and the temperature.

In addition to the aforementioned contributions, we also provide three supplementary pieces of research. For the first one, we contribute to the development of a model for representing

the process of ablation from resin pyrolysis to char ablation in a unified flow-material approach where the Volume-Averaged Navier-Stokes equations are solved. This model is implemented in the high-fidelity numerical code Argo from the research center Cenaero and verified by a code-to-code comparison. The second one pertains to the development of an uncertainty-quantified pyrolysis model including competitive mechanisms for the pyrolytic decomposition of the PICA material. Finally, the third one concerns the calibration of material properties and environmental conditions in a Bayesian inference framework using post-flight data of the Mars Science Laboratory mission.

Keywords Atmospheric entry, thermal protection system, carbon/phenolic composite, Bayesian inference, chemical kinetics, Markov chain Monte Carlo

Quantification des incertitudes pour la simulation aérothermique fluide-matériau des systèmes de protection thermique ablatifs légers

Coheur Joffrey

Résumé

Essentiel aux missions spatiales impliquant une entrée dans l'atmosphère, le système de protection thermique (SPT) protège le véhicule spatial et sa charge utile contre les contraintes aérothermiques importantes. Les matériaux composites de type carbone/phénol à faible densité ont connu un regain d'intérêt dans leur utilisation comme matériaux de protection thermique (MPT) ablatifs. Ces matériaux peuvent supporter les taux de chauffe élevés et les contraintes thermiques rencontrées lors de l'entrée dans l'atmosphère, à des vitesses hypersoniques, en absorbant une partie de la chaleur grâce à des transformations physico-chimiques. L'un des principaux processus endothermiques est la pyrolyse du composé résineux, qui libère des produits volatiles et laisse un résidu carboné sur les fibres.

Alors que de nouvelles données expérimentales ont déjà été publiées pour caractériser la décomposition des matériaux carbone/phénol légers, ces données restent encore à être exploitées pour l'inversion des modèles physico-chimiques. En outre, la question de la quantification des incertitudes, nécessaire pour évaluer la fiabilité du modèle numérique et des modèles physico-chimiques, n'a pas encore été abordée. Par conséquent, l'objectif global de cette thèse est de contribuer au développement d'une modélisation numérique avec quantification des incertitudes sur l'ablation de ces matériaux composites poreux et à l'analyse de l'impact de l'incertitude sur la conception des SPT. Dans ce but, nous abordons dans un premier temps le développement et la caractérisation de l'incertitude de modèles physico-chimiques pour la pyrolyse de la résine sur base des nouvelles données expérimentales relatives à la décomposition pyrolytique de la résine phénolique. Ensuite, nous analysons l'impact de l'incertitude des paramètres des modèles physico-chimiques sur la modélisation numérique de l'ablation des MTP au moyen de méthodes stochastiques non intrusives.

La contribution centrale de cette thèse est d'inférer à partir de ces nouvelles données expérimentales un modèle de pyrolyse avec quantification des incertitudes. Nous adoptons une approche probabiliste bayésienne pour tenir compte des incertitudes dans l'identification du modèle et de ses paramètres. Nous utilisons une fonction de vraisemblance approximative impliquant une distance pondérée entre les prédictions du modèle et les données expérimentales dépendant du temps. Pour échantillonner le postérieur, nous utilisons des méthodes de Monte-Carlo par chaînes de Markov informées par gradient avec une sélection adaptative des paramètres numériques. Nous développons un code polyvalent pour effectuer la caractérisation de l'incertitude en utilisant des outils d'inférence bayésienne sur des problèmes d'ingénierie dans lesquels les méthodes proposées sont mises en œuvre. Afin de sélectionner les mécanismes de décomposition du modèle de pyrolyse, nous procédons par augmentation progressive de la complexité du modèle jusqu'à ce qu'un ajustement satisfaisant aux données soit obtenu. Afin d'améliorer le temps de calcul, nous dérivons une solution semi-analytique au modèle de pyrolyse faisant intervenir des réactions parallèles de composants multiples, à la fois pour le cas d'une décomposition à température constante et le cas d'un taux de chauffe linéaire. Le modèle de pyrolyse ainsi obtenu implique six réactions et comporte 48 paramètres.

Une deuxième contribution est l'évaluation de l'impact des incertitudes sur la réponse du matériau d'un SPT ablatif, pertinente pour la prédiction des performances en vol. En utilisant

le modèle de pyrolyse à six réactions, nous démontrons son utilisation dans une simulation numérique de la récession de la surface du bouclier thermique lors d’une entrée dans l’atmosphère de Mars. Nous fournissons des projections probabilistes de la récession de la surface, de la production d’espèces gazeuses à la surface et de la température.

En plus des contributions susmentionnées, nous apportons également trois travaux de recherche supplémentaires. Pour le premier, nous contribuons au développement d’un modèle représentant le processus d’ablation, à partir de la pyrolyse de la résine jusqu’à décomposition de la matière carbonisée, suivant une approche unifiée fluide-matériau dans laquelle les équations de Navier–Stokes moyennées en volume sont résolues. Ce modèle, implémenté dans le code numérique haute-fidélité Argo du centre de recherche Cenaero, est vérifié au moyen d’une comparaison des résultats numérique avec ceux d’un autre code. Le second concerne le développement d’un modèle de pyrolyse avec quantification des incertitudes, incluant des mécanismes compétitifs pour la modélisation de la décomposition pyrolytique du matériau PICA. Enfin, le troisième concerne la calibration des propriétés du matériau et des conditions environnementales dans un cadre d’inférence bayésienne en utilisant les données post-vol de la mission Mars Science Laboratory.

Mots-clés Entrée atmosphérique, système de protection thermique, matériau composite carbone/phenol, inférence bayésienne, cinétique chimique, Monte-Carlo par chaînes de Markov

Acknowledgements

This work is the result of a fruitful collaboration between several institutions, namely the University of Liège, the University of Louvain, the von Karman Institute for Fluid Dynamics (VKI), Cenaero, and NASA Ames Research Center. It would be difficult to thank in a few lines all the people who crossed my path and who helped me in one way or another; I hope these people will recognize themselves, and to all of them, I would like to address my most sincere gratitude.

First and foremost, I would like to greatly acknowledge my three thesis supervisors. It is thanks to your confidence and your support during all these years that I could reach the end of this adventure. Having three supervisors is quite unusual, and more than once it was difficult to get us all together. But I have learned a lot from each of you, and each of you brought different points of view and distinct ways of working, which gives the additional richness to this work. Thierry, you introduced me to the fantastic world of space engineering, integrated me in your dynamic and enthusiastic team at the VKI, and allowed me to develop many collaborations. Philippe, you always motivated me to go forward, even when nothing was working, and to look at things in a positive way. Your relevant remarks always helped me to improve the quality of my research. Maarten, the time you spent on supervising this work is uncountable and I deeply thank you for all your rigorous feedback, your patience and your persistence.

I would like to acknowledge the members of the Jury, Pietro Congedo, Jean Lachaud, Miltiadis Papalexandris, and Vincent Terrapon, for their interest in this work, for the time they spent reviewing this manuscript, and for the constructive scientific exchanges we had. I would also like to thank Gaetan Kerschen for having been in my PhD thesis committee during these five years. I am also grateful to Olivier van Cutsem, Jeremie Meurisse and Pierre Schrooyen, for their review of parts of this manuscript, which helped me to increase its quality and readability. I would particularly like to mention the contribution of Kevin Bulthuis for the thorough reading of the manuscript; your help has been truly valuable and highly motivating.

I would like to thank my direct colleagues at the University. Of course, my office mates during four years, Kim Liegeois and Kevin Bulthuis, with whom it has always been a pleasure and helpful to discuss research. And more recently, my two new office mates, Thomas and Martin, who helped me to feel less alone during these long hours of writing. I would like to thank my colleagues from the A&M department (Arnaud, Adrien, Dominik, Paola, Juju), the Automotive Engineering Research group (Pierre, Maxime, Simon, Eduardo, Pablo, Denis), the technicians (Antonio, Antoine, Mathieu), and the other members of the Shell-Eco Marathon team (Asmaa, Laura, Pilou).

The VKI is a big family, and it would be difficult to thank all the people who helped me in one way or another, and my thanks go to all the members of this great team. I would like to deeply thank Alessandro Turchi, who has always been available for me and for your timely advice; Francisco Torres-Herrador for the many discussions, advice, new fancy tools and for including me in your collaborations; Anabel del Val, the only VKI person who could understand my UQ-related problems; Domenico Fiorini and Andrea Fagnani for collaborating on new topics with me.

During the first years of my thesis, I worked intermittently at Cenaero. I am grateful to Koen Hillewaert who welcomed me in his research team and for his guidance. I would like to express special thanks to Pierre Schrooyen, who always followed me and provided me with sound advice,

and always with benevolence. The time you dedicated to me was very meaningful to me and helped me tremendously during my early days as a researcher.

I would like to gratefully acknowledge my colleagues from NASA Ames Research Center. I am indebted to Dr. Nagi Mansour for providing me this amazing opportunity to carry out my work at NASA Ames and for his supervision. Your experience and knowledge have been a real inspiration for me. I would also like to acknowledge Jeremie Meurisse for the support and the numerous discussions and advice, and John Thornton for the help on the numerical simulations and data analysis. I am grateful to Steven Yoon, the NASA Advances Supercomputing Division branch chief, who welcomed me during 6 months. It was an extraordinary opportunity to be able to carry out my research in such an extraordinary environment with a rich history. Of course, this scientific stay would not have been the same without the people I met there. I would like to thank Francesco Panerai and Laura Villafañe for the warm welcome and for making my first visits to the US an amazing journey. I would like to mention my colleagues and weekend companions, Francisco, Andrea, Thanos, David, with a special dedication to this p.o.s. Quentin Wagnier, and I am glad for the long-lasting friendship that have grown since then. I am also grateful to the Rainbow Mansion community, where I lived during 6 months, and the amazing people I met there.

I would like to thank the master thesis students and students from the VKI short training program who have made a short, but important contribution to my work and allowed me to develop new management skills: Anthony Bosco, for your initial work on the Hamiltonian Monte Carlo applied to pyrolysis equations; Martin Lacroix, for your incredible implementation work within Pybitup; Jennifer Dania and Thomas Chavet, for giving me the opportunity to expand my field of applications.

Finally, I would like to thank my family and friends who has always been there to support me. This manuscript is dedicated to my loved ones who left too early to see the completion of this work, but who have always been proud of me.

Contents

1	Introduction	3
1.1	Research context: atmospheric entry and thermal protection systems for spacecraft	3
1.2	Atmospheric entry flows and thermal protection material decomposition	4
1.2.1	Thermal protection materials	5
1.2.2	Importance of the coupling between the flow and material ablation	8
1.3	Uncertainty quantification for TPS modeling	9
1.3.1	Assessing uncertainty margins on the predictions	10
1.3.2	Enhancing the reliability of the models	11
1.4	Towards a robust characterization of uncertainty in aerothermal flow-material simulations	11
1.5	Objectives and contributions of the thesis	12
1.6	Overview of the manuscript	13
I	Modeling of thermal protection systems for atmospheric entry	17
2	Physics of atmospheric entry flows	19
2.1	Hypotheses for the modeling of aerothermal flows and ablative porous material response	19
2.1.1	Continuum versus rarefied flow regimes	20
2.1.2	Flow-material interactions	21
2.1.3	Numerical strategies for flow-material simulations	23
2.2	Fundamentals of aerothermal flows	24
2.2.1	Thermodynamic properties	25
2.2.2	Homogeneous chemical reactions	26
2.2.3	Navier–Stokes equations for multi-component reactive flows	27
2.3	Multi-component reactive flows in porous medium	29
2.3.1	Volume-averaged Navier–Stokes equations	29
2.3.2	Heterogeneous chemical reactions	31
2.3.3	Momentum equation for material-response codes	32
2.3.4	Plain fluid-porous media interface modeling	32
2.4	Conclusions	33
3	Pyrolysis models for the decomposition of thermal protection materials	35
3.1	Introduction	35
3.2	Experimental characterization of pyrolysis	36
3.2.1	Isothermal characterization	37

3.2.2	Linearly increasing temperature	39
3.2.3	Pyrolysis of PICA	40
3.3	Multi-component pyrolysis models	41
3.3.1	Single-stage kinetic equation	41
3.3.2	Multi-component kinetic equations	42
3.3.3	Analytical solution of the multi-component model with the one-parameter Šesták-Berggren equation	44
3.4	Competitive pyrolysis reactions	48
3.4.1	Simple competitive kinetic scheme	48
3.4.2	Generalized competitive kinetic model	49
3.5	Examples	50
3.5.1	Six-equations multi-component model and sensitivity analysis	50
3.5.2	Model of Goldstein	54
3.6	Summary of the chapter	55
4	Numerical simulations of ablation experiments using a unified flow-material approach	57
4.1	Introduction	57
4.2	Flow in pyrolyzing reactive porous media	58
4.2.1	Multi-component single-step pyrolysis model	60
4.2.2	Model for char ablation	61
4.2.3	Porous material properties	62
4.3	Results and discussion	63
4.3.1	Verification: Ablation Workshop test case	64
4.3.2	Pyrolyzing material subjected to arc-jet conditions	67
4.4	Conclusion	70
II	Uncertainty Quantification: Theory and Methods	75
5	Uncertainty quantification methods for computational models	77
5.1	Introduction	77
5.2	Fundamentals of probability theory and statistics	79
5.2.1	Probability space	79
5.2.2	Conditional probability and Bayes' theorem	80
5.2.3	Random vectors and multivariate distributions	80
5.2.4	Stochastic processes	82
5.2.5	Markov chains	82
5.2.6	Distinction between random and stochastic differential equations	83
5.3	Forward and inverse problems	84
5.3.1	Forward problem	84
5.3.2	Inverse problem	85
5.4	Markov chain Monte Carlo methods	88
5.4.1	The Metropolis–Hastings algorithm	88
5.4.2	Random-walk Metropolis–Hastings algorithm	89
5.4.3	Adaptive Metropolis–Hastings algorithm	90
5.4.4	Hamiltonian Monte Carlo	91
5.4.5	Algorithm based on an Itô stochastic differential equation	93
5.4.6	Change of variables based on the support of the posterior	95
5.4.7	Checking and monitoring convergence	96
5.5	Propagation of uncertainties	96
5.5.1	Monte Carlo sampling method	97

5.5.2	Polynomial chaos method	97
5.6	Summary	99
6	Development of a numerical toolbox for Bayesian inference and uncertainty propagation	101
6.1	Brief review of existing software	101
6.2	Purpose of the present code	102
6.3	Implementation	103
6.3.1	Dependencies	103
6.3.2	Input/output	103
6.3.3	The model class	103
6.4	Sampling	104
6.4.1	Sampling from a Bayesian posterior	105
6.4.2	MetropolisHastings class	105
6.4.3	GradientBasedMCMC	106
6.4.4	Gradient of the log-likelihood	106
6.5	Illustration: spring model	107
6.5.1	Model file	108
6.5.2	Input files	108
6.5.3	Run file	110
6.6	Uncertainty propagation	111
6.7	Future developments	112
III	Quantification of uncertainty in thermal protection material de- composition	113
7	Bayesian identification of pyrolysis model parameters for thermal protection materials	115
7.1	Introduction	116
7.2	Bayesian approach to inverse problem for parameter calibration	118
7.2.1	Prior density function	119
7.2.2	Choice of the approximate likelihood function	119
7.2.3	Change of variables	120
7.3	Adaptive-gradient informed sampling algorithms	121
7.3.1	Adaptive procedure for the estimation of the covariance matrix	122
7.4	Application: parameter inference of multi-component pyrolysis models	123
7.4.1	Uncertainty analysis for the production of CH_4	123
7.4.2	Bayesian parameter inference for PICA devolatilization at a high heating rate	130
7.5	Application: parameter inference of competitive pyrolysis models	133
7.5.1	Competitive reaction scheme for PICA	133
7.5.2	Bayesian parameter inference for the competitive pyrolysis model	139
7.5.3	Extrapolation to other heating rates	141
7.6	Conclusions	143
8	Bayesian inverse analysis of the Mars Science Laboratory atmospheric entry data for aeroheating and material response estimation	147
8.1	Introduction	147
8.2	Methodology for model parameter calibration and rebuilding of surface conditions from flight data	150
8.2.1	Calibration and propagation	150
8.2.2	Material-response solver	151

8.2.3	Computational domain	152
8.3	Results for the calibration and propagation of material thermal property uncertainty	153
8.3.1	Calibration of material thermal properties	153
8.3.2	Posterior predictive checks	156
8.4	Rebuilding of the environment conditions	160
8.4.1	Results of the calibration and posterior predictive checks	161
8.5	Summary and perspectives for future work	164
9	Propagation of uncertainty in pyrolysis model parameters for a Martian atmospheric reentry	169
9.1	Introduction	169
9.2	Numerical model	171
9.2.1	Pyrolysis model in PATO	171
9.2.2	Pyrolysis gas chemistry	172
9.2.3	Mesh and issues regarding the in-depth resolution	173
9.2.4	Comparison with a two-equation pyrolysis model	174
9.3	Results of the propagation of uncertainty	176
9.3.1	Temperature and gas elemental composition	176
9.3.2	Recession and pyrolysis production	179
9.4	Conclusion and future work	180
IV	Conclusions and directions for future work	183
10	Conclusions and perspectives	185
10.1	Summary and conclusions	185
10.2	Suggestions and perspectives for future work	186
10.2.1	Pyrolysis model inference	186
10.2.2	Inferring the input parameter distributions	187
10.2.3	Directions in uncertainty propagation	188
10.2.4	Validation of the models	188
V	Appendices	191
A	Analytical solutions and link between competitive and multi-component parallel pyrolysis models	193
A.1	Computation of derivatives for multi-component parallel models	193
A.2	Discussion on the existence of a solution for generalized competitive mechanisms	194
A.3	From generalized competitive models to multi-component models	195
A.3.1	Linear multi-component, multi-step competitive reactions	195
A.3.2	Link with multi-component parallel models	196
B	Additional implementation details	197
B.1	Effect of the reparametrization on the Metropolis–Hastings algorithm	197
C	Link with Bayesian statistics, additional graph of convergence for activation energy and table of input parameter values	199
C.1	Motivation for the expression of the approximate likelihood function	199
C.1.1	Link with Bayesian statistics	199
C.1.2	Likelihood function for time series measurements	200
C.1.3	Interpretation using the notion of push-forward measure	201
C.2	Influence of the initial covariance matrix	202

C.3	Graph of convergence for \mathcal{E}	204
C.4	Parameter value table	204
D	Tables for the probabilistic predictions for gas elemental composition	207

List of Figures

1.1	MSL entry, descent and landing sequence of events.	4
1.2	Physics of atmospheric entry flows and TPM response.	5
1.3	Atmospheric entry flight paths on a velocity-altitude map.	6
1.4	Reusable vs. ablative thermal protection systems.	7
1.5	Summary of high temperature gas-surface interactions.	9
1.6	Roadmap to uncertainty quantification.	10
2.1	The atmospheric entry is a multiscale phenomenon.	21
2.2	Pyrolysis and ablation of low-density carbon/phenolic materials.	22
2.3	Porous medium with a single solid phase with cylindrical-shaped fibers.	30
3.1	Gas chromatography analyses (Sykes [1967]).	38
3.2	Sample mass and mass yields (Wong et al. [2015b]).	39
3.3	Sample mass and mass yields at 366 K min^{-1} (Bessire and Minton [2017]).	40
3.4	Single-step competitive reaction.	48
3.5	Generalized pyrolysis scheme with competitive mechanisms.	49
3.6	Comparison of the analytical solution and a numerical integration.	52
3.7	Comparison of analytical sensitivities with finite differences.	53
4.1	Porous medium with multiple solid phases.	59
4.2	Computational domain of the Ablation Workshop test case.	64
4.3	Temperature, material recession and mass blowing rate (PATO vs. Argo).	65
4.4	Species partial pressures for the Ablation Workshop test case.	66
4.5	Computational domain and mesh for the iso-Q test.	68
4.6	Dimensions and mesh of the samples for the iso-Q test case.	69
4.7	Shape change for the iso-Q test case after 40 s.	70
4.8	Species partial pressures along the stagnation line.	71
4.9	Temperature field of the iso-Q test case after 40 s.	72
4.10	Total pressure field of the iso-Q test case after 40 s.	72
4.11	Streamlines around and within the material.	73
5.1	Comparison of the sampling of a Gaussian using RWHM and HMC.	93
5.2	Comparison of the position-momentum map for the ISDE and HMC.	95
7.1	Influence of the weights on the uncertainty in \mathcal{A} and \mathcal{E}	124
7.2	Bivariate posterior densities for the parameters \mathcal{A} and \mathcal{E}	126
7.3	Influence of the parameter ϑ on the three-equation model.	127
7.4	Convergence of the mean and variance for \mathcal{A} for varying MCMC strategies.	128
7.5	MCMC iterations for \mathcal{E} with the adaptive ISDE algorithm.	129

7.6	Convergence of the mean and variance for \mathcal{A} with different n_{end}	130
7.7	Convergence of the mean and variance for \mathcal{A} for ISDE for varying f_0	131
7.8	MCMC chains for selected parameters after 10^5 burn-in samples.	134
7.9	Marginal and bivariate posterior distributions for selected parameters.	135
7.10	Mass yields with $\vartheta = n_t$	136
7.11	Mass yields with $\vartheta = 0.7n_t$	137
7.12	Proposed competitive scheme for PICA at 10 and 366 K min ⁻¹	138
7.13	Comparison experiments vs. calibrated model at 10 and 366 K min ⁻¹	139
7.14	Mass loss and production rate curves with uncertainty ranges.	140
7.15	Correlation matrix.	142
7.16	Evolution of the different solid densities with their uncertainty intervals.	143
7.17	Extrapolation to other heating rates.	144
8.1	MSL heat shield and MEDLI sensor locations.	148
8.2	Temperature data and the MISP locations.	149
8.3	Parameter inference and uncertainty propagation using Dakota.	151
8.4	PATO 1D computational model with and w/o the environment.	153
8.5	MCMC chains for the effective thermal conductivity model parameters.	155
8.6	Posterior PDFs for the effective thermal conductivities model parameters.	157
8.7	Thermocouple data posterior predictive checks.	158
8.8	Comparison of thermal conductivities as a function of temperature.	159
8.9	A high surface heating can lead to thermocouple burn-out.	160
8.10	MCMC chains for the environment model parameters at MISP3.	162
8.11	1D and 2D posterior PDFs at MISP3.	163
8.12	Comparison of CFD and inferred boundary conditions at MISP3.	164
8.13	MCMC chains for the environment model parameters at MISP4.	165
8.14	1D and 2D posterior PDFs at MISP4.	166
8.15	Posterior thermocouple data predictive checks at MISP3 and MISP4.	167
9.1	Numerical mesh for the MSL heat shield.	173
9.2	Total pyrolysis production rate for increasing mesh cells.	174
9.3	Comparison of the nominal temperature for the two pyrolysis models.	175
9.4	Pyrolysis production rate for the two pyrolysis models.	176
9.5	Heat shield embedded temperature probes with uncertainty ranges.	177
9.6	Surface gas elemental composition with uncertainty ranges.	179
9.7	Recession of the heat shield at the surface probe positions.	180
9.8	Advective pyrolysis heat flux and total pyrolysis production.	181
C.1	Comparison of posterior sampling for different initial covariance matrices.	202
C.2	Convergence of the mean and variance for \mathcal{E} for varying MCMC strategies.	204
C.3	Convergence of the mean and variance for \mathcal{E} for different n_f	206
C.4	Convergence of the mean and variance for \mathcal{E} for different f_0	206

List of Tables

1.1	Summary of the content of the thesis and indication of the contributions.	15
3.1	Parameter values for the six-reaction model (Torres-Herrador et al. [2019b]).	51
3.2	Parameters values for two-reaction model (Goldstein [1969]).	54
4.1	Species mass fractions for the pyrolysis gas (Sykes [1967]).	65
4.2	Arrhenius coefficients for oxidation reactions.	68
7.1	Nominal parameter values for the three-reaction pyrolysis model for CH_4	125
7.2	Nominal parameters values from optimization algorithm.	139
7.3	Calibrated hyperparameters from Bayesian inference.	140
8.1	Values of the bounds of the prior support for MISP3.	163
9.1	Nominal elemental composition of the pyrolysis gas.	172
9.2	Heat shield embedded probe positions.	174
9.3	Heat shield surface probe positions.	178
9.4	Probabilistic predictions of the gas elemental composition (nose).	178
C.1	MAP values for the six-equation pyrolysis model for $\theta = n_t$ and $\theta = 0.7n_t$	205
D.1	Probabilistic predictions of the gas elemental composition (wind side).	207
D.2	Probabilistic predictions of the gas elemental composition (lee side).	208
D.3	Probabilistic predictions of the gas elemental composition (middle side).	208

Acronyms

aPC	arbitrary polynomial chaos
AMH	adaptive random-walk Metropolis–Hastings
ANOVA	analysis of variance
BL	boundary layer
BLT	boundary layer transition
CFD	computational fluid dynamics
CF/PR	carbon fibers/phenolic resin
CV	coefficient of variation
Dakota	Design Analysis Kit for Optimization and Terascale Applications
DGM	discontinuous Galerkin method
d.o.f	degree of freedom
DPLR	Data-Parallel Line Relaxation
DSMC	Direct Simulation Monte-Carlo
DTA	differential thermal analysis
EDL	entry, descent, and landing
ESA	European Space Agency
FIAT	Fully Implicit Ablation and Thermal Response
GC	gas chromatography
HMC	Hamiltonian Monte Carlo
ISDE	Itô stochastic differential equation
i.i.d.	independent and identically distributed
LAURA	Langley Aerothermodynamic Upwind Relaxation Algorithm
LTE	local thermodynamic equilibrium
MALA	Metropolis-adjusted Langevin
MAP	maximum a posteriori
MCMC	Markov chain Monte Carlo
MEDLI	Mars Science Laboratory Entry, Descent and Landing Instrumentation
MISP	MEDLI Integrated Sensor Plug
MSL	Mars Science Laboratory

NASA	National Aeronautics and Space Administration
ODE	ordinary differential equation
OLS	ordinary least-squares
PATO	Porous-Material Analysis Toolbox Based on OpenFOAM
PCE	polynomial chaos expansion
PDF	probability density function
PDE	partial differential equation
PS	project science
PICA	Phenolic Impregnated Carbon Ablator
QoI	quantity of interest
Queso	Quantification of Uncertainty for Estimation, Simulation, and Optimization
RWMH	random-walk Metropolis–Hastings
SDE	stochastic differential equation
TACOT	Theoretical Ablative Composite for Open Testing
TC	thermocouple
TGA	thermogravimetric analysis
TPS	thermal protection system
TPM	thermal protection material
UQ	uncertainty quantification
VANS	volume-averaged Navier–Stokes
VKI	von Karman Institute

Introduction

1.1 Research context: atmospheric entry and thermal protection systems for spacecraft

“The exploration of space will go ahead [...], and it is one of the great adventures of all time.”

John F. Kennedy, 1962

More than half a century after the first Apollo missions, the adventure of space exploration has become even more of a reality as humans are now preparing to go back to the surface of the Moon and further on Mars. The number of space missions that will need to enter into the atmosphere of a planet and safely deliver their payload is going to increase in the next decade ([Bousquet et al. \[2015\]](#)). With the Artemis program, the National Aeronautics and Space Administration (NASA) is planning to send astronauts to the Moon and increase the presence of humans in space, which will require to safely return on Earth. Missions that will collect samples on the surface of Mars or other celestial bodies will have to bring back the samples to Earth, such as for instance the Mars Sample Return campaign ([Joffre et al. \[2018\]](#)), a joint effort of the European Space Agency (ESA) and NASA. Future missions, increasingly more ambitious in scope, with heavier payloads and targeting more extreme environments, call for the development of new technologies and robust methods for the design of all parts of spacecraft.

One critical part of space exploration missions is the atmospheric entry, descent, and landing (EDL) phase. It begins when the spacecraft reaches the uppermost layer of the atmosphere, after its separation from the cruise spacecraft, until touchdown. For instance, Fig. 1.1 illustrates the EDL sequence of events for the Mars Science Laboratory (MSL), whose entry capsule successfully delivered the Curiosity Rover on the surface of Mars on the 5th August 2012 ([Way et al. \[2007\]](#); [Prakash et al. \[2008\]](#); [Vasavada et al. \[2012\]](#)).

First, during the atmospheric entry, the spacecraft follows a ballistic descent. As it reaches a hypersonic flow regime, a strong *shock wave* is created in front of the body and leads to a sharp increase of the temperature of the surrounding flow, typically more than 10000 K in the *shock layer* ([Anderson \[2006\]](#)). The spacecraft loses most of its kinetic energy and is exposed to a severe aerothermal environment. A *thermal protection system (TPS)* is essential to protect the spacecraft from these harsh conditions. For instance, the Galileo probe was designed to survive a peak heating of 30 kW cm^{-1} and a total heat load of 300 kJ cm^{-1} during its atmospheric entry on Jupiter while transmitting data until the pressure and temperature completely destroyed the spacecraft ([Milos \[1997\]](#)). Second, during the descent, parachutes are released to reduce further




Figure 1.1: MSL entry, descent and landing sequence of events (artist's concept). Image credit: NASA/JPL-Caltech.

the speed of the spacecraft. A series of parachutes can be deployed, and the heat shield can be separated from the main part. Eventually, the spacecraft operates its final descent and, for some missions, operates its landing, in which case the spacecraft also includes lander in its payload for the touchdown, such as for instance the MSL (Prakash et al. [2008]), or more recently the NASA's Perseverance mission to Mars. In all cases, the spacecraft has to go through the severe atmospheric entry and requires to be protected by a heat shield.

1.2 Atmospheric entry flows and thermal protection material (TPM) decomposition

Atmospheric entry flows are intimately related to hypersonic flows. There is no general criterion to define hypersonic flows, but it can be identified by the appearance of several complex multi-physics phenomena that become dominant as the Mach number increases (Anderson [2006]). The effects of high temperature flows are by far the most dominant effects that are responsible of two typical mission-killers, namely the *gas-surface interactions (GSIs)* and the *radiation*. Some of these effects are illustrated in Fig. 1.2 and described below.

Because of the presence of the spacecraft entering the atmosphere at high-speed velocities, the density of the flow changes drastically across the shock wave: the ambient air in front of the body is strongly compressed and intense frictional dissipation creates high temperatures. Part of the kinetic energy of the spacecraft is transferred to the internal energy of the flow molecules, which increases the gas temperature. The internal energy modes are progressively excited due to the rise of the translational temperature through the shock and molecules start to vibrate, dissociate and even ionize (and thus is commonly referred as a *plasma flow*) for sufficiently high temperatures. The gas can also emit radiation when going from an excited state to a less excited one, adding *radiative heating* to the surface of the TPS. The importance of the radiation varies from one mission to another. For low-speed entries, the radiative heat flux is low but increases fastly with velocity (Duffa [2013]). It also depends strongly on the atmosphere considered:




Figure 1.2: The hypersonic flow results in high enthalpy conditions with a bow shock and a rarefied region behind the capsule. The figure shows the temperature in the symmetry plane and flow streamlines. Image adapted from <https://boltzplatz.eu/>.

for example, the atmosphere of Titan is largely dominated by nitrogen (98.6%) with a small amount of methane (1.4%). During the atmospheric entry of the Huygens probe, the formation of cyanide (CN) molecules, which is known to be a strong radiator, was particularly prevailing in that case (Magin et al. [2006]).

The gas in the shock layer is hot because of the strong shock wave. In the *boundary layer (BL)* close to the surface of the TPS, where viscous effects are important, the high viscous dissipation can also generate high temperatures. Chemical reactions occur inside the gas and the surface of the vehicle is covered by a *chemically-reacting boundary layer* that interacts with the surface of the TPS. This chemically-reacting boundary layer leads to high heat-transfer rates to the TPS. The aerodynamic heating of the surface of the TPS induced by the hot boundary layer is called the *convective heating*. At the surface, the *recombination* of the atoms can be triggered. These phenomena strongly influence the behavior of the flow around the spacecraft and the interaction with its surface, and depending on the type of material used for the TPS, additional phenomena can occur. In any case, the high heat flux transmitted to the surface of the TPS results in the heating of its surface and the increase of its inner temperature, which can eventually cause the recession of its surface. Hence, a careful design of the heat shield is of utmost importance for the integrity of the vehicle and its payload, and for the safety of the crew in the case of manned missions.

1.2.1 Thermal protection materials

The type of materials used in TPS depends on the entry conditions. The TPS has to sustain the peak heating and all the heat loads encountered over the entry trajectory. Examples of entry trajectories expressed under the form of a velocity-altitude map are shown in Fig. 1.3 for certain reentries into Earth's atmosphere. Spacecraft coming from extra-terrestrial bodies (Mars Sample Return, Apollo missions) usually enter on Earth's atmosphere at a relatively high speed (> 11 km/s), and the aerobraking will generate very high heat fluxes that are transferred to the heat shield. Yet, for moderate velocity entries, such as for the Space Shuttle, the heat load on the spacecraft is lower.

Therefore, TPMs are classically divided into two main categories (Williams and Curry [1992]). On the one hand, non-ablative materials are used for reusable TPSs. They absorb the incoming




Figure 1.3: Atmospheric entry flight paths on a velocity-altitude map for the Space Transportation System (STS), or Space Shuttle, Apollo capsule and Mars Sample Return (MSR) missions. The peak heatings are indicated by dots. Real gas effects as dissociation (more than 10% after the shock) and ionization are drawn in dotted lines. For the comparison with a supersonic aircraft, the Concorde trajectory is also represented. Figure modified from [Howe \[1989\]](#).

heat from the shock and boundary layers and re-radiate the energy away from the surface towards the gas. They are more subjected to recombination because of the nature of the material used (metals, ceramic). Reusable TPSs are preferred for moderate speed entries, typically below 7.5 km/s and heat fluxes up to 1 MW/m². A well-known example of reusable TPS was implemented on the Space Shuttle, for which the underside part of the TPS was made up of silica ceramic tiles. Two other examples are the ESA Intermediate eXperimental Vehicle (IXV), shown in Fig. 1.4(a), and the Space Rider ([Bernard et al. \[2019\]](#)). The TPS tiles are made up of a thin outer layer of carbon fiber reinforced carbon-silicon carbide (C/SiC) material designed to withstand the mechanical loads and resist to the heat fluxes ([Pichon et al. \[2006\]](#); [Panerai \[2012\]](#)), while layers of insulating material are inserted underneath for thermal insulation (alumina blankets close to the outer surface and silica aerogels close to the cold structure). The Space Rider is planned to be used up to six times, with minimum maintenance between two flights ([Bernard et al. \[2019\]](#)).

On the other hand, ablative materials can accommodate the high heating rates and heat loads encountered during the reentry through phase change and mass loss. They lose their mass when heated, due to chemical reactions such as oxidation or nitridation ([Helber et al. \[2017\]](#)), or mechanical erosion which can further fasten the recession process. Generally ablative TPMs are filled with resin that also sublimates when subjected to high heat fluxes, called *pyrolysis reactions*. They are largely used for entry speeds higher than 8 km/s and heat fluxes exceeding 1.5 MW/m², typically for extra-orbital high-velocity entry missions. The AVCOAT material (density $\approx 510 \text{ kg m}^{-3}$) has been used for the TPS of the Apollo missions for direct lunar return at $\approx 11 \text{ km/s}$ ([Erb et al. \[1970\]](#)) and is planned to be used for the multi-purpose crew vehicle Orion ([Harris et al. \[2018\]](#)).

Lightness is an important factor to account for in order to reduce the energy costs associated to the propulsion of the spacecraft and for maximizing its payload. Yet TPSs must be strong enough to resist to mechanical erosion (spallation) and other possible mechanical failures of the material. By the end of the 80's, these requirements led to the development of a new class of low-density carbon/phenolic ablative materials, like the Phenolic Impregnated Carbon Ablator (PICA) material (density $\approx 280 \text{ kg m}^{-3}$) by NASA ([Tran et al. \[1997b\]](#)) and later on to




Figure 1.4: Examples of reusable and ablative thermal protection systems. (a) Illustration of the ESA IXV atmospheric entry (reusable TPS). (b) MSL heat shield made up of PICA ablative TPMs. Credit: NASA/JPL-Caltech/Lockheed Martin.

the development of Asterm by Airbus DS for ESA missions ([Ritter et al. \[2011\]](#)). These materials are typically made up of a carbon fiber matrix in which a resin component is impregnated. The fibers provide mechanical resistance and the resin absorbs the heat transmitted to the material, which makes these kind of material a suitable choice for thermal protection systems. Many types of resins have been studied (see a review in [Sharpe and Wright \[2009\]](#)), and phenolic resins, used for instance in the PICA material, are widely used. Under the increasing heat encountered during the reentry, decomposition of this reinforced fiber/resin composite materials occurs mainly in two steps ([Duffa \[2013\]](#)). First, the resin is progressively pyrolyzed, producing pyrolysis gases that percolate through the material and leave a carbonaceous residue on the fibers, the char. These pyrolysis gases add a blockage effect on the boundary layer gases reducing further the overall heating of the material by acting as a barrier against the incoming flow. Then, the charred material is ablated simultaneously by heterogeneous chemical reactions, sublimation, and spallation. The high porosity of such materials is not only efficient in terms of weight, but also reduces the effective thermal conductivity of the material and thus helps to keep the substructure at a relatively low temperature for the payload to be safe.

In order to understand the physical phenomena at stake, experimental tests are necessary. Studies on ablative materials date back to the 60's with the development of early material response codes ([Kendall et al. \[1968\]](#); [Chen and Milos \[1999\]](#)) based on the experiments of [Sykes \[1967\]](#) and [Goldstein \[1969\]](#). Several experiments have been carried out more recently to get new insight into the pyrolysis of the resin ([Trick and Saliba \[1995\]](#); [Trick et al. \[1997\]](#); [Wong et al. \[2015b, 2016\]](#); [Bessire et al. \[2014\]](#); [Bessire and Minton \[2017\]](#)), char ablation ([Panerai et al. \[2014, 2016\]](#)), or both pyrolysis and ablation of the material such as the experiments in the von Karman Institute (VKI) Plasmatron facility ([Helber et al. \[2016b,a\]](#)). The Zeram material ([Natali et al. \[2016\]](#); [Reimer et al. \[2018\]](#)) is now being tested in order to have open access materials properties for material simulations. Then, mathematical models are elaborated and numerical simulations are used to validate them based on the experimental results and to perform predictions. To predict the material thermal decomposition, current models are still based on these early experiments, although nowadays numerical solvers allow the use of more accurate models.

Low-density ablative thermal protection materials have proven to be successful during atmospheric entries of several missions such as Stardust ([Willcockson \[1999\]](#); [Kontinos and Stackpoole \[2008\]](#); [Stackpoole et al. \[2008\]](#); [Trumble et al. \[2010\]](#)) the Mars Science Laboratory ([Wright et al.](#)

[2011, 2014]), or more recently the SpaceX Dragon capsule (Seedhouse [2016]). Currently, only ablative materials allow to sustain the high heat fluxes encountered during high-speed entries. These materials will therefore enable future sample return missions and planetary space exploration programs.

1.2.2 Importance of the coupling between the flow and material ablation

The interactions of the hypersonic flow and the reactions occurring at the surface and within the heat shield are manifold and there is a need to provide a detailed modeling for a correct estimations of the heat fluxes transmitted to the spacecraft. Figure 1.5 summarizes the most dominant physical phenomena.

High heat fluxes are transmitted to the material by convection from the boundary layer and radiation from the shock layer, heating up the surface, and melting and sublimation of the surface can be observed. Conduction within the material increases the heat shield temperature. Species from the chemically-reacting boundary layer can diffuse towards the surface of the TPS and react, for instance, through oxidation or nitridation (heterogeneous) reactions that lead to the recession of the surface. Additional material can be removed by mechanical erosion (spallation) at the surface. For low-density porous fiber/resin materials, species can even diffuse below the surface, leading to the ablation of the bulk of the material (Lachaud et al. [2010]; Panerai et al. [2014]; Schrooyen [2015]). The increase of temperature induces the pyrolysis of the resin of the material, which loses its mass. The gaseous products of ablation and pyrolysis are blown into the boundary layer where they may react with the boundary layer gases coming from the shock layer. Species produced in-depth from volume ablation and pyrolysis reactions percolate towards the surface and add to the mixture. For strong mass blowing rates, the boundary layer can be blown off of the surface, significantly reducing the thermal gradients at the surface and decreasing the convective heat flux. Depending on their type and concentration, ablation and pyrolysis gases may block a significant amount of the incoming radiation as well, through absorption. This process, known as radiation blockage, is still not fully characterized, due to the complexity of the phenomena leading to the creation of the ablation and pyrolysis gases. To further complicate matters, certain ablation products such as CO, CN, and C₃ are known to be strong radiators, leading to the possibility that ablation may actually increase the radiative flux to the vehicle through spontaneous emission of these species as they are heated in the boundary layer (Scoggins [2017]).

Finally, as the flow moves across the surface of the ablator, the boundary layer thickness tends to grow. For larger vehicles, the laminar boundary layer flow can transition to turbulent, which may significantly increase the convective heating rate to the surface. Shock and boundary layer gases are convected towards the after-body of the vehicle (see Fig. 1.2). As the flow wraps around the leading corner of the vehicle, a significant decompression and separation occurs, generating a large, turbulent wake and leading to rarefied regions where low-density effects are important. As the wake cools, previously dissociated atmospheric gases and ablation products may recombine, adding radiative heat flux to the after-body. Although relatively smaller compared with the fore-body heating, a significant radiative heat flux can occur on the backshell due to recombination and exceeds the convective heating. For Martian entries, a significant radiative heat flux can occur on the backshell of the vehicle due to recombination of CO₂ which is a strong emitter in the infrared region (Brandis et al. [2020]). This phenomenon may become increasingly important for future missions aiming to send increasingly large vehicles to Mars, and for which it is therefore important to quantify the species that are present in the boundary layer and their associated uncertainty.




Figure 1.5: Summary of high temperature gas-surface interactions.

1.3 Uncertainty quantification for TPS modeling

In many engineering domains, the ability to predict the outcome of a physical system is critical to understand and advance the state of science and technology, and is essential for engineers throughout any design process. Predictive science is at the confluence of computational, physical, and mathematical sciences. The field of uncertainty quantification (UQ) has recently emerged at the intersection of these three pillars and proposes to investigate, based on the probability theory and the use of statistical tools, the role played by uncertainty in the model and the data on the behavior of complex science and engineering systems (Le Maître and Knio [2010]; Smith [2014]; Arnst and Ponthot [2014]; Ghanem et al. [2017]). Due to the interdisciplinary nature of the field, UQ encompasses a broad range of problems driven by real challenges of practical importance involving the interplay between physics, modeling, computational hardware, algorithmic complexity, and decisions.

Computational studies are crucial in many fields that study complex physical systems or when experimental campaigns are expensive and difficult to perform. Major advancements have been made in the field of mathematical modeling, scientific computing, and data science that have progressively improve the reliability of numerical predictions. The goal of uncertainty quantification in numerical analyses is to characterize the impact of all the sources of uncertainty, that can stem from the model inputs (such as parameter values, or initial and boundary conditions), the difference between the model and the physical system (model discrepancy), numerical approximations and errors, on the model output, with the ultimate objective to provide more reliable predictions for practical problems.

We can think of the simulation tool as a black-box model that provides some output values, or quantity of interests (QoIs), based on some input parameter values for the model, as illustrated in Fig. 1.6. The first step toward the quantification of uncertainty in the QoIs is the characterization of input uncertainties by means of probability density functions (PDFs) that rely on experimental data; these PDFs can then be propagated through the numerical model to obtain the uncertainty in the QoIs.

In present context of TPS design, the purpose of an uncertainty quantification study is twofold. Firstly, it can be used to assist the design of TPS by performing predictions on the environment and the material-response by using computational software and statistical projections for future mission design. By providing a statistical characterization of the output quantities of interest, one can rigorously assess uncertainty margins on the predictions. Secondly, post-flight data acquired from on-board instrumentation during an actual atmospheric entry of a spacecraft can be used for inverse analysis in order to reconstruct unknown quantities that can be compared to the pre-flight analyses. By comparing the two, one can assess the validity of the




Figure 1.6: Roadmap to uncertainty quantification and illustration of uncertainty propagation. In this illustration, the input uncertainty is characterized by samples drawn from the input probability distribution. The samples are propagated through the computational model by means repeated simulations of the computational model in order to obtain the distribution of the values for the quantity of interest. The central picture illustrates the computational model, which represents the recession of the surface of the MSL heat shield during its atmospheric entry on Mars (from [Meurisse et al. \[2018\]](#)).

current predictions, which will in turn ultimately enhance the reliability of the prediction and help reducing the margins for future missions thanks to the reduction of uncertainty.

1.3.1 Assessing uncertainty margins on the predictions

The typical process for designing the thermal protection system is to calculate a nominal thickness by performing numerical simulations using predictive tools, apply mission margins to the thickness, and verify the results with appropriate experimental tests ([Wright et al. \[2014\]](#)). For design purposes, numerical tools have been mostly used for simple predictions that are assigned conservative safety margins. For instance, the nominal prediction for the Stardust¹ mission requires a thickness of 4.83 cm ([Olynick et al. \[1999\]](#)), and a safety margins of 20.5% was added, resulting in a baseline thickness of 5.82 cm ([Tran et al. \[1997a\]](#)). With the constant improvement of computer capabilities, more sophisticated models can be implemented ([Bianchi et al. \[2012\]](#); [Martin \[2013\]](#); [Martin et al. \[2012\]](#); [Bianchi et al. \[2013\]](#); [Chen and Gökçen \[2014\]](#); [Martin and Boyd \[2015\]](#); [Schrooyen et al. \[2016\]](#); [Lachaud et al. \[2017\]](#)) and quantification of the uncertainties in the model parameters and in the predictions can be performed ([Sanson et al. \[2018\]](#); [Rivier et al. \[2019\]](#)). Although baseline predictions can nowadays be performed using high-fidelity simulation tools, the estimation of their associated uncertainties has received less attention. Therefore, the specific mission margins remain high.

Predictive tools are developed based on models and calibration parameters that can be obtained from small-scale ground-based experiments in a controlled environment, and from which the calibration parameter PDFs can be inferred. In this context, the uncertainty represents our lack of knowledge of the underlying physics of the problem and measurement errors, and results in a variability on the observed data. This lack of knowledge is due, for instance, to the inherent variability of real-world manufacturing processes which lead to an unknown variability in the material properties. Indeed, the preparation of the sample to be tested, or the characteristics of the material at the microscopic level can be different from one experiment to the other and lead to variability in experimental results ([Panerai et al. \[2017\]](#)) which will ultimately lead to uncertainties on the inferred material properties. The experimental device itself can also lead to observation errors, and the mathematical model used in the simulation tool can have inaccuracies, that will also impact on the inferred model parameters. All these uncertainties are

¹Stardust (1999–2006) was a mission designed to collect extra-terrestrial materials from the coma of the comet Wild-2.

difficult to separate, and the approach taken in this work will consist in attributing them to the material parameter variability such that it can be propagated in numerical solvers. Therefore, there is a need for a robust characterization of the parameter uncertainty and estimation of the propagated uncertainty.

1.3.2 Enhancing the reliability of the models

Identifying inaccuracies in models can be tough as the atmospheric reentry is a complex multi-physics problem that calls for identifying models for, just to name a few, the atmospheric conditions, the chemistry (equilibrium or finite-rate), the radiation, turbulence, or the flow regime (Navier–Stokes equations in continuum regime or Boltzmann equations in the rarefied regime). The input model parameters that have an impact on the aerothermodynamics are also various, such as for instance, the atmospheric density and its variability, or the wind speed ([Vasavada et al. \[2012\]](#)); flight parameters, such as the vehicle velocity, angle of attack or the reentry path; or parameters from the models used within the computational fluid dynamics (CFD) codes. Ground-based experiments are essential for the validation of computer models, but they cannot substitute for actual flight environment. Results from post-flight analyses are necessary to assess the impact of extrapolating the models developed from ground-based experiments to flight. Hence, the comparison of post-flight data with the pre-flight predictions will ultimately help identifying whether the current models used in the simulations accurately reproduce the physics.

Post-flight data provide a global response of the real-size TPS and an adequate instrumentation on the heat shield can provide useful measurements about the physical phenomenon encountered. For instance, the MSL mission flight data consist of pressure and temperature probes were embedded within the heat shield, hence only allowing an indirect observation of the conditions at the surface. Such real-world experiments being rather scarce, a rigorous reconstruction of the environment using material response codes and the available flight data while taking into account the measurement uncertainty is thus of interest. Post-flight analyses, such as the reconstructed environment from the measurements of the thermocouple, can be ultimately compared to the pre-flight analysis and CFD predictions to demonstrate the simulation capabilities ([White et al. \[2013\]](#)) and in order to improve the numerical models and the knowledge of their parameters and ultimately reducing their uncertainties.

1.4 Towards a robust characterization of uncertainty in aerothermal flow-material simulations

Thermoanalytical studies involve using experiments to determine the reaction process provided a given temperature increase and characterize the species that are produced through the reaction. New experimental studies have recently been reported relevant to the pyrolytic decomposition of the phenolic resin used in carbon/phenolic composites, namely, in [Wong et al. \[2016\]](#) and [Torres-Herrador et al. \[2021\]](#), under several isothermal conditions, and in [Bessire and Minton \[2017\]](#) for linearly increasing temperatures. In these experimental studies, the mass of a sample is monitored as a function of temperature using thermogravimetric analysis (TGA), and the identification of the volatile decomposition products are measured by means of a mass spectrometer (in [Wong et al. \[2016\]](#)) or gas chromatography (in [Bessire and Minton \[2017\]](#) and [Torres-Herrador et al. \[2021\]](#)).

The class of pyrolysis models adopted represents the decomposition mechanisms in terms of a system of chemical kinetic equations with reaction rates that obey an Arrhenius-like generic model. Based on the observations, a mathematical model can be deduced (model inference) and the parameters of the model can be inferred. Several works already considered the inference of the kinetic parameters in order to update pyrolysis models based on recent experiments. Most of

the inference studies are performed in a deterministic framework, mainly by using optimization methods to obtain best-fit estimates (Lautenberger and Fernandez-Pello [2011]). In a previous work by Torres-Herrador et al. [2019b], the new experimental data of Bessire and Minton [2017] for PICA TPM have already been used to estimate parameters of a pyrolysis model, but the model identification was formulated as a deterministic inverse problem, and an estimation of associated uncertainties is thus still required.

There have been several efforts to infer kinetic parameters of Arrhenius-like chemical reaction models from experimental data with a probabilistic approach, especially for combustion applications, see, for instance, Najm et al. [2014]; Cheung et al. [2016]; Khalil and Najm [2018]; Urzay et al. [2012], in which a Bayesian approach (Tarantola [2005]; Stuart [2010]; Kaipio and Somersalo [2005]) is adopted in order to characterize the model parameter uncertainties. The Bayesian approach consists in assigning a prior PDF to the parameters by using the current knowledge of the parameters and to update this probability distribution by incorporating the physical model and the new information (measurements) with their uncertainty through the likelihood function, resulting in the so-called posterior PDF (Gelman et al. [2014]; Smith [2014]). Investigations of pyrolysis experiments by using a Bayesian approach have been performed for high-impact polystyrene, bisphenol-A polycarbonate, and polyvinyl chloride at two different heating rates for fire applications in Bruns [2015]. For TPMs decomposition, whereas new experimental data have already been published and new experimental procedure are still being developed, they are yet to be exploited for the inversion of physico-chemical models in a probabilistic framework.

In addition, the issue of uncertainty quantification, as required to assess the reliability of the numerical model and the physico-chemical models, is yet to be addressed. Numerical studies on thermal protection systems usually assign standard probability distributions to characterize the input uncertainties (Wright et al. [2007]; Copeland et al. [2012]; Rivier et al. [2019]), either from expert knowledge or from pre-flight ground-based experiments, which do not take into account the potential correlations between input parameters. It is therefore desirable to infer the probability distribution of the inputs from available experimental data in a stochastic framework, then to propagate the input probability distribution through high-fidelity numerical codes to obtain a probability distribution for quantities of interest. A correct estimation of those uncertainties will eventually lead to a robust estimation of errors on numerical simulations for better assessing safety margins.

1.5 Objectives and contributions of the thesis

The new experimental data that have recently been published for the characterization of the pyrolysis of carbon/phenolic TPMs are yet to be exploited for the inversion of the physico-chemical models, and the issue of uncertainty quantification is yet to be addressed. With the overarching goal of contributing to the development of an uncertainty-quantified numerical model for ablation of new porous fiber/resin composite materials, this thesis targets the following objectives:

1. the inference from new experimental data of an uncertainty-quantified pyrolysis model,
2. the uncertainty propagation through high-fidelity numerical models for in-flight predictions for an atmospheric entry on Mars.

For the first objective, we formulate the problem as a Bayesian probabilistic inverse problem, in which we use an approximate likelihood function involving a weighted distance between the model predictions and the experimental data. This formulation provides a probability distribution over the model parameters, the posterior, which assigns higher probabilities to parameter values that correspond to predictions that better fit the experimental data. In general, computing the posterior thus obtained from Bayes' theorem is not feasible and sampling algorithms

are usually required in order to obtain samples from the probability distribution (as sketched on the right in Fig. 1.6). To sample from the posterior, we investigate gradient-informed sampling methods and two improvements are proposed in order to enhance the efficiency of the sampling: a reparameterization of the pyrolysis equations, and an adaptation of the approximation of the so-called covariance matrix. In order to have a versatile and flexible toolbox for implementing and testing the proposed improvements, we develop a numerical code for performing Bayesian inference and uncertainty quantification in the Python language, named hereafter **Py-bitup** (Chapter 6). Hence, a first contribution of this thesis is to propose a new pyrolysis model for PICA decomposition with uncertainty-quantified parameters characterized from the samples drawn from the posterior distribution, which is representative of the experimental observations. This contribution and the proposed improvements are exposed in the first part of Chapter 7.

For the second objective, we provide probabilistic projections of the recession of the surface of an ablating TPS, the production of gaseous species at the surface and the temperature. The new pyrolysis model and the model parameter distribution inferred using the Bayesian approach are used for the propagation of uncertainty through a high-fidelity code for material-response modeling. The posterior PDF can be used as input parameter PDF and the samples can be used for direct propagation. Hence, a second contribution of this thesis is the assessment of the impact of the uncertainty in pyrolysis model parameters on the response of an ablating TPS relevant to in-flight performance prediction, and this is exposed in Chapter 9.

In addition to these two objectives and the foregoing contributions, we also contribute to three supplementary research works that are also presented in this manuscript:

1. The high-fidelity numerical code Argo is being constantly improved by a research team at Cenaero. The code is able to perform multi-dimensional simulations of complex flow-material interactions including the ablation of thermal protection materials. In close collaboration with Dr. P. Schrooyen and following the work initiated in [Schrooyen \[2015\]](#) and subsequently in [Coheur \[2016\]](#), we contribute to the implementation within Argo of a model for the resin pyrolysis and char ablation in a unified flow-material approach. To illustrate the novel approach for solving aerothermal flows around and through ablative porous materials, we carried out simulations of a pyrolysis-ablation experiment of a TPM subjected to a high enthalpy flow, which results are presented in Chapter 4.
2. Within the NASA Supercomputing Division at NASA Ames Research Center, researchers from the team of Nagi N. Mansour are developing the material-response code PATO to study the response of ablative thermal protection systems. We contribute to the analysis of the post-flight entry data of the Mars Science Laboratory in order to improve the modeling of the thermal protection system using PATO and the Dakota code for performing uncertainty quantification, which results are presented in Chapter 8.
3. With F. Torres-Herrador from the von Karman Institute and J. Blondeau from the *Vrije Universiteit Brussels*, a novel model for PICA decomposition featuring a competitive mechanism was introduced and presented in the co-authored work [Torres-Herrador et al. \[2019a\]](#). The focus of the research is on establishing the relevance of accounting for competitive mechanisms in a new pyrolysis model for PICA, to which we contribute to the inference of model parameters in a Bayesian probabilistic framework. These results are presented in the second part of Chapter 7.

1.6 Overview of the manuscript

The manuscript is divided into five parts with a total of 10 chapters, including this introduction, and four appendices. The three first parts constitute the main content of the thesis. Part I (Chapters 2–4) addresses the modeling of atmospheric entry problems and pyrolysis, and presents deterministic simulations of ablation experiments. Part II (Chapters 5–6) is concerned with the

review of theory and methods for performing uncertainty quantification and introduces the toolbox developed in this work for Bayesian inference. Part III (Chapters 7–9) presents the methodology developed in this work and its application to the quantification of uncertainty in model parameters, and presents the probabilistic projections for an atmospheric entry problem. Conclusions are drawn in Part IV (Chapter 10) and the appendices in Part V (Appendices A–D) provide additional information about this work. In Table 1.1, we summarize the content of the different chapters and we highlight whether a chapter is a review of the literature or presents new material.

Chapter 2 presents a review of the physics of atmospheric entry flows and TPM decomposition (pyrolysis and ablation), and gives an overview on the past and current computational tools used to simulate the flow-material interaction. We present the governing conservation equations of mass, momentum and energy and we introduce basic concepts related to high temperature flows. Then, the volume-averaged Navier–Stokes equations are reviewed, which allow to conciliate the flow and the material response within a single domain of computation (unified approach) and are used to simulate the flow within and around porous ablative TPMs.

Chapter 3 focuses on the models that are used to simulate the resin pyrolysis of thermal protection materials and presents both a review of current modeling and new contributions of this thesis. Pyrolysis experiments and the recent experiments on PICA are briefly reviewed. To simulate the resin pyrolysis, two types of model are investigated. Firstly, multi-component (parallel) reaction models are presented. For these models, we contribute to derive a semi-analytical solution that significantly speeds-up the computational time for the simulation of pyrolysis experiments, and we perform a sensitivity analysis of the production rate on the model parameters. This solution and the results of the sensitivity analysis are presented in the publication [Coeur et al. \[2021b\]](#). Secondly, a more detailed modeling for pyrolysis that includes competitive mechanisms is presented. Such mechanisms are already widely used for biomass pyrolysis, but were only introduced recently for TPM pyrolysis. The basics for competitive mechanisms and a pyrolysis model including a competitive mechanisms for PICA are presented in the co-authored work [Torres-Herrador et al. \[2019a\]](#).

Chapter 4 presents the use of the volume-averaged Navier–Stokes equations and of the unified flow-material approach to solve pyrolysis-ablation problems. We first compare with a state-of-the-art material response code the results of the simulation of a one-dimensional pyrolyzing material for verification using a multi-component mixture for pyrolysis gas production. We present an axisymmetric simulation of a material subjected to a high enthalpy flow, that undergoes pyrolysis and ablation with a six-species mixture with simplified pyrolysis gas production. The test problems are set up within the multidimensional tool Argo, developed at Cenaero.

Chapter 5 is an introduction to the theory and the methods for uncertainty quantification in computational models. We present the Bayesian framework that we use to characterize input parameter distributions, and we describe sampling algorithms that allow us to generate samples from the posterior distribution. We also discuss more advanced sampling algorithms that use the gradients of the posterior distribution to improve the convergence of sampling algorithms towards the right solution. For uncertainty propagation, we present the direct Monte Carlo propagation method and non-intrusive polynomial chaos methods, which aim at replacing the expensive computational model by a (cheaper) surrogate model in order to reduce the computational cost of uncertainty propagation.

Chapter 6 describes **Pybitup**, a numerical toolbox written in Python for performing Bayesian inference that was developed during this thesis. It implements the state-of-the-art sampling algorithms described in Chapter 5 as well as the proposed improvements later described in Chapter 7. We describe the structure of the code and provide examples. We discuss the future

developments sought for the code regarding uncertainty propagation, surrogate modeling and sensitivity analysis.

Chapter 7 focuses on the characterization of input parameter uncertainties, and in particular the parameters related to the kinetic model for the pyrolysis of the material, for which the values and related uncertainties are deduced (inferred) from the recent dedicated on-ground experiments within a Bayesian framework using an approximate likelihood function. We describe a methodology that seeks to improve the convergence of gradient-informed sampling algorithms by adapting tuning parameters of the algorithms. We demonstrate the application of the uncertainty quantification methods for parameter inference to the two type of pyrolysis model described in Chapter 3. First, the algorithms with the adaptation of the tuning parameters are applied to the parallel pyrolysis models, which are cheaper in terms of computational time. The methods and these results are presented in the manuscript [Coheur et al. \[2021a\]](#). Then, we apply the probabilistic framework with Bayesian methods to the inference of competitive pyrolysis model parameters, and investigate the influence of the competitive mechanisms applied to PICA decomposition. This second contribution is presented in the co-authored manuscript [Torres-Herrador et al. \[2019a\]](#).

Chapter 8 presents the results for the inference of material properties and the reconstruction of the aerothermal conditions at the surface of the heat shield. The post-flight data of the MSL mission are used and we apply a Bayesian probabilistic method to the material-response code PATO to investigate how the uncertainty in the post-flight data affect the reconstructed material and environment properties. We discuss the comparison of such results with pre-flight analyses.

Chapter 9 addresses the uncertainty quantification of the material response of the heat shield during a Mars atmospheric entry. We provide probabilistic projections on the heat shield recession, gas composition at the surface and temperature from the uncertainties in the pyrolysis model parameters.

Chapter 10 gives a summary and draws the general conclusions of the thesis. Directions for future research in atmospheric entry modeling and uncertainty quantification are also provided.

Keywords	Chapter	Content
Atmospheric entry / Numerical modeling	2	Review
Pyrolysis / Modeling	3	Review / new
Pyrolysis-ablation / Unified approach	4	New
UQ / MCMC / PCE	5	Review
Numerical toolbox / Sampling / UQ	6	Review / New
Inference / MCMC / Pyrolysis	7	New
Post-flight / Inference / Aeroheating	8	New
Atmospheric entry / Predictions	9	New

Table 1.1: Summary of the content of the thesis and indication of the contributions.

Part I

Modeling of thermal protection systems for atmospheric entry

Physics of atmospheric entry flows

This chapter is an introduction to the modeling of aerothermal flows and the decomposition of thermal protection materials. We first discuss the underlying hypotheses of this work and briefly review the approaches to model atmospheric entry flows and ablative material response (Section 2.1). We then present some fundamentals of high temperature gas effects, in particular thermodynamic properties and chemistry (Section 2.2) for flows in thermal equilibrium. Next, we present the Navier–Stokes equations for compressible, viscous reactive flows that are used to model the external flow around the spacecraft (Section 2.2.3). Lightweight ablative materials involve some porosity, and species from the external flowfield can penetrate inside the TPS. The description of the flowfield within the pores can be required for an accurate modeling of the ablation of the material. Hence, to model the flow within the porous medium, we present the volume-average theory (Section 2.3). By spatially smoothing the fluid mechanics equations, valid in the plain-fluid region, over the whole porous domain, their validity is extended to the whole domain. The spatially-smoothed version of the Navier–Stokes equations for compressible, viscous reactive flows, the so-called volume-averaged Navier–Stokes (VANS) equations, are then stated. With an appropriate treatment of the fluid-material interface, the VANS equations can be used to simulate in the same computational domain both the external and the internal flows, named hereafter as the unified approach. The focus is given here on the VANS equations because they are implemented within the numerical solvers used in this work, namely the code Argo, that features a unified approach, and the Porous-Material Analysis Toolbox Based on OpenFOAM (PATO) code, that focuses on the material response.

2.1 Hypotheses for the modeling of aerothermal flows and ablative porous material response

Atmospheric entry flows are by essence high-temperature flows. The high temperatures affect the thermodynamic properties of the flow (such as its energy, enthalpy, specific heat), the transport properties (e.g. viscosity, thermal conductivity), promote chemical reactions and lead to radiation to and from the gas. If the temperature is high enough to cause ionization, the gas becomes a plasma and electromagnetic forces could be included. As we mainly focus on the part of the flow that is close to the vehicle, electromagnetic effects will be neglected, but these effects can be important downstream of the shock layer and significantly influence the flow properties.

We start our discussion on high-temperature flows by discussing the flow regime that permits the description of the flow as a continuum. Next, the flow-material interactions considered are briefly described and finally we shortly review strategies used to model atmospheric entry flows and their interaction with the TPS. The next section (Section 2.2) will delve into the properties

of the flow and chemical reactions.

2.1.1 Continuum versus rarefied flow regimes

A gas is a collection of particles (molecules, atoms, ions, electrons, etc) that move in a more or less random motion in a volume V (m^3). The gas particles can interact with their neighboring particles due to intermolecular forces, that are created by the electronic structures of these particles. A *perfect gas* is a gas where intermolecular forces can be neglected. For perfect gases, the pressure p (Pa), temperature T (K) and density $\rho = m/V$ (kg m^{-3}), with m (kg) the mass of the gas, are related through the perfect gas equation

$$p = \frac{\rho \mathcal{R} T}{M}, \quad (2.1)$$

where $\mathcal{R} = 8.314$ is the universal gas constant ($\text{J g}^{-1} \text{K}^{-1}$) and M the molar weight of the gas (kg mol^{-1}). It can be expressed equivalently by

$$p = nk_{\text{b}}T, \quad (2.2)$$

where n is the number of particles per unit volume, k_{b} is the Boltzmann constant, defined as the ratio $\mathcal{R}/N_{\text{A}} = 1.38 \times 10^{-23}$ (J K^{-1}), with N_{A} the Avogadro's number which represents the number of particles per moles. Except at very high pressures or low temperatures, intermolecular forces can be neglected and in the present context of aerothermal flows, the perfect gas assumption can be made.

Following its reentry trajectory, a spacecraft will generally encounter different flow regimes for which the flowfield dynamics need to be properly addressed. At low altitudes, typically below 80 km on Earth, the flow is considered to be in the *continuum* regime: the collisions among the particles are important and the flowfield can be described by the Navier–Stokes equations. It is the flow regime for most of the practical aerodynamic applications. However, at higher altitudes, the flow starts to be *rarefied*, the gas density is decreasing and collisions among particles occur less often. In order to describe the physics happening at those higher altitudes (above 120 km on Earth), the gas must be considered at the molecular level.

The characterization of the flow as a continuum is based on the assumption that the mean free path λ , that is, the average distance traveled by a single particle between two collisions, is small compared to a characteristic length L_{c} . As a first approximation, the mean free path for a single-species gas is given by

$$\lambda = \frac{1}{\sqrt{2}\pi d^2 n}, \quad (2.3)$$

where d is the kinetic diameter of the particle. The quantity πd^2 is frequently called the *collision cross section*. The ratio $1/\pi d^2 n$ can be seen as the mean free path of a particle without accounting for the velocity of the other particles, and the factor $\sqrt{2}$ is a correction that accounts for the relative speeds (Vincenti and Kruger [1965]). Using Eq. (2.2) to estimate the number n of particles per unit volume, we can see that the mean free path is proportional to the temperature and inversely proportional to pressure. Then, the validity of the continuum regime can be verified by computing the Knudsen number, defined as the ratio

$$\text{Kn} = \frac{\lambda}{L_{\text{c}}}. \quad (2.4)$$

For $\text{Kn} < 0.1$, the continuum assumption can be made and the Navier–Stokes equations are applicable. For $0.1 < \text{Kn} < 10$, slip effects are important (the traditional no-slip condition used at the interface between the fluid and a solid surface begins to fail). For $\text{Kn} > 10$, the flow is considered to be rarefied. When considering rarefied flow regimes, individual impacts on the




Figure 2.1: Sketch of an atmospheric entry and zoom on the porous region of an ablative TPM. The illustration of the porous medium is adapted from [Schrooyen \[2015\]](#).

surface are important and the aerodynamics and heat transfer are approached from a statistical point of view. The Navier–Stokes equations are no longer valid and the Boltzmann equation is used to describe the evolution of the velocity of the particles in the gas, and methods such as the direct Direct Simulation Monte-Carlo (DSMC) can be used to solve it ([Bird \[1994\]](#)).

During the first part of the reentry, at high altitudes, there are less particles in the gas, the pressure is low and the mean free path of the particles is usually large. The flow is in the rarefied regime, but the aerobraking of the spacecraft and the heating of the TPS are lower. In this work, flight conditions that correspond to later parts of reentry trajectories are considered, in which continuum equations are valid at the scale of the heat shield ($L_c \sim 1$ m) and in which the ablation of the TPM can be substantial. For example, typical pressure and temperature at the peak heating trajectory point for the reentry of a spacecraft in the Earth atmosphere are 0.1 atm and 10000 K, respectively, and assuming a kinetic diameter for air equal to 35×10^{-10} m (oxygen and nitrogen have similar kinetic diameters), this leads to $\text{Kn} = 2.5 \times 10^{-7}$.

2.1.2 Flow-material interactions

Within the heat shield, the characteristic length is much smaller ($L_c \sim 10^{-6}$ m), see Fig. 2.1, of the order of the length of the fibers that constitute the composite material. Physico-chemical reactions happen at the microscopic scale within the heat shield, such as pyrolysis reactions, or oxidation of the carbonaceous char layer. During the atmospheric entry phase, part of the heat flux is transferred inside the TPS, leading to a gradual increase of the temperature of the material. With this increase, large temperature gradients are established and the fiber/resin material is thermally degraded and removed by two physico-chemical phenomena: pyrolysis and ablation. Four characteristic zones can be identified within the material:

- **Virgin zone:** region where the material is still in its virgin state and no reaction occurs yet.
- **Pyrolysis zone:** region where the pyrolysis reactions and mass loss occur.
- **Char zone:** region where all the disposable constituents from the resin have already been consumed. No more pyrolysis reaction can occur.
- **Ablation zone:** subregion of the char layer in which the ablation of the charred material is found to occur. Its thickness depends on the Thiele number.

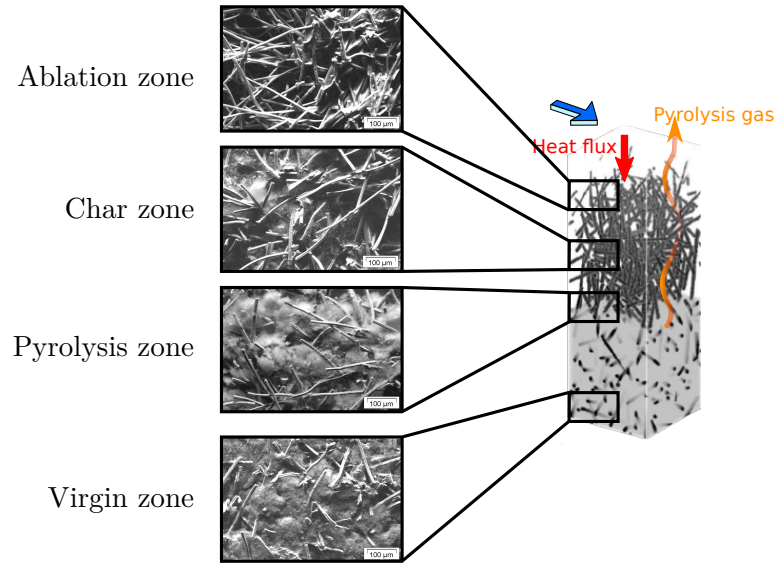


Figure 2.2: Illustration of the decomposition zones during the ablation of the PICA thermal protection materials from the virgin zone (bottom) to the ablation zone (top). The images showed on the left are scanning electron microscope micrographs of PICA (from [Lawson et al. \[2010\]](#)). On the right, result of a 3D microscopic simulation of the oxidation of PICA (from [Lachaud et al. \[2010\]](#)). Adapted from [Schrooyen \[2015\]](#).

In the pyrolysis zone, the polymeric resin is thermally degraded and is losing mass while releasing pyrolysis gases in the form of water, hydrogen and hydrocarbons. This thermal degradation can be separated into three (overlapping) temperature regions characterized by the reactions occurring in each range, which may be described by four dominant heterogeneous reactions ([Trick and Saliba \[1995\]](#); [Trick et al. \[1997\]](#)) and the production of pyrolysis gas increases the internal pressure inside the sample. [Lachaud et al. \[2015\]](#) showed numerically that these pressure gradients led to a flow that separates into two streams: one going toward the upper surface and blocking the gases from the boundary layer and one going toward the shoulder of an iso-flux sample. The process of pyrolysis decomposition will be addressed in more detail in Chapter 3.

In the ablation zone, the ablation of the char is caused by heterogeneous chemical reactions (oxidation, nitridation), phase change (sublimation) and/or mechanical erosion (spallation). Most of the current ablation models consider that ablation is a surface phenomenon (all the mass is lost at the surface) ([Li et al. \[2015\]](#); [Lin \[2007\]](#)). However, for lightweight porous ablative materials, the decomposition of the resin and the ablation of the fibers should also be treated within the porous medium, referred to as *volume ablation*, for accurate predictions of the heat shield temperature and overall surface recession. Indeed, because of the high porosity of the material, the reacting gas from the outer flow, mostly oxygen, diffuses inside the layer and reacts in the volume of the porous material with the carbon fibers. For example, post-flight analyses of the Stardust mission, for which the TPS was made up of PICA material, show that the models with surface ablation over-predict the recession of the surface of the TPS and surface density, while in-depth density profiles are in good agreement with the predictions ([Stackpoole et al. \[2008\]](#)), suggesting that close to the surface, subsurface oxidation may occur. Experimental evidence ([Panerai et al. \[2014\]](#)) also shows that for certain pressure and temperature conditions, flow species can diffuse within the material and react in-depth with the fibers. When such subsurface ablation occurs (sublimation, oxidation, erosion), it leaves the carbon fiber under the surface unprotected, which could ultimately lead to the removal of a more reactive (and fragile) matrix.

Therefore, for the study of volume ablation, [Lachaud et al. \[2010, 2011b\]](#) suggests to spatially

average the mass balance equations within the pores of the material (at the microscopic scale) in order to derive equations at the macroscopic scale of interest, thus allowing subsurface species diffusion and account for in-depth fiber ablation. Following this approach, [Schrooyen \[2015\]](#) performed the spatial smoothing of the Navier–Stokes equations, valid in the fluid phase, over the porous medium to obtain the so-called VANS equations, that are valid everywhere ([Whitaker \[1999\]](#)). By doing so, information regarding the complex structure of the porous medium is scaled-up, thus providing a macroscopic view of the local information. The volume-average framework also allows to implement a unified flow-material approach for aerothermal simulations, in which both the porous material and the surrounding fluid are simulated within the same computational domain. As it will be discussed in the next section, the unified approach in turn avoids the need for deriving boundary conditions at the interface and allows to have a time accurate flow-material interaction.

Thus, ablation occurs in a region that extends from the surface of the material to the beginning of the char zone. The ablation of the char occurs due to heterogeneous chemical reactions (oxidation, nitridation), phase change (sublimation) and/or mechanical erosion (spallation). Mechanisms of ablation depend on the reentry conditions. For moderate temperature and pressure, typically for $T < 3000$ K, $p < 55$ atm, chemical ablation dominates ([Palaninathan and Bindu \[2005\]](#)). The key parameter for this in-depth ablation zone is the Thiele number, defined as the ratio of the reaction rate and the diffusion rate:

$$\text{Th} = \frac{L_c}{\sqrt{D_{\text{eff}}/S_f k_f}}, \quad (2.5)$$

where D_{eff} is the effective diffusion inside the porous medium, S_f the specific surface and k_f is the reaction rate coefficient of the fibers. The characteristic length L_c can be, for instance, the depth of the char zone. At high Thiele numbers ($\text{Th} > 50$), ablation is mostly a surface phenomenon because the diffusion process is not fast enough. Conversely, at a low value of the Thiele number, the diffusion is much faster than the reaction rate and the thickness of the ablation zone becomes larger and volume ablation is promoted.

The value of the Thiele number depends on the flow conditions. For a typical plasma test experiment in the VKI Plasmatron ([Helber et al. \[2016b\]](#)), the Thiele number was found to be higher than 300, indicating a surface limited ablation ([Schrooyen et al. \[2017\]](#)). For an ablation experiment in a flow tube reactor ([Panerai et al. \[2014\]](#)), performed at a lower pressure, a value of the Thiele number around 2 was found, indicating a volume ablation regime.

2.1.3 Numerical strategies for flow-material simulations

Because of the complexity of the physical phenomena that need to be accounted for in the modeling and the differences in the length scales, most of state-of-the-art codes for the computation of aerothermal flows and TPM response are often decoupled in the sense that the code used to perform CFD simulations for the vehicle aerothermodynamics (the environment) is separated from the material-response code. The TPS decomposition is then studied by either providing boundary conditions at the fluid-material interface (fully decoupled), or by exchanging the boundary conditions iteratively between the fluid and the material solvers that run separately (weak coupling). Moreover, mass and heat exchange are usually considered as surface phenomena whereas, as mentioned previously, in certain flow conditions and/or for highly-porous ablative materials, the ablation also needs to be considered below the surface (volume ablation).

We distinguish in the common literature three strategies to represent the decomposition of TPMs, which depend on the type of equations that are solved and the boundary conditions that are applied at the flow-material interface. The strategy to solve ablation has evolved over time along with the increase of computational power enabling the use of more detailed models. A comprehensive literature review of ablative-material response codes can be found in [Lachaud et al. \[2011a\]](#).

The first strategy (Strategy I) is concerned with the in-depth response of the material and the interaction with the flowfield is specified through boundary conditions, and are referred to as *material-response codes*. For instance, the CMA program was first developed in the 1960s and was solving an energy equation within the material, and mass and energy balances at the surface (Kendall et al. [1968]). Since then, many improvements of such material-response codes have been performed, for instance with the Fully Implicit Ablation and Thermal Response (FIAT) (Chen and Milos [1999]) (fully implicit, thus more stable), TITAN (Chen [2001]) (FIAT's two-dimensional extension), or PATO (Lachaud and Mansour [2014]; Lachaud et al. [2015]). For the latter, Lachaud et al. [2010, 2011b] implemented the volume-averaged mass and energy conservation equations to model pyrolysis-ablation problems within material-response codes. Using this framework, multi-dimensional effects were shown to be important for ablative composite materials: Weng and Martin [2014] simulated pyrolysis gas transport inside charring ablative materials (using the KATS-MR solver). Weng and Martin [2014] and Lachaud et al. [2015] performed simulations with pyrolysis gas flow within the material and showed the importance of taking into account of multidimensional effects. Finally Meurisse et al. [2018] performed a full-scale numerical simulation with PATO.

Strategy II are usually *flow solvers* that are concerned with the simulation of the flowfield whereas the ablation of the surface is specified by means of boundary conditions. Those are generally thermochemical nonequilibrium hypersonic Navier–Stokes solvers like LeMANS (Martin et al. [2012]), GIANTS (Chen and Milos [2005]), the Data-Parallel Line Relaxation (DPLR) code (Wright et al. [2009]), or the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA) code (McNeil and Peter [1996]) for atmospheric entry applications. For solid rocket nozzles applications, see, for instance, Turchi et al. [2013]; Bianchi et al. [2013], in which steady-state ablation approximation is usually assumed.

Both Strategy I and Strategy II implement the material response or the freestream conditions by means of adequate boundary conditions. In order to improve the fidelity of the model employed and to better capture the interaction between the flow and the material, in particular when dealing with transient flows and material response, Strategy I and II are most often coupled together in an iterative way and the boundary conditions are exchanged dynamically (Milos and Chen [2010]; Chen and Gökçen [2014]; Martin and Boyd [2015]; Bianchi et al. [2012]; Schrooyen et al. [2018]).

To understand the complex physics of the interaction of the gas with both the surface and the inner part of the material, Schrooyen [2015] developed a numerical solver that implements the VANS equations in order to have equations that are valid both in the fluid and the material. Thus, this third approach (Strategy III) allows to get rid of the necessity to impose a boundary condition at the interface between the material and the fluid, and further allows to have a time accurate flow-material interaction and to consider volume ablation. It considers the flow and the material in the same domain of computation, and is often referred to as a *unified approach*, or *unified solver*. Conservation laws can be obtained by adjusting continuously the momentum law to fit Darcy's law in the porous domain (Weng and Martin [2017]) or can be derived from the volume-averaged theory (Whitaker [1999]). The derivation of the VANS equations for multi-component reactive flow in porous media was performed by Schrooyen et al. [2016] and will be reviewed in Section 2.3. The solid part in the computational domain (the material) is modeled by means of a variable that tracks the averaged solid density, which continuously decreases from its maximum value inside the material to zero in the fluid region. The main advantages of Strategy III are that the flow-material dynamic is time-accurate and there is no need for a coupling between a CFD solver and a material solver. It therefore provides an accurate way to analyze the importance of the species transport into the boundary layer due to pyrolysis blowing and the heat transfer. In Duzel [2020], a similar approach is used in which the VANS equations are derived and the capabilities of the code are shown on several test cases for ablative materials decomposition.

2.2 Fundamentals of aerothermal flows

From the discussions above, many phenomena happen during an hypersonic entry: different flow regimes are encountered (continuum vs. rarefied flow, or equilibrium vs. nonequilibrium), and different numerical strategies are used to simulate the flow-material interactions. This section focuses first on the modeling of the fluid part, and the next section will address the interaction between the flow and the material using the volume-average theory.

Energy transformation and temperature changes are important for hypersonic flows and some fundamental thermodynamic concepts are reviewed below. References for high-temperature gas dynamics are provided in [Anderson \[2006\]](#), see also [Bird \[1994\]](#); [Ferziger and Kaper \[1972\]](#); [Vincenti and Kruger \[1965\]](#). High-temperature gases and their interaction with ablative thermal protection materials are provided in [Duffa \[2013\]](#).

We focus on flows in chemical nonequilibrium that are in thermal equilibrium, meaning that excitation/deexcitation processes of the energy level are in equilibrium and the distribution of the energy levels for a given species follows a Boltzmann distribution. This is the case when the flow timescale is much larger than the kinetic timescale. For flows in thermal nonequilibrium, the energy level population depart from the Boltzmann distribution and a description of the states of the particles needs to be provided. They are not discussed here, but a review can be found in [Scoggins \[2017\]](#). We will also focus on flows that are in the continuum regime and for which the classical Navier–Stokes equations are valid. See for instance [Bányai et al. \[2013\]](#) for DSMC simulations for atmospheric entry in the rarefied regime or the book of [Iain D. Boyd \[2017\]](#) for high-temperature nonequilibrium gas in a more general framework.

2.2.1 Thermodynamic properties

Internal energy and enthalpy

Let us consider a *pure gas* composed of the same particles, or species, S . The particles have several thermal degrees of freedom (d.o.f) that add up to form the internal energy. For particles such as molecules, they move through space in a certain random motion because of the translational kinetic energy, they can vibrate back and forth along the molecular axis, or they can rotate around an axis. The motion of the electron around the particle defines also its electronic state. The sum of all the translational, rotational, vibrational and electronic energies forms the internal energy. Statistical thermodynamics and quantum mechanics allow to compute thermodynamic properties for systems at equilibrium from the state variables (pressure and temperature). For instance, the specific internal energy (in J kg^{-1}) for diatomic molecules is given by ([Vincenti and Kruger \[1965\]](#))

$$e = \underbrace{\frac{3}{2}RT}_{\text{Translational energy}} + \underbrace{RT}_{\text{Rotational energy}} + \underbrace{\frac{h\nu/k_bT}{\exp(h\nu/k_bT) - 1}RT}_{\text{Vibrational energy}} + \underbrace{e_{\text{el}}}_{\text{Electronic energy}}, \quad (2.6)$$

where $h = 6.626 \times 10^{-34}$ (J s) denotes Planck's constant, ν is the fundamental vibrational frequency of the molecule (s^{-1}), and $R = \mathcal{R}/M$ denotes the specific gas constant. The results for the translational and rotational energies can be understood from classical mechanics and the theorem of equipartition of energy, which states that each thermal degree of freedom contributes $RT/2$ to the specific energy. The translational motion of a diatomic molecule contributes three thermal d.o.f and the rotational motion contributes two d.o.f, hence leading to the two first terms on the right-hand side of Eq. (2.6). The expression for the vibrational and the electronic energies are obtained from quantum considerations. Note that for monoatomic gases, rotational and vibrational energies are absent. The internal specific enthalpy (J kg^{-1}) is a quantity related

to energy that writes

$$h = e + \frac{p}{\rho}, \quad (2.7)$$

where the second term expresses the work done by the gas to occupy its volume (per unit mass).

Consider now a *mixture* of perfect gases composed of N_g different species S_i , $i = 1 \leq N_g$. Extensive thermodynamic properties, such as the mixture energy and enthalpy, are the sum of the pure species thermodynamic properties weighted by their mass fractions. Therefore, mixture energy and enthalpy read as

$$h = \sum_{i=1}^{N_g} Y_i h_i, \quad e = \sum_{i=1}^{N_g} Y_i e_i, \quad (2.8)$$

and h_i and e_i and obtained from Eq. (2.7) and Eq. (2.6), respectively. Y_i is the mass fraction of species i that expresses as

$$Y_i = \frac{m_i}{m} = \frac{\rho_i}{\rho}. \quad (2.9)$$

The pressure of the gas mixture follows Dalton's law of partial pressures and reads

$$p = \sum_{i=1}^{N_g} p_i, \quad (2.10)$$

where p_i is called the *partial pressure* of species i , which is the pressure that would exist in the system if all of the other species were removed at volume V and temperature T constant. For perfect gases, each of the partial pressure follows Eq. (2.1) individually.

The enthalpy from Eq. (2.8) can be defined either for sensible or absolute enthalpy. The sensible enthalpy is defined relative to a formation enthalpy Δh_f^{ref} , defined at a reference temperature which can be the absolute zero or the standard temperature (equal to 298.16 K), while the absolute enthalpy contains both sensible and formation enthalpies, i.e. $h_{\text{abs}} = h_{\text{sens}} + \Delta h_f^{\text{ref}}$. In particular, at equilibrium, sensible enthalpy is obtained from statistical mechanics as in Eq. (2.6), which can also be defined as

$$h_{\text{sens}}(T_1) = \int_{T_{\text{ref}}}^{T_1} c_p(T) dT \quad (2.11)$$

where c_p is the specific heat at constant pressure. One often finds in the literature the relation $c_p = \partial h(T, p = \text{cte}) / \partial T$. Values for Δh_f^{ref} can be found in references tables as NASA SP-3001 (McBride et al. [1963]).

Total (stagnation) conditions

We consider a fluid element passing through a given point at a velocity \mathbf{u} where the local pressure is p . If we were able to decrease isentropically (adiabatic and reversible deceleration) the fluid velocity to rest, the value of P would change accordingly and this value and is called the total pressure. The total energy E and the total enthalpy H are defined in the same manner as

$$E = e + \frac{1}{2} \mathbf{u}^2, \quad H = h + \frac{1}{2} \mathbf{u}^2. \quad (2.12)$$

2.2.2 Homogeneous chemical reactions

The strong shock wave upstream of the vehicle excites the rotational and vibrational states, induces dissociation and even ionization. Thus, chemical reactions play a significant role on

the composition of the mixture, that evolves along the body and alters the thermodynamics properties. Additionally, the absorption and radiation behaviors of mixture component depend on their chemical composition. Accounting for chemistry therefore plays a major role in the accurate description of hypersonic flows.

The homogeneous production term for each species i is given by

$$\dot{\omega}_i^{\text{hom}} = M_i \sum_{k=1}^{N_r} (\nu_{ik}'' - \nu_{ik}') \left(k_{f,k} \prod_{j=1}^{N_s} \tilde{\rho}_j^{\nu_{jk}'} - k_{b,k} \prod_{j=1}^{N_s} \tilde{\rho}_j^{\nu_{jk}''} \right), \quad (2.13)$$

where $\tilde{\rho}_j = \rho_j/M_j$ is the molar density, $k_{f,k}$ and $k_{b,k}$ are the forward and backward reaction rates, $\nu_{i,k}$ are the stoichiometric coefficients and the superscripts ' and '' represent the reactants and products, respectively.

The rate constants are usually determined experimentally. Although methods from kinetic theory exist for their theoretical estimation, they do not provide sufficient accuracy. For most cases in hypersonic applications, the forward reaction rate coefficients are fitted to a modified Arrhenius law which reads

$$k_f(T) = \mathcal{A}T^m \exp \left\{ \frac{-\mathcal{E}}{\mathcal{R}T(t)} \right\}, \quad (2.14)$$

where T^m is an additional pre-exponential temperature factor that is used to improve the accuracy of the fitting. The whole pre-exponential factor $\mathcal{A}T^m$ is interpreted as the frequency of molecular collisions, which depends on temperature, and \mathcal{E} is the energy that must be provided for the reaction to occur. Thus, the exponential term (called the Boltzmann factor) specifies the fraction of collisions with energies greater than the activation energy.

The production rates of species depend on the temperature and reaction mechanisms that are considered in the mixture. For instance, below 9000 K, ionization is likely to be present. Above 9000 K, where ionization of the atoms takes place, reaction mechanisms with electrons and ions have to be included. Tables for kinetic mechanisms and rate constants for high-temperature air can be found in [Bussing and Eberhardt \[1989\]](#).

2.2.3 Navier–Stokes equations for multi-component reactive flows

The Navier–Stokes equations are obtained by expressing the conservation of momentum of a fluid element. Together with the continuity equation (mass conservation of the fluid element) and the conservation of energy, the resulting system of partial differential equations (PDEs) is sufficient to describe the behavior of compressible viscous flow. See, for instance, [Anderson \[1990\]](#) for a description of the equations. When considering further a chemically reacting gas, one must consider additional mechanisms due to the energy diffusion of chemical species. Below, we summarize the full set of Navier–Stokes equations in the case of multi-component reactive flows.

Mass, momentum and total energy conservation equations are written under the general formulation, in terms of conservative variables $\mathbf{U} = (\rho_i, \rho \mathbf{u}, \rho E)^T$

$$\frac{\partial \rho_i}{\partial t} + \nabla \cdot (\rho_i \mathbf{u}) = -\nabla \cdot (\mathbf{J}_i) + \dot{\omega}_i, \quad \forall i \in \{1, N_s\} \quad (2.15)$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u}) = -\nabla p + \nabla \cdot \bar{\bar{\boldsymbol{\tau}}}, \quad (2.16)$$

$$\frac{\partial \rho E}{\partial t} + \nabla \cdot (\rho H \mathbf{u}) = -\nabla \cdot \mathbf{q} + \nabla \cdot (\bar{\bar{\boldsymbol{\tau}}} \cdot \mathbf{u}) + \dot{\omega}_T, \quad (2.17)$$

where N_g is the number of species in the mixture, ρ_i is the species density which is related to mixture density expresses as $\rho = \sum_{i=1}^{N_g} Y_i \rho = \sum_{i=1}^{N_g} \rho_i$, and $\mathbf{u} = (u_1, u_2, u_3)$ is the hydrodynamic

velocity vector. \mathbf{J}_i is the species diffusion flux

$$\mathbf{J}_i = \rho Y_i \mathbf{V}_i = -\rho_i \sum_{j=1}^{N_g} D_{ij} \mathbf{d}_j, \quad (2.18)$$

where \mathbf{V}_i is the diffusion velocity of species i , related to the mass motion velocity of species i through $\mathbf{u}_i = \mathbf{u} + \mathbf{V}_i$, D_{ij} is the diffusion coefficient.

The viscous stress tensor is a force per unit area that causes the deformation of a fluid element due to viscous effects and expresses as

$$\bar{\boldsymbol{\tau}} = \mu \left(\left(\nabla \mathbf{u} + (\nabla \mathbf{u})^T \right) - \frac{2}{3} (\nabla \cdot \mathbf{u}) \mathbf{I} \right), \quad (2.19)$$

where μ is the dynamic viscosity ($\text{kg m}^{-1} \text{s}^{-1}$) and \mathbf{I} is the identity matrix. The off-diagonal terms of $\bar{\boldsymbol{\tau}}$ are the viscous shear stresses that act tangentially to the surface of a fluid element. The second term on the right-hand side is the bulk viscosity under the assumption that the bulk viscosity coefficient is equal to $-2/3\mu$ (Stokes hypothesis). The normal viscous stresses (e.g. τ_{11}) are meaningful if the gradient velocities (e.g. $\partial_x u_1$) are very large, such as across a shock wave, and in which case a viscous-induced force normal to the surface of a fluid element acts to compress or expand the fluid element, hence the link to the divergence of the velocity $\nabla \cdot \mathbf{u}$ of the bulk viscosity.

The heat flux \mathbf{q} for a chemically reacting gas is given by

$$\mathbf{q} = -k \nabla T + \sum_{i=1}^{N_g} \rho_i \mathbf{V}_i h_i + \mathbf{q}_r, \quad (2.20)$$

where k is the mixture thermal conductivity ($\text{W m}^{-1} \text{K}^{-1}$). The first term is the energy flux caused by thermal conduction. The second term is the energy flux caused by the diffusion of all species, that carry their enthalpy h_i . The third term is the energy flux caused by radiation. For the treatment of the coupling between the flow, ablation, and radiation phenomena, see Scoggins [2017].

The last term in Eq. (2.17), $\dot{\omega}_T$, is the energy production coming from the chemical reactions. Its presence in the energy equation depends on whether the enthalpy in Eq. (2.8) is expressed in terms of absolute enthalpy or the sensible enthalpy (see Chapt. 12 of Anderson [2006] for the mathematical details of the derivation of the energy equation in both case). For the latter, an additional source term appears in Eq. (2.20) due to the change in the chemical enthalpy of the species, which is equal to

$$\dot{\omega}_T = - \sum_{i=1}^{N_g} \dot{\omega}_i \Delta h_{f,i}^0. \quad (2.21)$$

In a compact form, the Navier–Stokes equations can be written as

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}^c = \nabla \cdot \mathbf{F}^d + \mathbf{S}, \quad (2.22)$$

where \mathbf{F}^c , \mathbf{F}^d and \mathbf{S} are the convective flux, diffusive flux, and source term, respectively, with

$$\mathbf{F}^c = \begin{pmatrix} \rho_i \mathbf{u} \\ \rho \mathbf{u} \otimes \mathbf{u} + p \mathbf{I} \\ \rho \mathbf{u} H \end{pmatrix}, \quad \mathbf{F}^d = \begin{pmatrix} -\mathbf{J}_i \\ \bar{\boldsymbol{\tau}} \\ \bar{\boldsymbol{\tau}} \cdot \mathbf{u} - \mathbf{q} \end{pmatrix}, \quad \mathbf{S} = \begin{pmatrix} \dot{\omega}_i \\ \mathbf{0} \\ \dot{\omega}_T \end{pmatrix}. \quad (2.23)$$

Many practical 3D applications can be expressed as 2D axisymmetric problems. For example, sample tests in the Plasmatron or arc-jet experiments are usually assumed to be axisymmetric.

In these cases, the flow is usually described in cylindrical coordinates (Magin [2011]; Lani [2008]). Axisymmetry implies that the derivatives in the azimuthal coordinate θ cancel, i.e. $\partial(\cdot)/\partial\theta = 0$, and azimuthal velocity u_θ is also neglected. In Argo, the 2D axisymmetric problem is expressed similarly to Eq. 2.22 and the source term is adapted in order to account for azimuthal components

$$\frac{\partial \mathbf{U}}{\partial t} + \left(\frac{\partial \mathbf{F}_x^c}{\partial x} + \frac{\partial \mathbf{F}_y^c}{\partial y} \right) = \left(\frac{\partial \mathbf{F}_x^d}{\partial x} + \frac{\partial \mathbf{F}_y^d}{\partial y} \right) + \mathbf{S} + \mathbf{S}^*, \quad (2.24)$$

where x and y stands now respectively for axial and radial directions, and \mathbf{S}^* is an additional source term that accounts for 3D effects, see Schrooyen [2015] for more details. In PATO, the system of equations is solved in the full 3D Cartesian coordinates.

Finally, we note that the number of mass conservation equations can be reduced by assuming local thermodynamic equilibrium (LTE). In this case, it is assumed that the characteristic time scales of the flow match or fall below the kinetic times, hence collisions and chemical reactions have the time to go to completion locally and at each location in the flowfield. For a mixture, the thermodynamic properties are functions of temperature, pressure and the elemental composition of the gas. Thus, it is possible to consider the continuity equations only for the elements rather than for all the species, and to compute the local equilibrium gas compositions. The equilibrium species composition can be computed, for instance, by minimizing the Gibbs free energy (Scoggins and Magin [2015]). Denoting \mathcal{U}_e the set of constitutive element of the gas mixture, the conservation equations for the element mass fractions Y_i^e , with $\rho_i^e = \rho Y_i^e$, reads as

$$\frac{\partial}{\partial t} (\rho_i^e) + \nabla \cdot (\rho_i^e \mathbf{u}) = -\nabla \cdot \mathbf{J}_i^e, \quad \forall i \in \mathcal{U}_e. \quad (2.25)$$

We note that the source term disappears, as a gas species element is conserved through homogeneous reactions.

All the thermodynamic properties, chemical reaction rates, as well as equilibrium composition are implemented and can be computed in the `Mutation++` library (Scoggins et al. [2020]). The LTE assumption is not implemented in Argo and only nonequilibrium chemistry is considered, but it is implemented in PATO.

2.3 Multi-component reactive flows in porous medium

The previous Navier–Stokes equations describe the evolution of a fluid element within a single phase. The presence of a solid surface can eventually be imposed through appropriate boundary conditions. However, it would be impracticable to solve the resulting system of equations within the porous medium over the whole domain. Instead, equations that are valid at the macroscopic scale both in the pure fluid and in the porous medium are derived. The methodology is based on the volume averaging theory applied to the Navier–Stokes equations (Whitaker [1999]). An important consequence of the spatial smoothing is that the boundary conditions due to the presence of the solid surface are joined to the governing equations.

The results of this section are based on the volume-averaged equations for multi-component Navier–Stokes equations derived in Schrooyen [2015] for a unified flow-material approach. The flow inside the pores is assumed to be in the continuum regime (low Knudsen number) and the Reynolds number based on the length of a pore is assumed to be small enough for the inertial term to be negligible. For a treatment of the Knudsen effects in the slip regime and the inertial effects at the pore-scale, see the recent work of Valdés-Parada and Lasseux [2021]. We recall here briefly the concept of the averaging operator and we present the full set of VANS equations derived in Schrooyen [2015].




Figure 2.3: Porous medium with a single solid phase with cylindrical-shaped fibers.

2.3.1 Volume-averaged Navier–Stokes equations

The volume averaging approach is used to derive continuum equations for multi-phase systems in the presence of discontinuities between the phases. Each phase is described by field quantities which are continuous within this specific medium but discontinuous in the whole volume. Solving the equations for each phase separately requires to account for the microstructure of the porous medium which can be quite complex and computationally expensive. Using the volume averaging theory, the equations describing one phase are averaged over a small elementary volume (see Fig. 2.3) such that the volume-averaged quantities are continuous in space. This results in a system of equations that are valid in the entire domain, either within the porous medium or in the plain-fluid region.

Let us consider that the porous medium is made up of a single solid phase, that are the carbon fibers, as sketched in Fig. 2.3. The intrinsic average of a quantity α on a general phase γ is defined as

$$\langle \alpha \rangle_\gamma = \frac{1}{dV_\gamma} \int_{dV_\gamma} \alpha \, dV, \quad (2.26)$$

and has to be distinguished from the superficial average

$$\langle \alpha \rangle = \epsilon_g \langle \alpha \rangle_\gamma = \frac{1}{dV} \int_{dV_\gamma} \alpha \, dV, \quad (2.27)$$

where dV is a small elementary volume and the porosity ϵ_g , defined as the ratio of the volume of fluid phase dV_g and the volume dV , can be computed as $\epsilon_g = 1 - \epsilon_s$. The subscript “s” is used to denote the solid phase and the subscript “g” the fluid (gas) phase.

The volume averaging theory then consists in applying the averaging operator defined in Eq. (2.27) to the whole set of Navier–Stokes Eqs. (2.15)–(2.17). The application of the volume-averaging operator to the system of PDEs leads to the calculation of volume-averaged derivatives, that are related to the derivatives of volume-averaged quantities through closure terms. Analogous to the temporal decomposition used in the study of transport phenomena, closure terms appeared due to point quantities that are decomposed as the sum of a mean and, here, a spatial deviation part, the integration of which needs to be modeled. For more details about the volume-averaging theory, see Whitaker [1999] and for the derivation of the VANS equations for multi-component reactive flows through porous media, see Schrooyen [2015] and references therein.

The evolution of the averaged variables describe the evolution of the macroscopic flow through the porous medium. The conservative variables used to describe an ablative material made of carbon preform are the averaged mass density of every species i in the mixture ($\epsilon_g \langle \rho_i \rangle_g$), the averaged momentum ($\langle \rho \mathbf{u} \rangle_g$), the total energy ($\langle \rho E_{\text{tot}} \rangle$) and the averaged solid density of the

fibers ($\langle \rho_s \rangle$). The VANS system of equations written in vectorial form is

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}^c = \nabla \cdot \mathbf{F}^d + \mathbf{S}, \quad (2.28)$$

with \mathbf{U} the vector containing the conservative variables

$$\mathbf{U} = \begin{pmatrix} \epsilon_g \langle \rho_i \rangle_g \\ \langle \rho \mathbf{u} \rangle_g \\ \langle \rho E_{\text{tot}} \rangle_g \\ \langle \rho_s \rangle \end{pmatrix}. \quad (2.29)$$

The expression for the convective and diffusive fluxes are

$$\mathbf{F}^c = \begin{pmatrix} \epsilon_g \langle \rho_i \rangle_g \langle \mathbf{u} \rangle_g \\ \langle \rho \rangle_g \langle \mathbf{u} \rangle_g \langle \mathbf{u} \rangle_g + \langle p \rangle_g \\ \epsilon_g \langle \rho \rangle_g \langle \mathbf{u} \rangle_g \langle H \rangle_g \\ \mathbf{0} \end{pmatrix}, \text{ and } \mathbf{F}^d = \begin{pmatrix} -\langle \mathbf{J}_i \rangle \\ \langle \boldsymbol{\tau} \rangle_g \\ \langle \boldsymbol{\tau} \cdot \mathbf{u} \rangle + K_{\text{eff}} \nabla \langle T \rangle - \sum_{i=1}^{N_s} h_i \langle \mathbf{J}_i \rangle \\ \mathbf{0} \end{pmatrix}, \quad (2.30)$$

where the averaged mass diffusive flux is

$$\langle \mathbf{J}_i \rangle = -\epsilon_g \langle \rho_i \rangle_g \frac{D_{i,m}}{\eta} \frac{W_i}{W} \nabla X_i + \epsilon_g \langle \rho_i \rangle_g \sum_{k=1}^{N_s} \frac{D_{k,m}}{\eta} \frac{W_k}{W} \nabla X_k, \quad (2.31)$$

where $D_{i,m}$ (and $D_{k,m}$) is the diffusion coefficient of the i -th species in the mixture. The quantity η is called the *tortuosity* which is a geometric factor that characterizes the ratio between the real distance covered by a particle between two points in the porous medium and a straight line.

The effective conductivity K_{eff} is a second-order tensor (matrix), in which off-diagonal terms are zero if the material microstructure is aligned with the principal directions of the physical property described by the tensor (orthotropic material). Anisotropy in the material thermal conductivity leads to diagonal terms that have different values. For isotropic materials, the material conductivity can be modeled as a single scalar quantity as

$$k_{\text{eff}} = \epsilon_s k_s + \epsilon_g k_g, \quad (2.32)$$

that is, the sum of the intrinsic conductivity of each phase weighted by their respective volume fraction. Isotropic material properties will be assumed in the following, but for carbon preform material, it was shown that while transverse effective conductivity was nearly isotropic (x-y plane), the through thickness (axial) direction was favored (Panerai et al. [2017]).

Finally, the source term in Eq. (2.28) expresses as

$$\mathbf{S} = \begin{pmatrix} \langle \dot{\omega}_i^{\text{het}}(\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle + \epsilon_g \dot{\omega}_i^{\text{hom}}(\langle T \rangle, \langle \rho_i \rangle_g) \\ \mathbf{F}_{\text{gs}} \\ - \sum_{i=1}^{N_s+1} (\langle \dot{\omega}_i^{\text{het}}(\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle + \epsilon_g \dot{\omega}_i^{\text{hom}}(\langle T \rangle, \langle \rho_i \rangle_g)) h_{f,i}^0 \\ \langle \dot{\omega}_i^{\text{het}}(\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle \end{pmatrix}, \quad (2.33)$$

where \mathbf{F}_{gs} is the momentum exchange term between the two phases. It is assumed that the closure of this term only includes the first order drag force and reads

$$\mathbf{F}_{\text{gs}} = -\mu \epsilon_g^2 \bar{\boldsymbol{\kappa}}^{-1} \langle \mathbf{u} \rangle_g, \quad (2.34)$$

where $\bar{\boldsymbol{\kappa}}$ is a second-order tensor that represents the permeability of the medium which measures the ability of a fluid to flow through the porous medium. The homogeneous production term $\dot{\omega}_i^{\text{hom}}$ is given by Eq. (2.13) where the molar densities are replaced by the averaged molar densities $\langle \bar{\rho}_i \rangle = \langle \rho_i \rangle / M_j$

Assuming LTE, the conservation equations for the element mass fractions Y_i^e , with $\langle \rho_i^e \rangle_g = \langle \rho \rangle_g Y_i^e$, read as

$$\frac{\partial}{\partial t} (\epsilon_g \langle \rho_i^e \rangle_g) + \nabla \cdot (\epsilon_g \langle \rho_i^e \rangle_g \langle \mathbf{u} \rangle_g) = -\nabla \cdot \langle \mathbf{J}_i^e \rangle + \langle \dot{\omega}_i^{\text{het},e} \rangle. \quad (2.35)$$

Again, there is no source term coming from the homogeneous chemical production terms, as a gas species element is conserved through these reactions. The last terms represent the elemental production stemming from heterogeneous chemical reactions, that are usually not accounted for in state-of-the-art material response codes (Milos and Chen [2010]; Lachaud et al. [2015]). However, if volume ablation is considered, the last term will be non-zero for carbon production.

2.3.2 Heterogeneous chemical reactions

The averaged heterogeneous production rate in Eq. (2.33) for a reacting species is modeled as

$$\langle \dot{\omega}_i^{\text{het}} \rangle = \frac{1}{dV} \oint_{\partial\Omega_g} -k_f^{i,\text{C(s)}} \langle \rho_i \rangle_{\text{gs}} dS, \quad (2.36)$$

where $k_f^{i,\text{C(s)}}$ denotes the reactivity of the carbon fibers for the reaction with the i -th species, and $\langle \rho_i \rangle_{\text{gs}}$ the intrinsic density of a reactant i at the interface. Assuming constant averaged density and constant reactivity over the interface, one can write

$$\langle \dot{\omega}_i^{\text{het}} \rangle = -S_f k_f^{i,\text{C(s)}} \langle \rho_i \rangle_g, \quad (2.37)$$

where $S_f = A_w/dV$ is the specific surface with A_w the contact surface of the solid with the fluid. In order to compute the heterogeneous reaction term, a model for S_f needs to be specified, which depends on the shape of the fibers and the evolution of their microstructure during the ablation process. Assuming cylindrical-shaped fibers (Fig. 2.3) and a uniform radial recession, the specific surface can be expressed as

$$S_f = \frac{2}{r_{f,0}} \sqrt{\epsilon_{s,0} \epsilon_s}, \quad (2.38)$$

where $r_{f,0}$ denotes the initial radius of the fibers and $\epsilon_{s,0}$ the initial volume fraction of the solid phase, which are the fibers. Other models, not investigated here, were studied in Schrooyen et al. [2016].

It is trivial to observe that if $\epsilon_g \rightarrow 1$ as it is the case in the unified approach when we progressively go from the porous medium to the plain-fluid region, then $\epsilon_s = 1 - \epsilon_g \rightarrow 0$, the heterogeneous source term disappears and the system of equations reduces to the classical Navier–Stokes equations for multi-species flows.

2.3.3 Momentum equation for material-response codes

When only the porous material is considered within the computational domain, the volume-averaged momentum equation can be further simplified to simulate the flow within the porous material. Assuming that the flow is incompressible, stationary, and that the Reynolds number is low, i.e. $\text{Re} \ll 1$ (Stokes flow), the VANS equations simplify to the Darcy’s law. Often, material-response codes consider a corrected version of Darcy’s law

$$\langle \mathbf{u} \rangle = -\frac{1}{\epsilon_g} \left(\frac{\bar{\bar{\kappa}}}{\mu} + \frac{\bar{\bar{\beta}}}{\langle p \rangle} \right) \cdot \nabla \langle p \rangle, \quad (2.39)$$

where the additional term $\bar{\bar{\beta}}$ is the Klinkenberg correction that is used to account for slip effects when the Knudsen number (at the pore scale) is in the transition regime (Lachaud et al. [2017]).

In the case where the Reynolds number is high, inertial effects within the pores can be accounted for by using the Forchheimer correction in Darcy's law ([Martin and Boyd \[2010\]](#)). Both effects are neglected in this work, and, at the pore scale, the Knudsen number is assumed to be in the continuum regime and the Reynolds number to be sufficiently low.

2.3.4 Plain fluid-porous media interface modeling

In the case where only the fluid around the TPS is simulated, there is no need to use the VANS equations and the full set of Navier–Stokes equations presented in Section 2.2.3 can be used. The presence of the TPS is communicated by means of specified boundary conditions accounting for catalytic recombination, convective and radiative heat exchanges, or the injection of ablation products in the boundary layer, just to name a few.

In a unified approach where the plain fluid and the porous medium are considered in the same domain of computation, the transition between the two regions requires a special treatment. The physical variables encounter strong but continuous variations from the plain fluid region to the porous domain. Discontinuities are smoothed over a transition region by a normalized hyperbolic tangent

$$H(\phi) = \frac{1}{2} \left(1 + \tanh \left(\frac{2\pi\phi}{\epsilon_\phi} \right) \right), \quad (2.40)$$

where ϕ is the initial distance from the porous medium and ϵ_ϕ is the grid distance on which the quantity is smoothed, which can be seen as the initial interface thickness. This continuum description is flexible enough that it can go smoothly from a plain fluid region to a receding porous medium. The material porosity is itself a variable smoothed on the grid during initialization and then computed to track the evolution of the reactive porous medium. In order to avoid spurious oscillations of the physical quantities at the interface, numerical tests have shown that the smoothing distance has to satisfy the following relation ([Pochet et al. \[2013\]](#)):

$$\epsilon_\phi > 3 \frac{h}{p}, \quad (2.41)$$

where h is the mesh size and p is the interpolation order.

In the case where only the material response is computed, the full set of VANS equations or the corrected Darcy's law from Eq. (2.39) can be used. In this case, the interface between the plain fluid and the solid is treated as a boundary condition. For ablative materials, the surface further recedes and needs to be accounted for. For instance, in the PATO software, the mesh is moving with the interface as the material is receding. The velocity of the moving interface is computed from the ablation rate and the density of the material at the surface ([Meurisse et al. \[2018\]](#)).

2.4 Conclusions

In this chapter, we introduce the modeling of the main physical phenomena during atmospheric entry that are accounted for in this work. We presented the Navier–Stokes equations and the volume-averaged theory used to simulate flow within porous media, and we presented the different numerical strategies used for simulating such flows. We emphasized the features of the two codes that will be used in this work, namely Argo and PATO, for which a comparison of the capabilities can be found in [Schrooyen \[2015\]](#) (in appendix). In Chapter 4, the unified approach implemented within the Argo code, that simulates both the external aertothermal flow and the material response, will be used to model the pyrolysis and char decomposition of a TPM. In Chapter 8 and Chapter 9 the material-response code PATO will be used, thus allowing faster computations for uncertainty quantification.

Pyrolysis models for the decomposition of thermal protection materials

“Essentially, all models are wrong, but some are useful.”

G. E. P. Box, 1987

After a brief review of experimental techniques for the characterization of pyrolysis decomposition, we describe the multi-component and competitive mechanisms for pyrolysis modeling. For multi-component models, we derive a computationally efficient semi-analytical solution to a one-parameter Šesták-Berggren kinetic model and its derivatives. The solution uses exponential-integral functions to express accurately the solution of the temperature integral both for constant temperature (isothermal) and for constant heating rate conditions. We show the interest of this solution on a six-equation model using parameters inferred in [Torres-Herrador et al. \[2019b\]](#) for the PICA material, and we show the computational gain and accuracy of the implemented analytical solution with a numerical solution. When more complex mechanisms are present in the pyrolysis reaction, multi-component models can fail to accurately simulate the material decomposition. In this case, competitive models are preferred, but for which no analytical solution was found. The solution to the multi-component models was presented in [Coheur et al. \[2021b\]](#) and part of the current chapter is based on this manuscript. Competitive reaction mechanisms applied to TPMs was presented in [Torres-Herrador et al. \[2019a\]](#).

3.1 Introduction

The literature on pyrolysis of organic compounds is vast ([Galwey and Brown \[1999\]](#); [Di Blasi \[2008\]](#); [Vyazovkin et al. \[2011\]](#)), and many different models have been proposed to simulate this process. A widely used formulation for the pyrolysis of TPMs is the multi-component (or multi-step) pyrolysis model ([Martin and Boyd \[2015\]](#); [Lachaud et al. \[2017\]](#); [Schrooyen et al. \[2016\]](#); [Coheur et al. \[2017\]](#); [Torres-Herrador et al. \[2019b, 2020\]](#)). When the material is heated, it is decomposed following independent reaction processes that occur in parallel, each of them representing the decomposition of a sub-phase, or component, in the solid. The use of multi-component models for pyrolysis is not limited to the application to TPMs. They are also widely used in fire safety ([Rein et al. \[2006\]](#); [Lautenberger and Fernandez-Pello \[2011\]](#); [Kolaitis and Founti \[2013\]](#); [Nguyen et al. \[2016\]](#); [Bruns \[2015\]](#); [Yuen et al. \[2018, 2020\]](#); [Zhang et al. \[2021\]](#)) or (less commonly) in biomass ([Torres-Herrador et al. \[2020\]](#)).

The reaction rate characterizing the reaction processes are usually parameterized in terms of

temperature and the extent of conversion (or advancement of reaction factor). The dependence on temperature, called the rate constant, is typically parameterized through the Arrhenius equation. The dependence on the extent of conversion is parameterized using a model reaction that depends on the type of the reaction (Galwey and Brown [1999]). In particular, the majority of model reactions can be generalized using the Šesták-Berggren equation (Šesták and Berggren [1971]; Málek and Criado [1992]; Brown et al. [2000]). One of the objectives of kinetic studies is to determine the kinetic parameters appearing in the parameterization of the rate constant and the model reaction.

Usually, multi-component parallel models are valid for a restricted range of heating rate conditions around which the models are calibrated. When considering highly different heating rates, the reaction mechanisms may vary and multi-component models may fail to accurately reproduce the physics. As the heating rate increases, the decomposition of the material can occur at higher temperatures due to thermal lag effects, which results in a shift towards higher temperatures of the mass loss curve and the pyrolysis production peak (Torres-Herrador et al. [2019b]). However, it was observed that for carbon/phenolic materials, the mass loss curve and the production peak shift towards lower temperatures (Jackson and Conley [1964]; Stokes [1995]; Bessire and Minton [2017]). In order to take into account this shift towards lower temperatures, pyrolysis models that include competitive mechanisms are therefore considered (Miller and Bellan [1997]; Blondeau and Jeanmart [2012]; Herrador et al. [2019]). The incorporation of competitive mechanisms in pyrolysis models allows for an initial solid reactant to decompose into different solid products in proportions that depend on the heating rate. Those solid products may become in turn reactants, resulting in different reaction paths that are naturally included in a single pyrolysis model. Competitive models can be seen as a generalization of the multi-component (parallel) reaction models that take into account competitive mechanisms and the dependence between the different pyrolyzing components. The use of competitive mechanisms for the pyrolysis of TPMs was initially proposed in Herrador et al. [2019] based on the work of Blondeau and Jeanmart [2012] for biomass pyrolysis.

Solving inverse problems to infer parameters of the kinetic model through optimization techniques or Bayesian inference methods for uncertainty quantification may require a large number of evaluations of the model and its sensitivities (derivatives with respect to parameters). Moreover, in the case of kinetic equations, the Arrhenius parameters can exhibit strong dependence that can require further model evaluations for an accurate parameter calibration. Therefore, for multi-component models, we propose a computationally efficient semi-analytical solution to a one-parameter Šesták-Berggren kinetic model and its derivatives that uses exponential-integral functions to express the solution of the temperature integral both for constant temperature (isothermal) and for constant heating rate conditions. The interest of this analytical solution is to reduce the computational cost while having high accuracy to perform parameter calibration from experiments and sensitivity analysis. Although competitive models allow to generalize the previous multi-component models, we do not know an analytical solution in this case.

This chapter is organized as follows. In Section 3.2, we provide an overview of experimental characterizations of the pyrolysis of low-density carbon/phenolic thermal protection materials. Then, in Section 3.3, we review the multi-component models and we derive the semi-analytical solution, while Section 3.4 is dedicated to the description of competitive mechanisms. Finally in Section 3.5, we provide an example of a general competitive reaction, an example that illustrates the advantage of the semi-analytical solution using a six-component model for the decomposition of PICA, and we finally present the two-equation pyrolysis model from Goldstein [1969] that is implemented in many state-of-the-art numerical codes. We further refer to Chapter 7 for practical applications of multi-component and competitive models to the pyrolysis of TPMs.

3.2 Experimental characterization of pyrolysis

We focus on experiments that have been carried out in order to characterize the thermal decomposition of the resin component of low-density carbon fiber/phenolic resin thermal protection materials. Phenolic resins are generally formed through a polycondensation reaction between phenol ($\text{C}_6\text{H}_5\text{OH}$) and formaldehyde (CH_2O) in the presence of a catalyst. Resin samples are then typically heated up in a furnace under controlled temperature and pressure environments. The mass of the sample and the species produced during the resin decomposition are measured as a function of time. The variation of the mass is usually monitored through TGA and the composition of the species produced is analyzed through mass spectrometry (MS) or gas chromatography (GC). Techniques such as differential thermal analysis (DTA) or differential scanning calorimetry provides information on the type of the pyrolysis reactions, either endothermic or exothermic, and allow to determine the enthalpy change due to the reactions.

Depending on the measurement techniques and the experimental set-up, the sample is usually exposed either to an isothermal temperature or to an increasing function of temperature. The former consists in heating the material up to a target temperature in a very short time and maintaining this temperature during a given time interval in which pyrolysis reactions are happening. The latter consists in increasing the temperature as a function of time (usually linearly) while pyrolysis is occurring until the target temperature is reached.

The actual heating observed during a reentry is neither a constant value nor exactly a linear function. It is a complex combination of the incoming heat from radiation in the shock layer, convection from the boundary layer, species recombination at the surface of the material and conduction within the material. Moreover, the heating rates vary dramatically spatially across the heat shield, both span-wise and in-depth, and in time during the entry trajectory. In practice, the exact same heating is never reproduced during ground experiments. Nevertheless, experiments under a controlled environment already provide an understanding of the reaction mechanisms occurring during the pyrolysis and the model developed from these experiments can still be used in numerical solvers.

3.2.1 Isothermal characterization

Sykes [1967] experiments were carried out during a large test campaign at NASA Langley Research center. The objective was to provide a comprehensive analysis of the decomposition mechanisms of ablative materials. The experimental data provided by Sykes [1967] are still largely used in the aerospace community for the determination of the pyrolysis gas composition.

In these experiments, a phenolic specimen of 7 mg is heated in the entrance port of a gas chromatograph that is used to collect the pyrolysis gases produced during a given time interval at a fixed temperature and to analyze the gas composition¹. The test temperature is reached over a 10-second interval and the pyrolysis of the specimen was conducted from 373.15 to 1273.15 K.

The measured species from the GC are shown in Fig. 3.1. Below 800 K, water is the most dominant species and is the result of thermal breakdown at the -OH groups attached to the ring structure of the polymer. The decomposition of the polymer continues up to 1000 K, where the major product of the reaction is hydrogen. The type of reaction was analyzed by means DTA² and a measure for the formation enthalpy was also obtained.

In Wong et al. [2015a] and Wong et al. [2015b], the production of pyrolysis gases are quantified using a GC technique similar to Sykes [1967]. A phenol-formaldehyde resin is heated in a reactor following a step-wise procedure with a temperature increment of 50 K from 320 K to 1290 K.

¹Gas chromatography determines the type of species by measuring its retention time in the chromatograph which amount is proportional to an electric signal provided by a thermal conductivity detector

²DTA measures the temperature difference between an inert material and the decomposing material under investigation, which value indicates whether the reaction is endothermic or exothermic

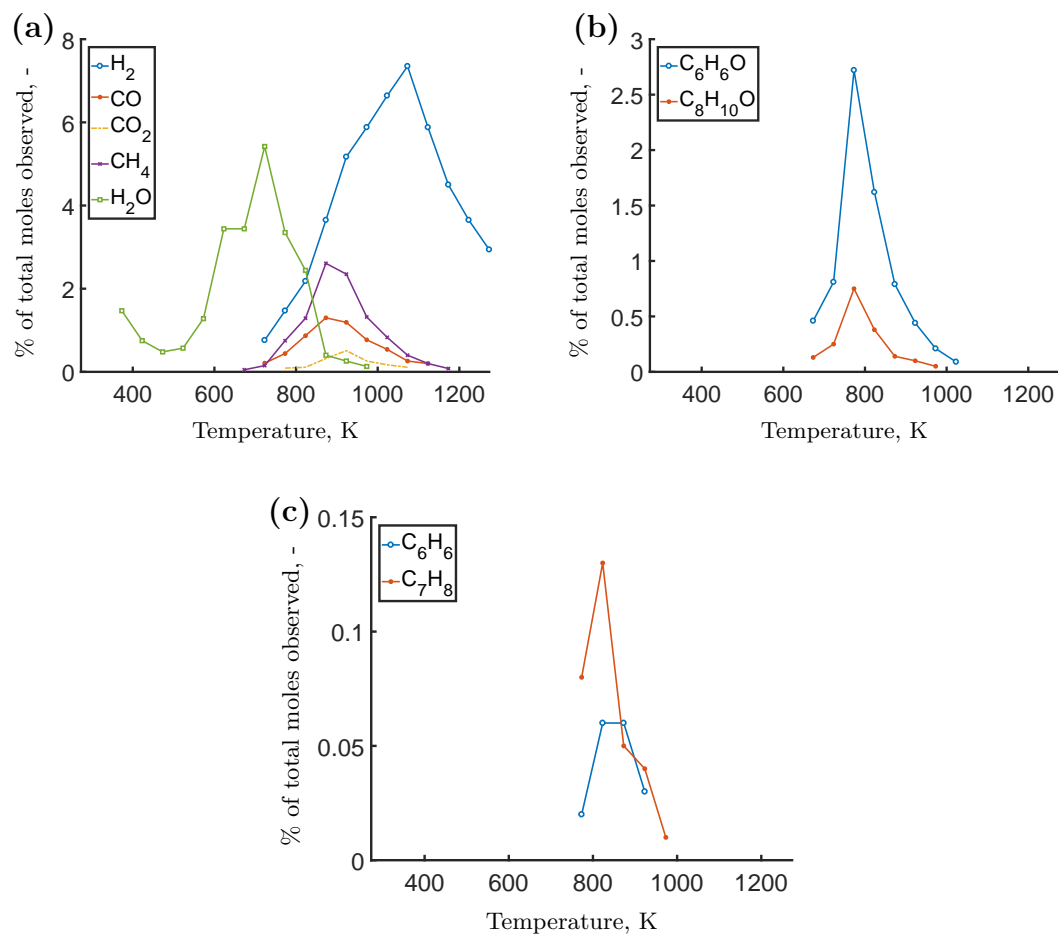


Figure 3.1: GC analyses from Sykes [1967] experiments. The species are classified in three families: (a) permanent gases and water; (b) phenol and derivatives and (c) aromatics.

At each step, the reactor is kept at the target temperature for 1 h in order to ensure that pyrolysis reaction at this temperature is near completion. The amount of pyrolysis gas products is collected during each interval. This allows the authors to identify the different gas products and quantify the mass of every species with minimal loss, but leaves time for secondary reactions to occur in the gas collector. An electronic balance is also used to weigh the reactor at the end of each temperature increment. The overall experimental procedure was repeated three times and a mean and a standard deviation for the measured quantities at each temperature were estimated.

Results of the GC analysis for the pyrolysis gas production are shown in Fig. 3.2(b) for permanent gases and water, and similar trends as in Sykes [1967] were observed. Finally, the experiment was also performed using a linear increasing temperature and the sample mass was recorded using TGA, as shown in Fig. 3.2(a).

Bessire et al. [2014] provided in-situ measurements of PICA pyrolysis product yields using a mass spectrometer in the range 375.15 to 1223.15 K. A PICA sample is heated in a vacuum chamber by means of an electrical current that passes through the sample. A target temperature is reached within a matter of seconds and the temperature is maintained at its value during a hold time of 15–20 min. To avoid saturation of the mass spectrometer due to the rapid increase

of pyrolysis gases after a temperature increment, the valve of the mass spectrometer is closed during the initial moments of each temperature increment. Data for a particular temperature are then collected during a 2.5 min period in the range between 12–15 min of each temperature step. Then the material is heated to the next temperature step with an increase of roughly 50 K until 1223.15 K.

The authors observed that H_2O is the most dominant species during a first decomposition process. Then CO and CH_4 species production becomes significant in a second stage and finally H_2 is the most dominant species in the last stage.




Figure 3.2: Sample mass (a) and mass yields for permanent gases and water (b) as a function of temperature from the experiments of Wong et al. [2015b] for a phenol-formaldehyde resin pyrolysis. The markers are the estimated means and the error bars represent the empirical standard deviations.

3.2.2 Linearly increasing temperature

Goldstein [1969] considered nylon 6-6, phenolic CTL-91LD (a phenolic novolac resin), and their composite. TGA is used to characterize the weight loss, but does not allow to identify the species produced. However, the kinetic parameters were obtained from a fitting procedure against the experimental results at heating rates of 3 and 18 K min^{-1} . For the pyrolysis of the phenolic resin, a two-step reaction mechanisms was found (two maximum reaction peaks were observed). Although nowadays materials are slightly different from those being used at that time, these kinetic parameters are still used in numerical models. In particular, charring mass fraction was observed to be 0.5 but recent experiments on PICA (Wong et al. [2016]; Bessire and Minton [2017]) show that this fraction could be higher.

Trick and Saliba [1995] identified three mechanisms occurring during the thermal degradation of a commercially available carbon/phenolic by means of Fourier Transform Infrared Spectroscopy, corresponding to the formation of additional crosslinks, breaking of the crosslinks and stripping of the aromatic rings. Trick et al. [1997] then proposed a kinetic model using isoconversional methods on TGA curves that consists in four dominant heterogeneous reactions occurring in overlapping temperature regions.

Bessire and Minton [2017] performed a second study in which the temperature is a linear function of time for different heating rates. The pyrolysis of a 0.441 g PICA sample was performed at significantly different heating rate values, namely 186, 366, 762 and 1500 K min^{-1} . In these experiments, the PICA sample is heated from 375 to 1400 K in a vacuum chamber by

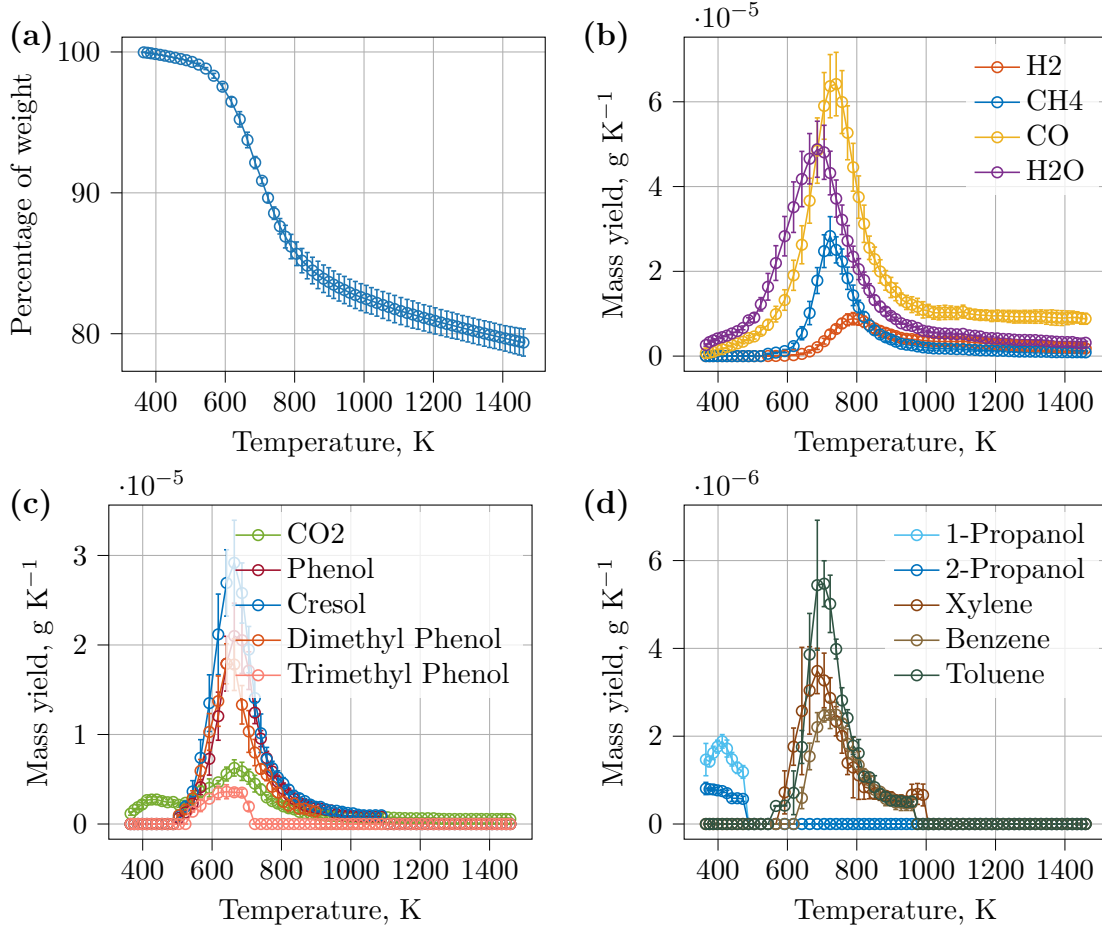


Figure 3.3: Total mass loss of the sample (a) and the 14 mass yields at a heating rate of 366 K min^{-1} from the experiments of Bessire and Minton [2017]. The markers are the estimated means and the error bars represent the empirical standard deviations. The mass yields shown are for (b) permanent gases and water, (c) phenol and derivatives, and CO₂, and (d) alcohol and aromatics.

means of an electrical current that passes through the sample. Monitoring the temperature using a thermocouple located in the center of the material and using a control system, the authors were able to achieve heating rates much higher than those of traditional TGA systems. In-situ measurements of PICA pyrolysis product yields are performed using mass spectrometry, which provides the product yields of 14 gaseous species as a function of the material temperature, shown in Fig. 3.3 at the heating rate of 366 K min^{-1} . This included permanent gases, water and hydrocarbons (up to C₈). A TGA-like mass loss versus temperature (or time) curves were reconstructed by integrating the measured gas production rates.

Torres-Herrador et al. [2020] performed TGA on ZURAM (Rothermel et al. [2014]; Reimer et al. [2018]), a carbon/phenolic TPM developed at DLR that has similar properties to PICA. Three different heating rates were considered and a four-step reaction mechanisms was used in order to numerically reproduce the results. The pyrolysis model was selected to be in agreement with results from Trick et al. [1997] and the experiments from Bessire and Minton [2017] and Wong et al. [2015b].

3.2.3 Pyrolysis of PICA


These previous works have permitted to highlight that the pyrolysis of phenolic resins, like in PICA, is typically divided into four overlapping stages, which is also in agreement with other studies on phenolic resins, see for instance [Ouchi and Honda \[1959\]](#). In the first stage, water observed at low temperatures around 400 K is usually attributed to humidity (outgassing) [Bessire and Minton \[2017\]](#). In stage 2, two competing major processes were identified at temperatures in the range 575–785 K: cross-linking reactions releasing H_2O while forming ether and carbon—carbon bonds, and scission of methylene bridges releasing phenol and its methylated derivatives. Stage 3 corresponds to the generation of CO and CH_4 due to reactions that, in the current literature, are not yet clearly understood: [Ouchi and Honda \[1959\]](#) identified, for instance, a two-step mechanism while [Trick and Saliba \[1995\]](#) identified a one-step mechanism. Finally, stage 4 occurs at high temperatures (≥ 1000 K), where H_2 is the major product due to the formation of stable char. The CO observed at high temperatures and other remaining gases still do not have a clear origin.

The two mechanisms in stage 2 are in competition with each other, meaning that the products resulting from these processes are generated in varying proportions depending on the heating rate ([Torres-Herrador et al. \[2019a\]](#)). Hence, a pyrolysis model that accounts for competitive mechanisms should ideally be considered for the decomposition of PICA when the heating rate value is varied significantly.

3.3 Multi-component pyrolysis models

3.3.1 Single-stage kinetic equation


A pyrolysis reaction occurs when an organic compound is submitted to high temperatures in the absence of oxygen. This reaction is typically assumed to be a non-reversible process whose evolution is governed by a reaction rate. In this process, the reactant R gets decomposed into different gaseous products S_j , with $1 \leq j \leq N_g$ and N_g the total number of species, and a residual solid known as char of pure carbon,



where the ζ_j , $1 \leq j \leq N_g$, are the stoichiometric coefficients and the superscripts denote the state of matter, either solid (s) or gaseous (g). This elementary process describes in detail the decomposition of each molecule but they are difficult to resolve with current experimental methods and global reaction schemes are the most common approach in pyrolysis modeling ([Di Blasi \[2008\]](#)). Mass coefficients γ_C and γ_j , $1 \leq j \leq N_g$, are therefore introduced

$$\gamma_C = \frac{\zeta_C M_C}{\sum_j^{N_g} \zeta_j M_j + \zeta_C M_C}, \quad \gamma_j = \frac{\zeta_j M_j}{\sum_j^{N_g} \zeta_j M_j + \zeta_C M_C}, \quad (3.2)$$

and we write the single-stage reaction in terms of mass coefficients as



The latter reaction can either refer to elementary processes or to global reactions.

It is common to describe the amount of decomposition which has taken place using the dimensionless fractional decomposition factor $\alpha(t)$ (or advancement of reaction factor, or extent of reaction, or extent of conversion) as a function of time. Here, we will take this factor as varying from 0 for the virgin reactant to 1 signifying the complete decomposition. This factor is related to the mass $m(t)$ of the solid substance as

$$\alpha(t) = \frac{m_0 - m(t)}{m_0 - m_\infty}, \quad (3.4)$$

where m_0 is the initial mass of the solid and m_∞ is the mass of the solid residue after complete decomposition (the char). Assuming that the pyrolysis reaction Eq. (3.3) reaches completion, we have $m_\infty = m_C = \gamma_C m_0$.

The dependence of the reaction rate on temperature is commonly described by the Arrhenius equation (White et al. [2011])

$$k(T(t)) = \mathcal{A} \exp \left\{ \frac{-\mathcal{E}}{\mathcal{R}T(t)} \right\}, \quad (3.5)$$

where $k(T(t))$ is the constant rate (s^{-1}), \mathcal{R} is the universal gas constant ($\text{J K}^{-1} \text{mol}^{-1}$), \mathcal{A} is the pre-exponential factor (s^{-1}) and \mathcal{E} is the activation energy (J mol^{-1}). In gas (homogeneous) reactions, \mathcal{A} was interpreted as the frequency of molecular collision and \mathcal{E} is the energy that must be provided for the reaction to occur. The use of the Arrhenius equation for solid-state reactions and the actual meaning of these kinetic parameters, contrary to gas reactions, is still subject to debate (White et al. [2011]). A possible interpretation for \mathcal{A} is the frequency factor or specific vibration of the reactant and for \mathcal{E} the energy barrier that must be surmounted during the bond redistribution occurring during the phase change (Galwey and Brown [1999]; L'vov [2007]).

Temperature is not the only factor influencing the rate of the pyrolysis process: it may depend on the fractional decomposition factor α and the pressure of volatile products (in case of gas-solid reaction) p_j . This leads to the following more general formulation for the evolution of the solid

$$\frac{d\alpha}{dt}(t) = k(T(t)) f(\alpha(t)) g(p_j(t)). \quad (3.6)$$

The dependence on the partial pressure is ignored in most kinetic studies (Vyazovkin et al. [2011]). The dependence on the fractional decomposition (the reaction model) can be generically formulated as $f(\alpha) = \alpha^l (1 - \alpha)^r [-\ln(1 - \alpha)]^r$ as proposed by Šesták and Berggren [1971] and Málek and Criado [1992]. Typical aerospace applications assume a reaction-order model, i.e. $l = 0, r = 0$, for the dependence on the advancement reaction: $f(\alpha) = (1 - \alpha)^r$, with r the order of reaction. There are two reasons here for considering this particular form for the reaction model. The first one is because the integration of this single-parameter reaction-order model results in a closed-form solution. For instance, the reaction model $f(\alpha) = \alpha^l (1 - \alpha)^r$, as in Zhang et al. [2021], would result in a solution that involves the incomplete beta function. The second one is the long legacy of previous pyrolysis modeling of TPMs that assumed this kinetic model and which provided satisfactory results.

This finally leads to the following ordinary differential equation (ODE) for pyrolysis with the one-parameter Šesták-Berggren kinetic model

$$\frac{d\alpha}{dt}(t) = \mathcal{A} \exp \left\{ \frac{-\mathcal{E}}{\mathcal{R}T(t)} \right\} (1 - \alpha(t))^r. \quad (3.7)$$


This ODE represents the advancement of the reaction for a point mass sample, which is a good approximation for many thermoanalytical experiments (Torres-Herrador et al. [2020]). It can be extended to more complex cases by considering an energy equation that include the thermal conductivity within the material or the heat capacity of the material, in which case the temperature is a function of space and time and the ordinary differential equation is replaced by a partial differential equation (Blondeau and Jeanmart [2012]; Lachaud et al. [2017]; Coheur et al. [2017]).

Sometimes, the pre-exponential factor is expressed as $\mathcal{A} = \mathcal{A}_0 T^m$, where \mathcal{A}_0 denotes a temperature-independent factor, in order to emphasize a potential dependence on temperature (Flynn [1997]). Previous work on parameter identification for thermal protection materials showed that the parameter m tends to zero (Torres-Herrador et al. [2020]), so in the following

we will set $m = 0$ and omit T^m for simplicity. Therefore, the kinetic triplet (\mathcal{A} , \mathcal{E} and r) is assumed to be independent of temperature and we assume in the following that $\mathcal{A} > 0$, $\mathcal{E} > 0$ and $r \geq 1$.

3.3.2 Multi-component kinetic equations

In most cases, solid state reactions are often too complex to be described using a single reaction (Vyazovkin and Wight [1997]; Brown et al. [2000]). It is therefore convenient to define multiple advancement of reaction factors $0 \leq \alpha_i \leq 1$, where $1 \leq i \leq N_p$ indexes independent pyrolysis reactions occurring in parallel and N_p is the number of pyrolysis decomposition processes present in the solid resin (or the number of reactants, or the number of steps). Each step represents a distinct reactant (or sub-phase (Lachaud et al. [2017]), or pseudo-component (Di Blasi [2008])) in the solid. The corresponding reaction scheme is represented schematically in Eq. (3.8):



where w_i is the relative weight, or mass fraction of compound $R_i^{(s)}$ from the virgin reactant R and $\sum_{i=1}^{N_p} w_i = 1$. The evolution of the N_p advancement of reaction factors given the reaction scheme Eq. (3.8) is therefore

$$\frac{d\alpha_i}{dt}(t) = \mathcal{A}_i \exp \left\{ \frac{-\mathcal{E}_i}{RT(t)} \right\} (1 - \alpha_i(t))^{r_i}, \quad 1 \leq i \leq N_p. \quad (3.9)$$

This expression is widely used in the aerospace community to model thermal protection materials pyrolysis (Goldstein [1969]; Trick et al. [1997]; Torre et al. [1998a,b]; Clayton [2001]; Bhatia and Roy [2010]; Chen and Milos [2013]; Lachaud and Mansour [2014]; Lachaud et al. [2015]; Coheur et al. [2017]; Lachaud et al. [2017]; Torres-Herrador et al. [2019b]). We refer to this process as multi-step, multi-component, or parallel pyrolysis reactions. The advancement of reaction factors are defined for the compounds $R_i^{(s)}$ of mass $m_i(t)$ as

$$\alpha_i(t) = \frac{m_{0,i} - m_i(t)}{m_{0,i} - m_{\infty,i}}, \quad 1 \leq i \leq N_p \quad (3.10)$$

with $m_{0,i} = m_0 w_i$ and we have $\alpha(t) = \sum_{i=1}^{N_p} w_i \alpha_i(t)$. The total mass of the sample can then be obtained by inserting the latter in Eq. (3.4). The parameters w_i are unknown as well as the $m_{\infty,i}$ that can also be unknown if the reaction processes did not reach completion at the end of the temperature program and therefore cannot be measured experimentally. For instance, the weight in Fig. 3.3(a) suggests that the global reaction has not yet reached completion. Considering that $m_{\infty,i} = \gamma_{C,i} m_{0,i}$, the total mass is finally given by

$$m(t) = m_0 \left(1 - \sum_{i=1}^{N_p} \alpha_i(t) w_i (1 - \gamma_{C,i}) \right). \quad (3.11)$$

The term $w_i(1 - \gamma_{C,i}) = w_i \sum_j \gamma_{ij}$ represents the mass fraction of the solid reactant $R_i^{(s)}$ that is converted into gas. The mass after completion of the reaction can be retrieved by computing $m_{\infty} = m_0(1 - \sum_{i=1}^{N_p} w_i(1 - \gamma_{C,i}))$.

Pyrolysis gas production

The advancement of reaction factors α_i cannot be observed directly from experiments; thus, for calibrating model parameters, we need to provide a mathematical expression for the model output that can be linked to the experimental observations. The mass of the sample is linked to the α_i through Eq. (3.11), but the pyrolysis species production allows to observe more accurately the independent reaction steps. Although production rates can also be obtained by differentiating the TGA curve, the differentiation introduces additional numerical noise in the experimental data (Conesa et al. [2001]; Torres-Herrador et al. [2020]).

During pyrolysis, part of the solid material is converted into gas and pyrolysis gas production can be obtained from mass conservation. It is assumed that char and gaseous species are the only products from pyrolysis. Because the char is non-volatile, mass variation of the sample is necessarily linked to the pyrolysis gas production and any sample mass decrease implies pyrolysis gas production. Thus, the pyrolysis gas production rate is given by

$$\pi(t) = -\frac{dm}{dt}(t) = \sum_{i=1}^{N_p} w_i(1 - \gamma_{C,i}) \frac{d\alpha_i}{dt}(t) m_0 := \sum_{i=1}^{N_p} \pi_i(t), \quad (3.12)$$

where $\pi_i(t)$ is the gas production rate coming from the decomposition of the reactant i . Considering that $(1 - \gamma_{C,i}) = \sum_{j=1}^{N_g} \gamma_{ij}$, we can write

$$\pi_i(t) = \sum_{j=1}^{N_g} \underbrace{\gamma_{ij} w_i}_{F_{ij}} \frac{d\alpha_i}{dt}(t) m_0, \quad (3.13)$$

where $F_{ij} := \gamma_{ij} w_i$ is the fraction of mass from the sample that decomposes into gas species j due to reaction i . In Eq. (3.13), the summands represent the amount of species j produced by the reaction i relative to the initial sample mass m_0 . We finally let

$$\nu_j(t) := \sum_{i=1}^{N_p} F_{ij} \frac{d\alpha_i}{dt}(t) m_0, \quad (3.14)$$

$$\text{or } \nu_j(T) := \sum_{i=1}^{N_p} F_{ij} \frac{d\alpha_i}{dT}(T) \frac{dT}{dt} m_0. \quad (3.15)$$

The quantity $\nu_j(T)/(dT/dt)$, measured in g K^{-1} is the quantity observed experimentally in Figs. 3.3(b)–3.3(d).

The introduction of the F_{ij} can help reducing the number of parameters to infer, which will be useful in Chapter 7, providing a few more assumptions that we state now. The value of m_∞ can be computed from the F_{ij} using the relation $m_\infty = m_0 - \sum_{i=1}^{N_p} \sum_{j=1}^{N_g} F_{ij} m_0$. To retrieve $w_i = \sum_{j=1}^{N_g} F_{ij} / (1 - \gamma_{C,i})$, one needs the value of each $\gamma_{C,i}$, which are unknown. One possibility to infer on these values is to assume that each reaction produces char in the same proportions, i.e. $m_\infty/m_0 = \gamma_{C,i}$, $1 \leq i \leq N_p$. Hence, $w_i = \sum_{j=1}^{N_g} F_{ij} / (1 - m_\infty/m_0)$ and finally $\gamma_{ij} = F_{ij}/w_i$.

Instead of assuming a constant char fraction produced by each reaction, the model implemented in PATO and in the `pyrolysis_general` code from VKI considers that the char is produced outside the gasification reactions, i.e. $\gamma_{C,i} = 0$ for $1 \leq i \leq N_p$, $\sum_{j=1}^{N_g} \gamma_{ij} = 1$, and has an additional non-pyrolyzing compound $w_c \text{R}_c^{(s)}$, which is pure carbon, such that $\sum_{i=1}^{N_p} w_i + w_c = 1$. The char is seen as being present initially in the resin. In this case, we have $w_c = 1 - \sum_{i=1}^{N_p} w_i = 1 - \sum_{i=1}^{N_p} \sum_{j=1}^{N_g} F_{ij}$ and finally $w_i = \sum_{j=1}^{N_g} F_{ij}$, $\gamma_{ij} = F_{ij}/w_i$.

3.3.3 Analytical solution of the multi-component model with the one-parameter Šesták-Berggren equation

An analytical solution does not necessarily exist when the temperature program is an arbitrary function of time. However, most experimental studies such as in TGA are carried out by using either a linearly increasing temperature $T = \beta(t - t_0) + T_0$ with β the heating rate in K min^{-1} (Bessire and Minton [2017]; Torres-Herrador et al. [2020]), t_0 the initial time and T_0 the initial temperature, or constant temperature $T = \text{const}$ (isothermal conditions) (Wong et al. [2015b]; Wong et al. [2016]; Hu et al. [2020]). The case of a linearly increasing temperature is widely used in thermal protection materials decomposition studies Goldstein [1969]; Trick and Saliba [1995]; Torre et al. [1998a]; Bessire and Minton [2017]; Torres-Herrador et al. [2020]; Olejnik et al. [2020]. For constant temperature, some experimental set-ups use a piecewise constant temperature program (Wong et al. [2015b]; Wong et al. [2016]), where the isothermal condition is maintained only during a given time interval. In these settings, the temperature $T(t)$ is maintained at a constant value $T^{(k)}$ during a time interval $A^{(k)} = [t_0^{(k)}, t_f^{(k)}]$ in order to collect and analyze the species products before increasing the temperature to its next value $T^{(k+1)}$. The temperature is thus assumed to be the piecewise constant function $T(t) = T^{(k)}$ for all $t \in A^{(k)}$ with $1 \leq k \leq n_{\text{step}}$ and n_{step} the number of temperature increase.

Therefore, we wish to find the analytical solution to the initial-value problem (IVP)

$$\begin{cases} \frac{d\alpha_i}{dt}(t) &= \mathcal{A}_i \exp \left\{ \frac{\mathcal{E}_i}{\mathcal{R}T(t)} \right\} (1 - \alpha_i(t))^{r_i}, & \text{for } t > t_0, \\ \alpha_i(t_0) &= \alpha_{i,0}, & i = 1, \dots, N_p \end{cases} \quad (3.16)$$

with either

$$T(t) = \text{const} \quad (\text{isothermal}) \quad (3.17)$$

or

$$T(t) = \beta(t - t_0) + T_0 \quad (\text{linearly increasing temperature}). \quad (3.18)$$

The system of equations in (3.16) is a system of independent first-order non-linear ODEs with constant coefficients in the case of (3.17) and with non-constant coefficients in the case of (3.18). We start by performing a simple change of variables as follows. Substituting $u_i(t) = 1 - \alpha_i(t)$ leads to

$$\frac{du_i(t)}{dt} = -u_i^{r_i}(t) \mathcal{A}_i \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T(t)} \right\}. \quad (3.19)$$

Dividing both handsides by $u_i^{-r_i}(t)$, assumed nonzero, and integrating both sides of the equation with respect to time, we have

$$-\int_{t_0}^t u_i^{-r_i}(\tilde{t}) \frac{du_i(\tilde{t})}{d\tilde{t}} d\tilde{t} = \int_{t_0}^t \mathcal{A}_i \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T(\tilde{t})} \right\} d\tilde{t}, \quad (3.20)$$

where the right-hand side is the integral of the Arrhenius equation and is named the temperature integral

$$I_i(t) = \int_{t_0}^t \mathcal{A}_i \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T(\tilde{t})} \right\} d\tilde{t}. \quad (3.21)$$

The case where $r_i = 1$ is left for the end of the section. Assuming that $r_i \neq 1$, the left-hand side can be integrated by substitution leading to

$$-\int_{u_i(t_0)}^{u_i(t)} \tilde{u}^{-r_i} d\tilde{u} = -\frac{1}{1 - r_i} (u_i^{1-r_i}(t) - u_i^{1-r_i}(t_0)). \quad (3.22)$$

Inserting this expression in Eq. (3.20) and finally expressing the result with the advancement of reaction factor, we obtain

$$u_i(t) = (u_i^{1-r_i}(t_0) - I_i(t)(1 - r_i))^{1/(1-r_i)}, \quad (3.23)$$

$$\alpha_i(t) = 1 - \left((1 - \alpha_i(t_0))^{1-r_i} - I_i(t)(1 - r_i) \right)^{1/(1-r_i)}. \quad (3.24)$$

In this last equation, we still need to provide an expression for the temperature integral $I_i(t)$, that depends on time through the temperature program.

Case 1: piecewise isothermal

In isothermal pyrolysis experiments, we assume the temperature to be constant on a given time interval $A = [t_0, t_f]$. Thus $I_i(t)$ is easily integrated on the time interval, leading to

$$I_i(t) = \mathcal{A}_i \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} (t - t_0). \quad (3.25)$$

For completeness, we consider the more general case of a piecewise constant temperature as a function of time; this procedure is applied, for instance, in Wong et al. [2015b] and Wong et al. [2016]. The overall experiment is seen as the succession of n_{step} smaller experiments with advancement of reaction factors $\alpha_i^{(k)}(t)$ where the initial condition $\alpha_i^{(k)}(t_0)$ depends on the state of the material at the end of the previous time interval. The temperature $T^{(k)}$ is maintained constant during the time interval $A^{(k)} = [t_0^{(k)}, t_f^{(k)}]$ and is equal to $T^{(k+1)}$ in the next interval. The transient part of temperature increase between each time interval is neglected because of its short duration compared to the whole process. The solution for the advancement of reaction factor as a function of time in the k interval is therefore

$$\alpha_i^{(k)}(t) = 1 - \left((1 - \alpha_i^{(k)}(t_0^{(k)}))^{1-r_i} - (1 - r_i) \mathcal{A}_i \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T^{(k)}} \right\} (t - t_0^{(0)}) \right)^{1/(1-r_i)}, \quad (3.26)$$

and the solution is $\alpha_i(t) = \alpha_i^{(k)}(t)$ for all $A^{(k)}$. Usually $\alpha_i^{(0)}(t_0^{(0)}) = 0$ in the first time interval, then the value for $\alpha_i^{(k)}(t_0^{(k)})$ in a given time interval is equal to $\alpha_i^{(k-1)}(t_f^{(k-1)})$ from the previous time interval. This expression simplifies further if we set $t_0^{(k)} = 0 \forall k$. If $n_{\text{step}} = 1$, then the solution is for a single isothermal condition.

Case 2: linearly increasing temperature

In this case, the temperature increases linearly with time until the end of the experiment, thus $T = \beta(t - t_0) + T_0$ and

$$I_i(t) = \int_{t_0}^t \mathcal{A}_i \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}(\beta(\tilde{t} - t_0) + T_0)} \right\} d\tilde{t}. \quad (3.27)$$

For the computation of $I_i(t)$, we apply the following change of variables: $v_i = -\mathcal{E}_i/(\mathcal{R}(\beta(t - t_0) + T_0))$ thus $dt = \mathcal{E}_i/(\mathcal{R}\beta v_i^2) dv_i$. The integral becomes

$$\begin{aligned} I_i(t) &= \int_{v_{i,0}}^{v_i} \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}\tilde{v}^2} \exp \{ \tilde{v} \} d\tilde{v} \\ &= \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \int_{v_{i,0}}^{v_i} \frac{\exp \{ \tilde{v} \}}{\tilde{v}^2} d\tilde{v} \\ &= \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \left(\left[-\frac{\exp \{ \tilde{v} \}}{\tilde{v}} \right]_{v_{i,0}}^{v_i} + \int_{v_{i,0}}^{v_i} \frac{\exp \{ \tilde{v} \}}{\tilde{v}} d\tilde{v} \right), \end{aligned} \quad (3.28)$$

where the last line results from the integration by parts. The second summand on the right-hand side is the exponential-integral function ([Abramowitz and Stegun \[1972\]](#))

$$\text{Ei}(v) = \int_{-\infty}^v \frac{\exp(\tilde{v})}{\tilde{v}} d\tilde{v} \quad (v > 0). \quad (3.29)$$

Thus, inserting the exponential-integral function into Eq. (3.28) leads to

$$\begin{aligned} I_i(t) &= \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \left[-\frac{\exp\{\tilde{v}\}}{\tilde{v}} + \text{Ei}(\tilde{v}) \right]_{v_{i,0}}^{v_i(t)} \\ &= \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \left(-\frac{\exp\{v_i(t)\}}{v_i(t)} + \frac{\exp\{v_{i,0}\}}{v_{i,0}} + \text{Ei}(v_i(t)) - \text{Ei}(v_{i,0}) \right) \\ &= \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \left(\frac{\mathcal{R}T(t)}{\mathcal{E}_i} \exp\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T(t)} \right\} - \frac{\mathcal{R}T_0}{\mathcal{E}_i} \exp\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} \dots \right. \\ &\quad \left. + \text{Ei}\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T(t)} \right\} - \text{Ei}\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} \right). \end{aligned} \quad (3.30)$$

This can be inserted into Eq. (3.24) to obtain the final expression of the advancement coefficient of each reactant i as a function of time.

Usually, the advancements of reaction are plotted against temperature instead of time for convenience. Gathering the terms that depend on temperature, we can write

$$\alpha_i = 1 - \left\{ (1 - r_i) \left[-\frac{\mathcal{A}_i}{\beta} T \exp\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} - \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \text{Ei}\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} + C_i \right] \right\}^{\frac{1}{1-r_i}}, \quad (3.31)$$

$$C_i = \frac{(1 - \alpha_i(T_0))^{1-r_i}}{1 - r_i} + \frac{\mathcal{A}_i}{\beta} T_0 \exp\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} + \text{Ei}\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} \frac{\mathcal{A}_i \mathcal{E}_i}{\beta \mathcal{R}}. \quad (3.32)$$

For the example in Section 3.5.1 we will consider the gas production from Eq. (3.15). Inserting the analytical solution gives

$$\begin{aligned} \nu_j &= \sum_{i=1}^{N_p} m_0 F_{ij} \left(\left\{ (1 - r_i) \left[-\frac{\mathcal{A}_i}{\beta} T \exp\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} \dots \right. \right. \right. \\ &\quad \left. \left. - \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \text{Ei}\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} + \frac{(1 - \alpha_{i,0})^{1-r_i}}{1 - r_i} + \frac{\mathcal{A}_i}{\beta} T_0 \exp\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} \dots \right. \right. \\ &\quad \left. \left. + \text{Ei}\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} \frac{\mathcal{E}_i \mathcal{A}_i}{\beta} \right] \right\}^{\frac{1}{1-r_i}} \right)^{r_i} \mathcal{A}_i \exp\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\}. \end{aligned} \quad (3.33)$$

Finally, computation of derivatives with respect to parameters are provided in Appendix A.1.

Particular case: linear Šesták-Berggren model

The case of $r_i = 1$ in Eq. (3.9) results in a linear system of uncoupled ODEs that is now addressed in this section for the two heating conditions.

We substitute $r_i = 1$ in Eq. (3.20) and the left-hand side leads to

$$\ln(u_{i,0}) - \ln(u_i(t)) = I_i(t). \quad (3.34)$$

For the (piecewise) isothermal case, with $I_i(t)$ from Eq. (3.25), this leads to

$$\alpha_i^{(k)}(t) = 1 - (1 - \alpha_{i,0}^{(k)}) \exp\left(-\mathcal{A}_i \exp\left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} (t - t_0^{(k)}) \right). \quad (3.35)$$

For the linearly increasing temperature, we substitute $I_i(t)$ from Eq. (3.30), thus leading to

$$\alpha_i = 1 - C_i \exp \left\{ -\frac{\mathcal{A}_i}{\beta} T \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} - \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} \right\}, \quad (3.36)$$

$$C_i = (1 - \alpha_{i,0}) \exp \left\{ \frac{\mathcal{A}_i}{\beta} T_0 \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} + \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} \right\}. \quad (3.37)$$

Matrix formalism

The case of $r_i = 1$ can be written using a matrix formalism thanks to the linearity of the equations. This particular case of $r_i = 1$ was already considered in the more general framework of competitive reaction mechanisms in [Herrador et al. \[2019\]](#) where the system of ODEs is no longer uncoupled.

Using the change of variables $u_i(t) = 1 - \alpha_i(t)$, the general reaction mechanisms Eq. (3.8) with Eq. (3.9) governing the evolution of the advancement and $r_i = 1$ can be written as

$$\frac{d\mathbf{u}}{dt} = L(t)\mathbf{u}, \quad (3.38)$$

where $L(t)$ is a square matrix whose coefficients are a linear combination of the reaction rates k_i and depend on time. In the case of a multi-component pyrolysis model such as the one presented here, the matrix L is diagonal with $l_{ii} = -k_i$.

3.4 Competitive pyrolysis reactions

In the previous section, the solid reactants were considered to be independent from each other. To include competitive reaction mechanisms, we assume that an initial solid reactant decomposes into pyrolysis gases and other solid products. These products, in turn, may become reactants¹ thus following different reaction paths that are in competition. The mathematical description of a general reaction scheme with competitive reactions applied to TPMs was introduced in [Torres-Herrador et al. \[2019a\]](#) and part of this section is based on this manuscript.

3.4.1 Simple competitive kinetic scheme

We consider a solid reactant $R^{(s)}$ of mass m_R that decomposes into generic products P_j of mass m_{Pj} with $1 \leq j \leq N_c$ and N_c the number of competitive reactions, as sketched in Fig. 3.4.




Figure 3.4: Illustration of a competitive mechanism. A solid $R^{(s)}$ reacts producing N_c products.

The evolution of each reaction j is controlled by its kinetic rate k_j that is modeled using the first order Arrhenius equation from Eq. 3.5. The solid reactant will degrade generating the products at a rate given by

$$\frac{dm_{Pj}}{dt}(t) = k_j(T(t)) (m_R(t))^{r_j}, \quad 1 \leq j \leq N_c \quad (3.39)$$

¹The variable is said to be activated, which is sometimes refer to as a depolymerization step.

where the exponent r_j is the reaction order. Such pyrolysis applications typically assume first order reactions ($r_j = 1$) (Tanaka [1995]; Galwey and Brown [2002]; Blondeau [2013]) due to the reduction on the number of parameters to be adjusted in the model and to the numerical advantages solving the system of equations. In the following, we assume first order reactions, and the parameter r_j will be omitted for simplicity and clarity.


Now, considering the N_c simultaneous reactions that the solid reactant may undergo, the total decomposition rate of the solid can be expressed as

$$\frac{dm_R}{dt}(t) = - \sum_{j=1}^{N_c} \frac{dm_{Pj}}{dt}(t) = - \sum_{i=j}^{N_c} k_j(T(t))m_R(t). \quad (3.40)$$

It is worth mentioning that the use of an advancement coefficient as defined in Eq. (3.10) is not meaningful here as an activated species can be both produced and consumed and would result in a non-monotonic advancement coefficient. Lautenberger [2007] used the cumulative production of a solid species in the definition of the advancement coefficient, but this is not investigated here.

3.4.2 Generalized competitive kinetic model

The generalization of the previous competitive mechanism to multi-component competitive reactions is represented in Fig. 3.5. We denote by S_i , with $1 \leq i \leq N$, all the species which can be either gaseous, with $i \in I_g$, or solid (reactant or product), with $i \in I_s$. The number of species in each set is $|I_g| = N_g$ and $|I_s| = N_s$ respectively with the total number of species $N = N_g + N_s$. Competitive mechanisms occur on the i -th component with mass m_i following the reaction rate $k_{i,j}$ with $1 \leq j \leq N_{ci}$ the index for its j -th competitive branch and N_{ci} the number of competitive reactions for the component S_i . Each reaction may generate more than one product species, S_l either solid or gaseous,



where $\zeta_{i,j,k}$ are the stoichiometric coefficients. Mass coefficients $\gamma_{i,j,k}$ are introduced as

$$\gamma_{i,j,k} = \frac{\zeta_{i,j,k} M_{i,j,k}}{\sum_{l=1}^N \zeta_{i,j,l} M_{i,j,l}} \quad 1 \leq j \leq N_{ci}, \quad 1 \leq i, k \leq N, \quad (3.42)$$

For $i \in I_s$, the variation the mass of a component S_i is given by the balance between the production and destruction of component S_i as

$$\frac{dm_i}{dt}(t) = - \underbrace{\sum_{j=1}^{N_{ci}} k_{i,j} m_i(t)}_{\text{destruction of } S_i} + \underbrace{\sum_{j=1}^N \sum_{l=1}^{N_{cj}} k_{j,l} \gamma_{j,l,i} m_j(t)}_{\text{production of } S_i}. \quad (3.43)$$

The first term on the right-hand side (RHS) is the destruction of S_i , which is the sum over all reaction branches from the solid component S_i towards products. This term is zero for gases since it is assumed that only solid components pyrolyze while gas phase reactions will be computed elsewhere, thus $k_{i,j} = 0$ for $i \in I_g$. The second term represents the production of S_i coming from all other components and their respective competitive reactions, except from $j = i$ for which $\gamma_{j,l,i} = 0$ (otherwise, there is no reaction).

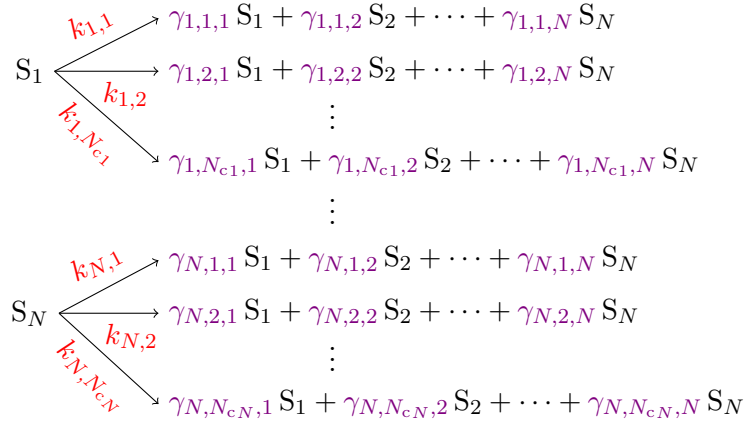


Figure 3.5: Generalized pyrolysis scheme with competitive mechanisms.

Matrix formalism

Eq. (3.43) defines a system of first order Ordinary Differential Equations (ODEs) that can be written under a matrix formalism as

$$\frac{d\mathbf{m}}{dt}(t) = L(t)\mathbf{m}, \quad (3.44)$$

where \mathbf{m} is a column vector with all the components m_i and L is the matrix of coefficients. The diagonal terms l_{ii} are given by the first term on the RHS of Eq. (3.43) while off-diagonal terms are given by the second term as $\gamma_{j,i,i} = 0$ for $j = i$. The matrix L can be reduced to a full rank matrix if all the dependent variables are removed, which are all the gaseous products and solid products that do not further react. In this case, the coefficient matrix is lower triangular due to the irreversibility of the pyrolysis process (Blondeau [2013]). In general, when considering multi-step non isotherm reactions, an analytical solution can not be found, see discussion in Appendix A.2.

In addition, it can be shown that the generalized model degenerates in a parallel model if the second term on the RHS (Eq. (3.43)) is zero (Herrador et al. [2019]) (see in Appendix A.3). In that case, the system becomes a diagonal matrix and analytical solutions can be found.

Generating different products from the same reactant provides great flexibility to the competitive reaction mechanisms. This means that different kinetic pathways may become dominant depending on the heating conditions leading to the possibility of “selecting” the pathway. This is intensively used in biomass pyrolysis where this selectivity allows to predict which products (gas, char or tar) will be mostly produced (Blondeau and Jeanmart [2012]).

Due to the nature of competitive schemes, a component S_i can be produced and consumed during the pyrolysis process (non diagonal matrix L). This makes that the advancement of reaction becomes irrelevant to describe the evolution of a particular reaction. A workaround to this is to describe the state of the reaction in terms of densities and if required define a global advancement of reaction from virgin to char (Blondeau [2013]; Lautenberger and Fernandez-Pello [2011]).

3.5 Examples

3.5.1 Six-equations multi-component model and sensitivity analysis


This example is used to illustrate the efficiency of the analytical solution from Section 3.3.3 by comparing it with a numerical solution of the non-linear system of ODEs. We consider the

Table 3.1: Parameter values for the six-reaction pyrolysis model from [Torres-Herrador et al. \[2019b\]](#), with the w_i parameters re-scaled in order to consider the pyrolysis of a pure resin.

R	w_i	$\log(\mathcal{A})$ (s^{-1})	\mathcal{E} (kJ mol^{-1})	n
1	0.0039	6.96	61.3	9.96
2	0.0258	6.59	77.6	5.65
3	0.0873	6.71	95.1	4.23
4	0.0804	6.67	103.0	4.38
5	0.0112	6.58	113.9	6.68
6	0.0254	6.35	175.2	8.85

pyrolysis scheme described in [Torres-Herrador et al. \[2019b\]](#) based on the experiments of [Bessire and Minton \[2017\]](#) that were described in Section 3.2. We verify the accuracy of the analytical solution and we show the computational gain achieved by using it in a case involving a linearly increasing temperature. We finally perform a sensitivity analysis to illustrate the use of the derivatives with respect to parameters.

[Torres-Herrador et al. \[2019b\]](#) performed a model calibration using a multi-component model from the results from the experiments of [Bessire and Minton \[2017\]](#) at the heating rate $\beta = 360 \text{ K min}^{-1}$. Six components ($N_p = 6$) were identified in the pyrolysis model by visual inspection of the maximum species production peaks in order to simulate the production of the 14 observed gas products. A model fitting procedure was then performed using a multi-objective genetic algorithm (MOGA) and the pyrolysis model was solved using the software PATO ([Lachaud and Mansour \[2014\]](#)) and integrated in time using a first-order implicit scheme. The 24 kinetic parameter and mass fraction values resulting from the fitting procedure of the six-reaction model are summarized in Table 3.1. We note that the relative weights w_i had to be rescaled compared to [Torres-Herrador et al. \[2019b\]](#) because the model for the mass loss and production from Eq. (3.11) and Eq. (3.15) is relative to the total mass of the sample, while in PATO it is relative to the mass of the resin. Therefore, we have to divide w_i by a factor equal to $(1 - m_{\text{fibers}}/m_0)$ which is around 0.43 (PICA contains a 43% of resin in mass). The sample mass and the mass production rates π_i from the decomposition of each component for the calibrated pyrolysis model are shown in Fig. 3.6. The six-reaction pyrolysis model with the calibrated species mass fractions is as follows:



The pyrolysis reaction model from above and the resulting mass production rates π_i from Fig. 3.6b can be related to the stages of pyrolysis described in Section 3.2.3. The production π_1 in Fig. 3.6b is related to the outgassing from stage 1, while π_3 is related to stage 3 and π_4 to stage 3. Both π_5 and π_6 are related to the production of H_2 at high temperature from stage 4. The CO observed at high temperatures does not have a clear origin, but it was included in π_6 . Remaining gases that do not have a clear reaction were included in a reaction corresponding to their production peaks, such as CO at low temperature (π_1) or xylene (π_3), or to a new reaction (π_2), such as for isopropanol.

Comparison of the analytical solution with a numerical solution

The solutions (both analytical and numerical) to the pyrolysis equations and their derivatives are implemented in the `pyrolysis_general` code, an in-house Python code developed at VKI. The exponential-integral function is computed using the `special.expi()` function from the `scipy` library. Its implementation is based on a fortran implementation of the series expansion of the exponential-integral $Ei(x)$ (with $x > 0$) (Abramowitz and Stegun [1972]). For the numerical solution, the integration of the system of ODEs is performed using the `integrate.solve_ivp()` function from the `scipy` library. In particular, it features an implicit multi-step Runge-Kutta method of the Radau IIA family of order 5 (Hairer and Wanner [1999]). Indeed, chemical reactions involving the Arrhenius equation can be stiff, i.e. there can be a rapid variation of the solution due to a fast degradation of a reactant compared to an other one, and for generality we decided to use a numerical solver capable of dealing with these different reaction scales (Curtiss and Hirschfelder [1952]). The varying time step is determined by the solver that keeps a relative tolerance, which controls the number of correct digits, that we set to 10^{-8} .

The numerical solution, the analytical solution and their derivatives are computed at 101 points in the temperature range with the parameter values from Table 3.1. The comparison of the analytical solution with the numerical solution is shown in Fig. 3.6. The solution is shown for the mass loss in Fig. 3.6a and the total gas production in Fig. 3.6b. The quantities on the y-axis are normalized by the total mass of the composite. Mass loss of the material is computed from Eq. (3.32) and Eq. (3.31) and the production of pyrolysis gas is computed from Eq. (3.33) and we can see from Fig. 3.6 that the two solutions match. The relative errors of the analytical solution compared to the numerical solution are computed as

$$\text{err}(T) = \left| \frac{y_{\text{analytical}}(T) - y_{\text{numerical}}(T)}{y_{\text{numerical}}(T)} \right|. \quad (3.45)$$

For the mass of the material, we have $\text{err}(T) \leq 2.5 \times 10^{-8}$ which is of the order of magnitude of the relative tolerance set to the numerical solver and show that the analytical solution is computed with high accuracy. For the total gas production, we have $\text{err}(T) \leq 6.2 \times 10^{-6}$, which is higher than for the mass because the reaction process rate needs to be recomputed from Eq. (3.16).

In Fig. 3.6b, we also show the production rate π_i coming from the decomposition of each reaction i with $\pi = \sum_{i=1}^{N_p} \pi_i$ to identify the different contributions of all reactions to the global production rate.

Derivatives are computed numerically using a first-order finite-difference formula by perturbing one by one the parameters about their nominal value. Each perturbed solution is computed using the analytical expressions Eqs. (A.1)–(A.10) (sensitivities of π and ν are the same here because we considered only one species). Results for the 24 derivatives are shown in Fig. 3.7. Once again, the two solutions match over the whole temperature range.

Computational gain for parameter inference

In order to illustrate the computational gain of this solution, we compared the execution time to compute 600 times the solution, which corresponds to the initial sample size for the MOGA algorithm in Torres-Herrador et al. [2019b]. Thus, the following analysis only characterizes the computational time for the initialization of the MOGA algorithm, but the total number of model evaluations is higher. Parameter values are kept constant, but we should keep in mind that varying them with the iterations could influence the computational time of the numerical solution. The computations are performed on a single CPU Intel Core i7-6820HQ using a single thread. The time to compute the 600 solutions numerically took 42.46 s while the analytical evaluations took 0.29 s, which is almost 150 times faster. Although they are not used in the MOGA algorithm, we can estimate the time that the derivatives would take to be computed. For

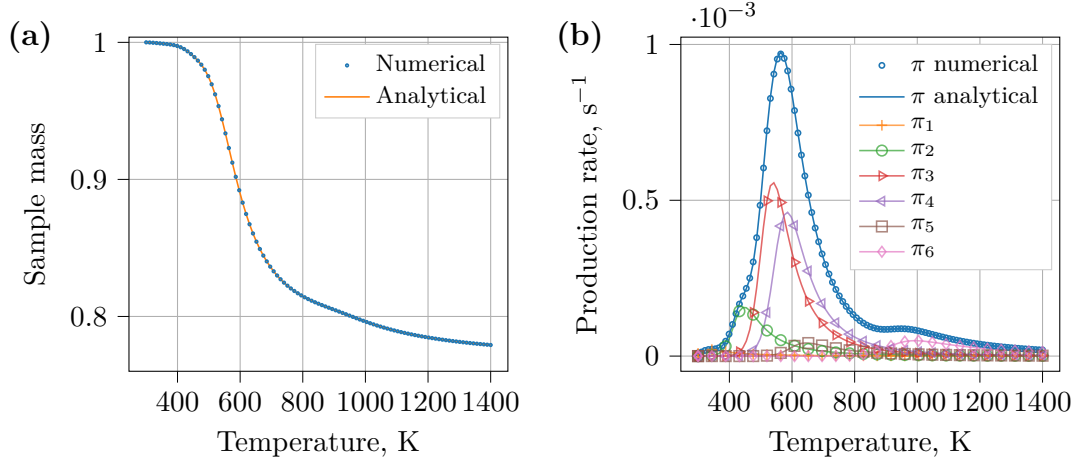


Figure 3.6: Comparison of the solution obtained from the analytical solution and the numerical integration for (a) the sample mass and (b) the total production rate.

the numerical computation of the derivatives, we need one additional model evaluation for each perturbed parameter for the finite difference evaluation that would be approximately 24×42.46 s (≈ 17 min), while the 600 analytical evaluations of the derivatives is taking approximately 3.35 s, which is approximately 300 times faster.

It is also important to note that the computational time increases proportionally with the number of equations in the model N_p (and so the complexity of the pyrolysis model). This increase in the computational time is more impacted by the number of equations with the numerical solution because of the adaptive time-step procedure and the time to compute the Jacobian matrix (derivative of the right-hand side of Eq. (3.16) with respect to t) needed for the implicit Radau IIA method, which was evaluated by the solver using finite differences.

Local sensitivity analysis

Finally, we perform a local sensitivity analysis from the graph of the sensitivities in Fig. 3.7. We can first observe the compensation effect between \mathcal{A} and \mathcal{E} (for both i) from Fig. 3.7(a) and 3.7(b): we see from the graphs that they have opposite effects on the solution (because \mathcal{E} appears with a negative sign in the exponential). This means that an increase of \mathcal{A} can be compensated by a decrease of \mathcal{E} while keeping the value of the solution π almost constant. This will result in a wide range of possible values for \mathcal{A} and \mathcal{E} when calibrating their values from experiments and explains the hurdle of calibrating them at the same time. We also see that the sensitivity of π with respect to \mathcal{E} is roughly three orders of magnitude higher than with \mathcal{A} and due to the fact that \mathcal{E} appeared in the exponential term.

Regarding the order of magnitudes of all the parameters, we see that the w_i parameters have the highest impact on the solution π , followed by n_i and finally \mathcal{E}_i and \mathcal{A}_i . w_i has the biggest impact on the production as it influences directly the production of char and is positive everywhere. Increasing wF_i will always result in an increase of the gas produced from the pyrolysis reaction and less char left at the end of the process. The shape of the curves for \mathcal{E}_i and n_i is almost the same, with a little shift in temperature and with three orders of magnitude difference. Finally, we note that reaction 1 has a high influence on the production rate despite its lowest contribution to the total mass ($w_1 = 0.00039$). Comparatively, reaction 5 and 6 have higher mass fractions, but $\partial\pi/wF_i$ (Fig. 3.7d) is lower than for reaction 1. This is due to the fact that changes in the mass fraction of the first reaction will impact the subsequent productions

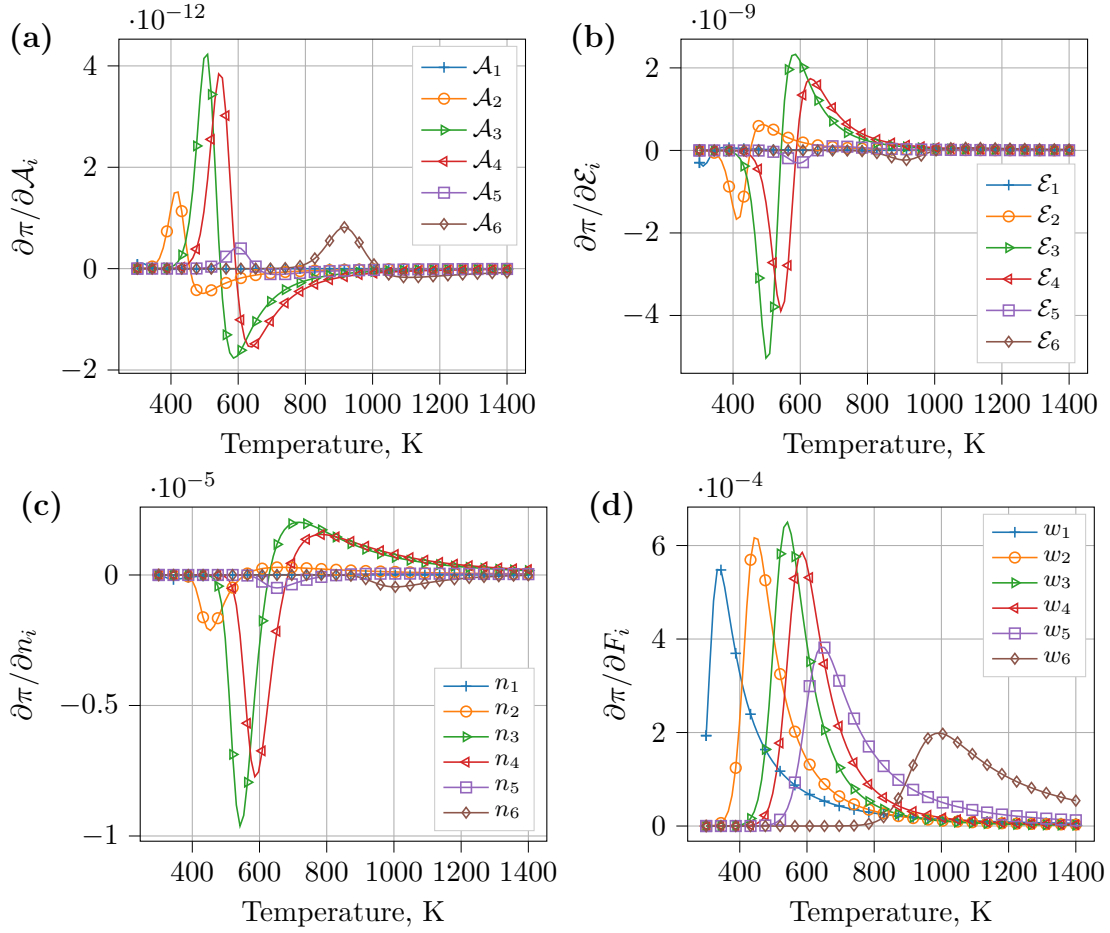


Figure 3.7: Comparison of the sensitivities computed analytically and using finite differences (numerical) with respect to the 24 parameters: **(a)** pre-exponential factors \mathcal{A}_i , **(b)** activation energies \mathcal{E}_i , **(c)** reaction orders n_i and **(d)** relative weight fractions w_i . The continuous and discontinuous lines are for the analytical solutions and the markers for the numerical solutions. For the clarity of the images, we displayed only 25 markers instead of the 101 points where the solution were computed.

rates.

This short analysis is valid locally about the nominal parameter values and provides a good overview of the relative influence of the parameters. We should also keep in mind that the analysis might be slightly different with a different combination of the nominal parameters. Finally, a global sensitivity analysis should be performed to give a more global overview of the impact of the parameters on the solution.

3.5.2 Model of Goldstein

Because of its common used in state-of-the-art numerical solvers for TPMs decomposition, the model of Goldstein [1969] is presented in this section. Goldstein [1969] calibrated the following multi-component parallel pyrolysis model

$$\frac{dm_i}{dT}(T) = -m_{0,i} \frac{\mathcal{A}'_i}{\beta} \exp\left(\frac{-\mathcal{E}_i}{\mathcal{R}T}\right) \left(\frac{m_i(T) - m_{\infty,i}}{m_{0,i}}\right)^{r_i}, \quad 1 \leq i \leq N_p, \quad (3.46)$$

Table 3.2: Kinetic parameters for the pyrolysis model of Goldstein [1969].

	\mathcal{A} (sec ⁻¹)	\mathcal{A}' (sec ⁻¹)	$\frac{\mathcal{E}}{\mathcal{R}}$ (K)	r	$\frac{m_{0,i}}{m_0}$	$\frac{m_{\infty,i}}{m_0}$
Reaction I	1.40×10^4	1.40×10^4	8560	3.0	0.25	0
Reaction II	4.9778×10^8	4.48×10^9	20450	3.0	0.75	0.5

with $N_p = 2$. The mass m_i can be expressed in terms of the extent of reaction $\alpha_i(t)$ from Eq. (3.10) as $m_i(T) = m_{0,i} - \alpha_i(T)(m_{0,i} - m_{\infty,i})$ $1 \leq i \leq N_p$ in order to retrieve Eq. (3.9) with $\mathcal{A}_i = \mathcal{A}'_i ((m_{0,i} - m_{\infty,i}) / m_{0,i})^{(r_i-1)}$. Parameters values for the model of Goldstein [1969] are provided in Table 3.2. The w_i correspond to $m_{0,i}/m_0$ with $\sum_{i=1}^{N_p} w_i = 1$. In PATO, we have $\sum_{i=1}^{N_p} w_i + w_c = 1$, thus $w_1 = 0.25$, $w_2 = 0.75 - 0.5 = 0.25$ and $w_c = 0.5$.

3.6 Summary of the chapter

We reviewed pyrolysis experiments that were performed on carbon/phenolic ablators in order to provide a better understanding of the chemical reactions occurring during the material decomposition. Typical models used for the simulation of pyrolysis decomposition were then reviewed. Multi-component models are widely used in current numerical solvers and are appropriate for single-step mechanisms without competitive reactions, for instance when considering the decomposition at a single heating rate. Pyrolysis models with competitive mechanisms can include multiple reaction paths in the reaction scheme, hence allowing us to extend the applicability of pyrolysis models to a wider range of heating rate when several reaction paths are present.

For the multi-component pyrolysis models, we derived a general analytical solution featuring a one-parameter Šesták-Berggren kinetic model. The mathematical derivation of the solution was performed for two typical heating conditions encountered in pyrolysis experiments, namely constant temperature and constant heating rate. The sensitivities of the solution with respect to the different model parameters were also calculated.

The analytical solution was illustrated on a six-reaction pyrolysis model for the linearly increasing temperature case. We showed the high accuracy of the proposed analytical solution by comparing it with a numerical solution, and we showed that the computational gain was significant. A local sensitivity analysis highlighted the well-known compensation effect and the relative importance of the mass fraction over the other parameters. However this analysis is only reflecting the local behavior about the nominal parameter values. In Chapter 7, this analytical solution will be useful for building a detailed pyrolysis models with a high number of equations and applying Bayesian inference algorithms to infer the parameter values for uncertainty quantification.

Numerical simulations of ablation experiments using a unified flow-material approach

This chapter presents the simulation results of the pyrolysis and char ablation of thermal protection materials subjected to high enthalpy flows. The flow in the porous medium is solved by means of the VANS equations in a unified flow-material approach in order to capture in the same domain of computation the interaction between of the chemically-reacting boundary layer and the decomposing porous medium. Multiple solid phases are considered in the porous medium, that are the fibers, the resin components and the char. The resin decomposition is modeled using a two-reaction (multi-component, single-step parallel) pyrolysis model, and the ablation of the char is modeled by assuming fibers of increased reactivity. Thermal and transport properties of the material are assumed to evolve smoothly between the fluid and the solid. The solver is verified against a one-dimensional test case with a state-of-the-art ablation-pyrolysis code, and the differences in the heat transfer are analyzed. Then, the pyrolysis and the char ablation of an iso-Q ablating material subjected to high enthalpy conditions similar to an arc-jet experiment are simulated. This work was carried out at Cenaero (Gosselies, Charleroi) in close collaboration with P. Schrooyen and under the direction of K. Hillewaert. Preliminary results were presented at the 47th AIAA Thermophysics conference and published in [Coheur et al. \[2017\]](#).

4.1 Introduction

The wide variety of physical phenomena occurring during an atmospheric entry requires the need for developing accurate numerical models that can take these numerous factors into account, both on the macroscopic and microscopic scales. For instance, the interaction between the boundary layer and pyrolysis and ablation gases needs to be correctly simulated to predict the blockage effect, and the decomposition of the porous medium needs to be represented to predict the gas production, effective material properties variations, or heat and mass transport within the material.

In an attempt to solve the complex response of an ablative TPM and the flow field through and around the material, [Schrooyen \[2015\]](#) developed the DGAblation branch of the Argo code, developed at Cenaero. The Argo code ([Hillewaert \[2013\]](#); [Carton de Wiart \[2014\]](#)) is a numerical solver based on a high-order discontinuous Galerkin method (DGM) ([Rivière \[2008\]](#)). The DGM combines the advantages of finite element and finite volume methods: high-order elements allow

a better accuracy on both structured and unstructured meshes with a higher computational efficiency and scalability due to the locality of the data in memory. The DGM also ensures the local conservation of the physical quantities (the net flux for each cell is in balance, as for finite volume methods) and achieves a low numerical dissipation and dispersion (Lv and Ihme [2014]). The DGAblation branch features a unified flow-material approach and solves the VANS equations, described in Section 2.3, in the whole computational domain, hence allowing to handle the porous material and the surrounding flow field within the same domain of computation. The code was shown to be able to simulate the ablation of carbon preform materials (Schrooyen et al. [2016, 2017]). In Schrooyen et al. [2016], the competition between surface and volume ablation was shown to be a key parameter for predicting the recession of the surface of porous rigid carbon fiber ablative materials. The rigid carbon fiber matrix, or carbon preform material, is the skeleton of carbon/phenolic composite ablative TPMs and has been widely studied in literature. Replication of Plasmatron experiments on carbon preform materials from Helber et al. [2016b] was performed and the comparison of the simulation results with the experimental results demonstrated the validity of the unified approach to simulate ablating materials subjected to atmospheric entry conditions (Schrooyen et al. [2017]). Finally, the code was further extended by Coheur [2016] to account for pyrolysis reactions in a porous medium. Pyrolysis reactions were implemented within DGAblation and the capability of such a pyrolysis model on a simplified case using a fictitious pyrolysis gas made up of a single species of carbon monoxide with conditions similar to a Plasmatron experiment on carbon/phenolic materials (Helber et al. [2016b]) were showed. However, the test case was less than one second of physical time and no char and fiber ablation was considered.


Therefore, in this chapter, we contribute to the development of the DGAblation branch of the Argo code in order to simulate the pyrolysis and char ablation of decomposing TPMs within the same unified flow-material approach. We implement a model for char ablation in the VANS framework and we perform a numerical simulation from the pyrolysis to char ablation of a carbon fibers/phenolic resin (CF/PR) TPM. A one-dimensional ablation test case is first simulated to verify the implementation of the char model. Then, benefiting of the continuous improvements performed by the Cenaero research team, the simulation of a 40 s pyrolysis and ablation experiment was made feasible. The results for the one-dimensional test case were published in Coheur et al. [2017].

This chapter is organized as follows. We first present the pyrolysis and char ablation models that are implemented within Argo (Section 4.2). Replication of ablation experiments of composite CF/PR materials are then simulated (Section 4.3). The first test case is a one-dimensional ablation test case and the results are compared with the state-of-the-art PATO code for verification (Section 4.3.1). The second one is a two dimensional axisymmetric test case subjected to arc-jet conditions (Section 4.3.2).

4.2 Flow in pyrolyzing reactive porous media

The VANS equations for non-charring reactive porous media, such as carbon fiber materials, were described in Section 2.3. Considering now a pyrolyzing porous medium, the formulation can be generalized to multiple solid phases inside the porous medium as sketched in Fig. 4.1. The medium is composed of a fluid phase, denoted by g , that fills the pores and of multiple solid phases σ_k . The elementary volume dV can therefore be written as $V = V_g + \sum_k V_k$ and the volume fraction for each phase is $\epsilon_{\sigma_k} = V_{\sigma_k}/V$.

We assumed that the porous medium is made up of three solid phases: the fibers σ_f , the pyrolyzing resin σ_r and the char σ_c . The averaged density of the overall solid phase is hence $\rho_s = \rho_f + \rho_c + \rho_r$. The fibers and the char follow heterogeneous reactions, which occur between the solid phase and the gas phase and depend on the average temperature, the local concentration of volatile reactants and the concentration of the solid reactant. In this case, the model

Figure 4.1: Porous medium with multiple solid phases σ_i .

reaction can be linked to the geometry of the solid phase through the specific surface introduced in Section 2.3.2. The global heterogeneous production of species i , $1 \leq i \leq N_g$ is denoted $\langle \dot{\omega}_i^{\text{het}}(\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle$ and will be further described in Section 4.2.2.

The resin is decomposed through pyrolysis reactions, which are a pure thermal process and depend on the average temperature and the average concentration of solid reactant. In Coheur [2016], a multi-component single-step pyrolysis reaction was implemented, similar to what was described in Section 3.3.2. The solid phase of pyrolyzing resin of average density $\langle \rho_r \rangle = \epsilon_{\sigma_r} \langle \rho_r \rangle_{\sigma_r}$ can be seen as being composed of several solid resin components (sub-phases) $\langle \rho_{r,I} \rangle$ with $1 \leq I \leq N_p$. The pyrolyzing polymer thus undergoes N_p independent pyrolysis reactions, each representing the decomposition of a fictitious solid component I composing the resin, $p_{r,I} \rightarrow \sum_j \gamma_{Ij} S_j$, $1 \leq j \leq N_g$, where γ is the mass fraction coefficient of species j produced from the reaction I and S_j denotes the gaseous species present in the mixture. The average matrix density can be written as

$$\langle \rho_r \rangle = \sum_{I=1}^{N_p} \langle \rho_{r,I} \rangle. \quad (4.1)$$

The conservation equation for a single resin component writes

$$\frac{\partial \langle \rho_{r,I} \rangle}{\partial t} = \langle \dot{\omega}_I^{\text{pyro}} \rangle, \quad 1 \leq I \leq N_p. \quad (4.2)$$

In the VANS equations from Eq. (2.28), the set of conservative variables, for $1 \leq i \leq N_g$ and $1 \leq I \leq N_p$, is given by

$$\mathbf{U} = (\epsilon_g \langle \rho_i \rangle_g, \langle \rho \mathbf{u} \rangle_g, \langle \rho E_{\text{tot}} \rangle, \langle \rho_f \rangle, \langle \rho_c \rangle, \langle \rho_I \rangle)^T. \quad (4.3)$$

The convective and diffusive fluxes express as

$$\mathbf{F}^c = \begin{pmatrix} \epsilon_g \langle \rho_i \rangle_g \langle \mathbf{u} \rangle_g \\ \langle \rho \rangle_g \langle \mathbf{u} \rangle_g \langle \mathbf{u} \rangle_g + p \\ \epsilon_g \langle \rho \rangle_g \langle \mathbf{u} \rangle_g \langle H \rangle_g \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}, \quad \mathbf{F}^d = \begin{pmatrix} -\langle \mathbf{J}_i \rangle \\ \langle \bar{\bar{\tau}} \rangle_g \\ \langle \bar{\bar{\tau}} \cdot \mathbf{u} \rangle + \lambda_{\text{eff}} \nabla \langle T \rangle - \sum_{i=1}^{N_g} h_i \langle \mathbf{J}_i \rangle \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}, \quad (4.4)$$

and finally the source term as

$$\mathbf{S} = \begin{pmatrix} \langle \dot{\omega}_i^{\text{het}} (\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle + \epsilon_g \dot{\omega}_i^{\text{hom}} (\langle T \rangle, \langle \rho_i \rangle_g) + \sum_{I=1}^{N_p} \gamma_{Ii} \langle \dot{\omega}_I^{\text{pyro}} (\langle T \rangle, \langle \rho_{r,I} \rangle) \rangle \\ \mathbf{F}_{\text{gs}} \\ - \sum_{i=1}^{N_g} \left(\langle \dot{\omega}_i^{\text{het}} (\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle + \epsilon_g \dot{\omega}_i^{\text{hom}} (\langle T \rangle, \langle \rho_i \rangle_g) \right) h_{f,i}^0 \dots \\ - \sum_{I=1}^{N_p} \left(\langle \dot{\omega}_I^{\text{pyro}} (\langle T \rangle, \langle \rho_I \rangle) \rangle h_{f,I}^0 + \sum_{j=1}^{N_s} m_{j,I} \langle \dot{\omega}_{j,I}^{\text{pyro}} (\langle T \rangle, \langle \rho_I \rangle) \rangle h_{f,j}^0 \right) \\ \sum_{i=1}^{N_g} \langle \dot{\omega}_i^{\text{het,f}} (\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle \\ \sum_{i=1}^{N_g} \langle \dot{\omega}_i^{\text{het,c}} (\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle \\ \langle \dot{\omega}_I^{\text{pyro}} (\langle T \rangle, \langle \rho_I \rangle) \rangle \end{pmatrix}. \quad (4.5)$$

When considering the element mass conservation (equilibrium chemistry), the mass continuity equation reads as

$$\frac{\partial}{\partial t} (\epsilon_g \langle \rho_i^e \rangle_g) + \nabla \cdot (\epsilon_g \langle \rho_i^e \rangle_g \langle \mathbf{u} \rangle_g) = -\nabla \cdot \langle \mathbf{J}_i^e \rangle + \langle \dot{\omega}_i^{\text{pyro,e}} \rangle. \quad (4.6)$$

The heterogeneous production term is usually neglected, but carbon production needs to be accounted for if volume ablation is considered. In Argo, only non-equilibrium chemical reactions are implemented and the source term from Eq. (4.5) is considered.

In the next sections, the models for the pyrolysis and char ablation are presented and are linked to the averaged production terms $\langle \dot{\omega}_I^{\text{pyro}} (\langle T \rangle, \langle \rho_I \rangle) \rangle$ and $\langle \dot{\omega}_i^{\text{het}} (\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle$.

4.2.1 Multi-component single-step pyrolysis model

This section focuses on the presentation of a model for pyrolysis and the expression of the pyrolysis decomposition rate $\langle \dot{\omega}^{\text{pyro}} \rangle$. The rate of decomposition of each resin compound $1 \leq I \leq N_p$ is assumed to follow an irreversible process as described in Eqs. (3.10)–(3.9), here expressed in terms of average densities

$$\frac{\partial \langle \rho_{r,I} \rangle}{\partial t} = -(\langle \rho_{r,0,I} \rangle - \langle \rho_{c,0,I} \rangle) k_{r,I} \left(\frac{\langle \rho_{r,I} \rangle - \langle \rho_{c,0,I} \rangle}{\langle \rho_{r,0,I} \rangle - \langle \rho_{c,0,I} \rangle} \right)^{r_I}, \quad (4.7)$$

where the forward reaction rate $k_{r,I}$ follows an Arrhenius law

$$k_{r,I} = \mathcal{A}_I \exp \left(\frac{-\mathcal{E}_I}{\mathcal{R} \langle T \rangle} \right). \quad (4.8)$$

In Eq. (4.7), $\langle \rho_{r,I} \rangle$ varies between its virgin density $\langle \rho_{r,0,I} \rangle$ and its char density $\langle \rho_{c,0,I} \rangle$. Because the char matrix is tracked into a distinct density variable, it is more convenient to work with pyrolysis density variables that vary from their maximum values at the beginning of the pyrolysis

process to zero at the end. Therefore, we introduce the translated resin density variable $\langle \rho'_{r,I} \rangle = \langle \rho_{r,I} \rangle - \langle \rho_{c,0,I} \rangle$, hence allowing us to rewrite Eq. (4.7) as

$$\frac{\partial \langle \rho'_{r,I} \rangle}{\partial t} = -\langle \rho'_{r,0,I} \rangle \mathcal{A}_I \exp\left(\frac{-\mathcal{E}_I}{\mathcal{R}\langle T \rangle}\right) \left(\frac{\langle \rho'_{r,I} \rangle}{\langle \rho'_{r,0,I} \rangle}\right)^{r_I}. \quad (4.9)$$

Following this, the compound-specific and the global pyrolysis advancement of reaction factors (see Eq. (3.10)) are expressed as

$$\alpha_I = \frac{\langle \rho'_{r,0,I} \rangle - \langle \rho'_{r,I} \rangle}{\langle \rho'_{r,0,I} \rangle}, \text{ and } \alpha = \sum_{I=1}^{N_p} w_I \alpha_I = \frac{\langle \rho'_{r,0} \rangle - \langle \rho'_r \rangle}{\langle \rho'_{r,0} \rangle}, \quad (4.10)$$

respectively. The parameters used here are the two-reaction pyrolysis model from Goldstein [1969] summarized in Table 3.2. In that case, the global volume-averaged solid density can be expressed as the following sum:

$$\langle \rho_s \rangle = \langle \rho_I \rangle + \langle \rho_{II} \rangle + \langle \rho_c \rangle + \langle \rho_f \rangle. \quad (4.11)$$

4.2.2 Model for char ablation

When heated, the resin is thermally degraded by pyrolysis and progressively transformed into a low-density carbon residue, the char. For this charred matrix, Lachaud et al. [2010] proposed to use an expression similar to the one for fiber ablation, where

$$\langle \dot{\omega}_i^{\text{het},c} \rangle = -S_c k_c^{i,C(s)} \langle \rho_i \rangle_g. \quad (4.12)$$

As a first approximation, the reactivity of the charred material $k_c^{i,C(s)}$ is assumed to be the same as the one of the fibers, i.e. $k_c^{i,C(s)} = k_f^{i,C(s)}$. However, it should be kept in mind that the reactivity of the carbonized matrix is theoretically higher. In fact, the structure of the charred matrix features many more defects than the carbon fibers, and an estimation of the reactivity could be for instance $k_c^{i,C(s)} = 10k_f^{i,C(s)}$ (Lachaud et al. [2010]).

Matrix surrounding the fibers

For the modeling of S_c , it is assumed that the matrix is surrounding the fibers (Lachaud et al. [2010]) and that during the ablation process, the charred matrix shrinks uniformly around the fibers. The motivation for using this model is because a model of uniform radial recession is already implemented within Argo for the carbon fibers recession. Hence, the implementation of such a model for the charred matrix recession only requires to define an equivalent fiber radius that includes the increase of carbon coming from the char residue. This equivalent fiber radius is defined as

$$r_e = r_{f,0} + e_c, \quad (4.13)$$

where e_c denotes the char thickness. Expressed in terms of solid fraction ϵ_s and initial fiber fraction $\epsilon_{f,0}$, the equivalent fiber radius writes

$$r_e = r_{f,0} \sqrt{\frac{\epsilon_s}{\epsilon_{f,0}}} \quad (4.14)$$

which has exactly the same formulation as the simple non-charring material made up only of carbon fibers, except that the equivalent radius can have values bigger than the initial fiber radius $r_{f,0}$. Indeed, the solid fraction ϵ_s can have a value higher than the initial fiber fraction

$\epsilon_{f,0}$ because of the presence of the char stemming from the resin pyrolysis. Following the same development as for non-charring materials, the specific surface is defined as

$$S_c = \frac{2}{r_{f,0}} \sqrt{\epsilon_{f,0} \epsilon_s} \quad (4.15)$$

which is valid for both the charred and preform materials.

For instance, the char could also be represented as a pore-filling matrix, in which the charred matrix is considered to fill homogeneously the pores between the fibers with a given porosity (Lachaud et al. [2010]). The matrix is not anymore located around the fibers but is still filling the void between the fibers as for the resin. The functional relation between the specific surface and the porosity, that depends on the topology of the medium, is not well known and experiments on the carbonized matrix would be needed to estimate this dependency.

Therefore, the global heterogeneous production is given by the sum of the production from the char and the fibers

$$\langle \dot{\omega}_i^{\text{het}} (\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle = \langle \dot{\omega}_{i,f}^{\text{het}} (\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle + \langle \dot{\omega}_{i,c}^{\text{het}} (\langle T \rangle, \langle \rho_i \rangle_g, \langle \rho_s \rangle) \rangle. \quad (4.16)$$

The reason of separating the two is because other model could be implemented, which does not necessarily have a common expression for the specific surface common to char and fiber ablation as in Eq. (4.15). However, the implementation and the study of different models for char ablation is left for future work.

Transition from char to fiber ablation

Heterogeneous reactions start only when pyrolysis is completed or near completion. There is no oxygen in the pyrolysis gas that is produced, and the pressure is high enough to prevent oxygen from the boundary layer to diffuse and react with the char and the fibers. However, when the char and the fibers are no longer protected by the resin, heterogeneous reactions are triggered. Assuming that the char matrix surrounds the fibers, the char density $\langle \rho_c \rangle$ is the first one to react. Then, fiber ablation starts when the char is itself ablated or near completion.

In order to decide whether heterogeneous reactions stem from the ablation of the char or from the ablation of the fibers, the transition between the char and the fiber heterogeneous reactions is treated smoothly by using an hyperbolic tangent function to prevent the use of “if” statement in the implementation, hence avoiding any abrupt variation in the source term and facilitating the analytical derivation of Jacobian matrices. We define the smoothing function as

$$s_c^- = \left[0.5 \left(1 - \tanh \left(\frac{2\pi}{\delta_\rho} (\langle \rho_c \rangle - \rho_{c,\text{trig}}) \right) \right) \right], \quad (4.17)$$

such that the heterogeneous term can be expressed as

$$\langle \dot{\omega}^{\text{het}} \rangle = \langle \dot{\omega}_c^{\text{het}} \rangle + \langle \dot{\omega}_f^{\text{het}} \rangle s_c^-. \quad (4.18)$$

The transition length depends on δ_ρ that should neither be too wide nor too small and is taken to be equal to 5×10^{-3} . The value of $\rho_{c,\text{trig}}$ determines when the fiber ablation starts with respect to the char density. A value of $\rho_{c,\text{trig}} = 10^{-2}$ is used, thus char and fiber ablation will be competing as $\langle \rho_c \rangle$ approaches $\rho_{c,\text{trig}}$ until the char has completely disappeared.

4.2.3 Porous material properties

The numerical simulations presented in this chapter uses the material properties of the Theoretical Ablative Composite for Open Testing (TACOT) material, a fictive and open-source material of known properties developed for enabling code-to-code comparison (Lachaud et al. [2012]; van Eekelen et al. [2014]). Some of the properties are coming from PICA, for which the properties

were calibrated on the FIAT model predictions to experimental arc-jet data (Covington et al. [2008]). These properties need to be extended in order to be valid with the unified approach.

Regarding the permeability, both the Carman-Kozeny and a linear model are implemented within Argo (Schrooyen [2015]). While Carman-Kozeny model is based on the definition of a specific surface area of the solid $S_{f,0}$ (Dullien [1991]) which is clearly defined for the fibers and for the char assuming that it surrounds the fiber, there is no rigorous definition of the specific surface for the resin matrix. While the permeability of carbon preform was investigated in Panerai et al. [2016] and Borner et al. [2017], the permeability for CF/PR materials has received little attention (White et al. [2016]). The solid microstructure is unknown and the permeability is therefore assumed to evolve linearly with porosity according to

$$\kappa^{-1} = \frac{1}{\kappa_0} \frac{(1 - \epsilon_g)}{\epsilon_{s,0}}, \quad (4.19)$$

where $\kappa_0 = 1.6 \times 10^{-11} \text{ m}^2$ is the initial permeability of the virgin matrix and $\epsilon_{s,0}$ the initial solid fraction with $\epsilon_{s,0} = \epsilon_{f,0} + \epsilon_{m,0} + \epsilon_{c,0}$.

Regarding the tortuosity, an initial value of 1.2 (virgin material) is adopted and evolves linearly with porosity between the freestream tortuosity ($\eta = 1$) and its initial value:

$$\eta = \frac{1 - \eta_0}{1 - \epsilon_{g,0}} (\epsilon_g - 1) + 1. \quad (4.20)$$

Note that using the formulation of Eqs. (4.19)–(4.20), both the char permeability and tortuosity from TACOT are not exactly retrieved within the material, with $\kappa_{c,\text{TACOT}} = 2 \times 10^{-11} \text{ m}^2$ and $\eta_{c,\text{TACOT}} = 1.1$ respectively.

The effective thermal conductivity appearing in Eq. (4.4) also evolves with porosity to go continuously from the pure fluid conductivity where the solid fraction is zero to its effective conductivity in the porous medium:

$$\lambda_{\text{eff}} = \epsilon_g \lambda_g + \epsilon_s \lambda_s. \quad (4.21)$$

The solid thermal conductivity λ_s is an average between its different constituents, which are in this case the fibers, the resin and the char. In practice however, only the thermal conductivity of the global virgin and char states together with the fibers is measured and instead, we consider the solid thermal conductivity to vary with the advancement of pyrolysis reaction

$$\lambda_s = (1 - \alpha) \lambda_{\text{virgin}} + \alpha \lambda_{\text{char}}. \quad (4.22)$$

Because TACOT provides the effective thermal conductivity for the virgin and char states and not the intrinsic conductivity, one must consider $\lambda_{\text{virgin}} = \lambda_{\text{TACOT,virgin}} / \epsilon_{s,0}$ and $\lambda_c = \lambda_{\text{TACOT,char}} / \epsilon_{s,c}$; yet it still includes the gas thermal conductivity. Nevertheless, we still use λ_{virgin} and λ_c , and by substituting Eq. (4.22) into Eq. (4.21), it is expected that the influence of the average gas conductivity $\epsilon_g \lambda_g$ will have little influence, as the average solid conductivity is almost one order of magnitude higher than the gas average conductivity. Ideally, virgin and char thermal conductivities in the void should be provided for the present unified approach.

Finally, regarding the internal energy and enthalpy of the resin, modified properties of TACOT were obtained in Coheur [2016] by subtracting from the properties of the fibers the virgin and char properties.

4.3 Results and discussion

We assume finite-rate ablation-pyrolysis chemistry, the flow field is assumed to be in thermo-chemical equilibrium, and we consider multicomponent diffusion and transient simulations.




Figure 4.2: Sketch of the domain and boundary conditions for the Ablation Workshop test case with no recession. Slip adiabatic walls are imposed on the sides to simulate a 1D domain within Argo.

4.3.1 Verification: Ablation Workshop test case

In this section, a code-to-code comparison with the Ablation Workshop test case #1 ([Lachaud et al. \[2011c\]](#)) is performed. The objective of this section is to show that the complete set of VANS equations together with the pyrolysis model lead to the same results as with state-of-the-art pyrolysis-ablation material codes.

The computational domain for the first test case is sketched in Fig. 4.2. It consists in a 5-cm high one-dimensional sample for which no recession of top surface is assumed (no fiber ablation). We choose to reproduce a test case with no recession because the ablating boundary conditions implemented in standard material-response codes can hardly be reproduced within this framework, as we will see in the next section. As the sample will be heated, the resin matrix will decompose because of the pyrolysis reactions occurring within the sample. The sample is at an initial temperature of 298 K and a pressure of 1 atm and is filled with a pure gas of CO. The sample is then heated from the top during 60 s, where the temperature is first increased gradually from 298 K to 1644 K during 0.1 s and secondly keeps its maximum value of 1644 K during the remaining 59.9 s, as sketched on the right in Fig. 4.2. The bottom wall is adiabatic with zero momentum (stick wall) and the side walls are also adiabatic with no velocity gradient towards the sides (slip wall). The pressure at the top is fixed at 1 atm with the same species mass fraction as the initial sample composition. We assume a simplified mixture of pyrolysis gas products composed of seven species from the finite-rate chemistry model of [Sykes \[1967\]](#) and the mass fractions are summarized in Table 4.1. Multicomponent diffusion within the mixture could be considered, but the collision integrals for most of the species of the pyrolysis gas are not yet computed. We therefore use a constant Lewis number equal to one to be as closed as possible to the test case conditions described in [Lachaud et al. \[2011c\]](#). Furthermore, we do not consider the species to react with each other neither to dissociate and the mixture is considered to be frozen. Note that for this case where there is no pure fluid region, one can consider directly $\lambda_{\text{eff}} = \lambda_{\text{TACOT}}$. The computational domain is made up of 201 quadrilateral elements with increasing size from the surface towards the bottom wall with the smallest cell size equal to 1.02×10^{-2} cm and the largest to 4.98×10^{-2} cm.

Table 4.1: Species mass fractions for the pyrolysis gas (Sykes [1967]).

Gas species	CO ₂	CO	C ₆ H ₆	C ₆ H ₅ OH	CH ₄	H ₂ O	H ₂
Initial mass fraction	0.0001	0.9994	0.0001	0.0001	0.0001	0.0001	0.0001
Pyrolysis gas mass fraction	0.03766	0.0882	0.04396	0.4581	0.08748	0.23	0.0546

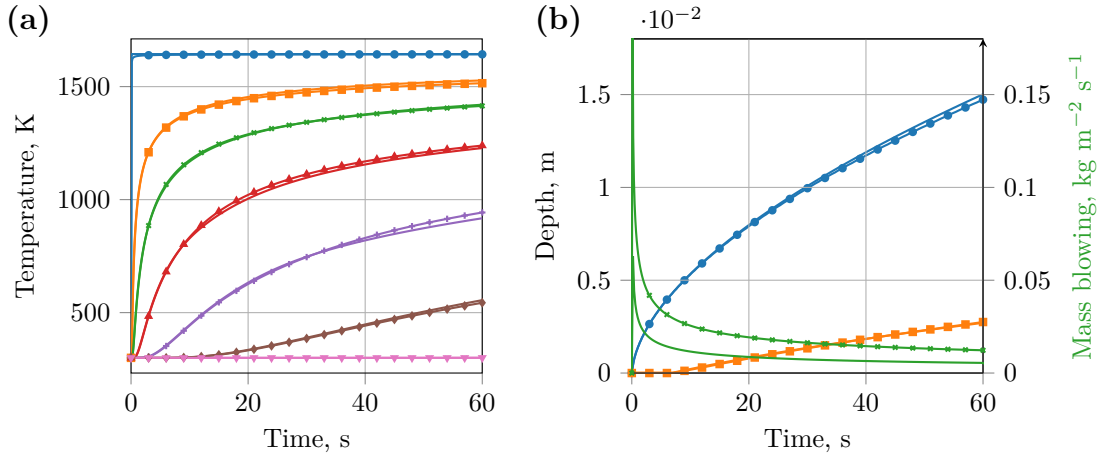


Figure 4.3: Comparison of the results of PATO (continuous lines) and Argo (lines with markers). (a) The temperature inside the sample as a function of time is shown at several thermocouple locations: surface (●), 1 mm (■), 2 mm (×), 4 mm (▲), 8 mm (+), 16 mm (◆), 50 mm (▼). (b) The 98% virgin (●) and 2% char (■) fronts, and mass blowing (×) rate are shown as a function of time.

Results

Temperature measurements at the thermocouple locations are plotted against the results obtained from PATO (Lachaud et al. [2015]) in Fig. 4.3(a). The results shown are in good agreement for all thermocouples, except for slight differences that can be observed after the initial temperature increase, notably at 1 mm, 4 mm and 8 mm. This difference is attributed to the different law for momentum conservation that is implemented in the two codes.

The mass blowing flux through the top surface is computed and compared to the results obtained with PATO. In this simple 1D case, the mass blowing flux through the surface is given by the density of the gas times the velocity through the surface. The blowing rate is shown in Fig. 4.3(b) and the results are again in close agreement, except that a higher value for the blowing is observed with Argo. This is explained by the higher molar density of the mixture of pyrolysis gas that is used in Argo, for which the molar density at 1 atm and 1644 K is equal 19.35 kg kmol⁻¹, computed using the Mutation++ library (Scoggins et al. [2020]), compared to the molar density of TACOT equal to 11.38 kg kmol⁻¹, read from tables within PATO. Therefore, there is slightly more mass that is ejected through the surface.

The position of the virgin and the char zones within the material are also compared in Fig. 4.3(b). The front zones are defined using a threshold value for the resin and charred densities given by

$$\langle \rho_{v,t} \rangle = (\langle \rho_{c,0} \rangle + \langle \rho_{f,0} \rangle) + 0.98 (\langle \rho_{s,0} \rangle - (\langle \rho_{c,0} \rangle + \langle \rho_{f,0} \rangle)), \quad (4.23)$$

$$\langle \rho_{c,t} \rangle = (\langle \rho_{c,0} \rangle + \langle \rho_{f,0} \rangle) + 0.02 (\langle \rho_{s,0} \rangle - (\langle \rho_{c,0} \rangle + \langle \rho_{f,0} \rangle)), \quad (4.24)$$

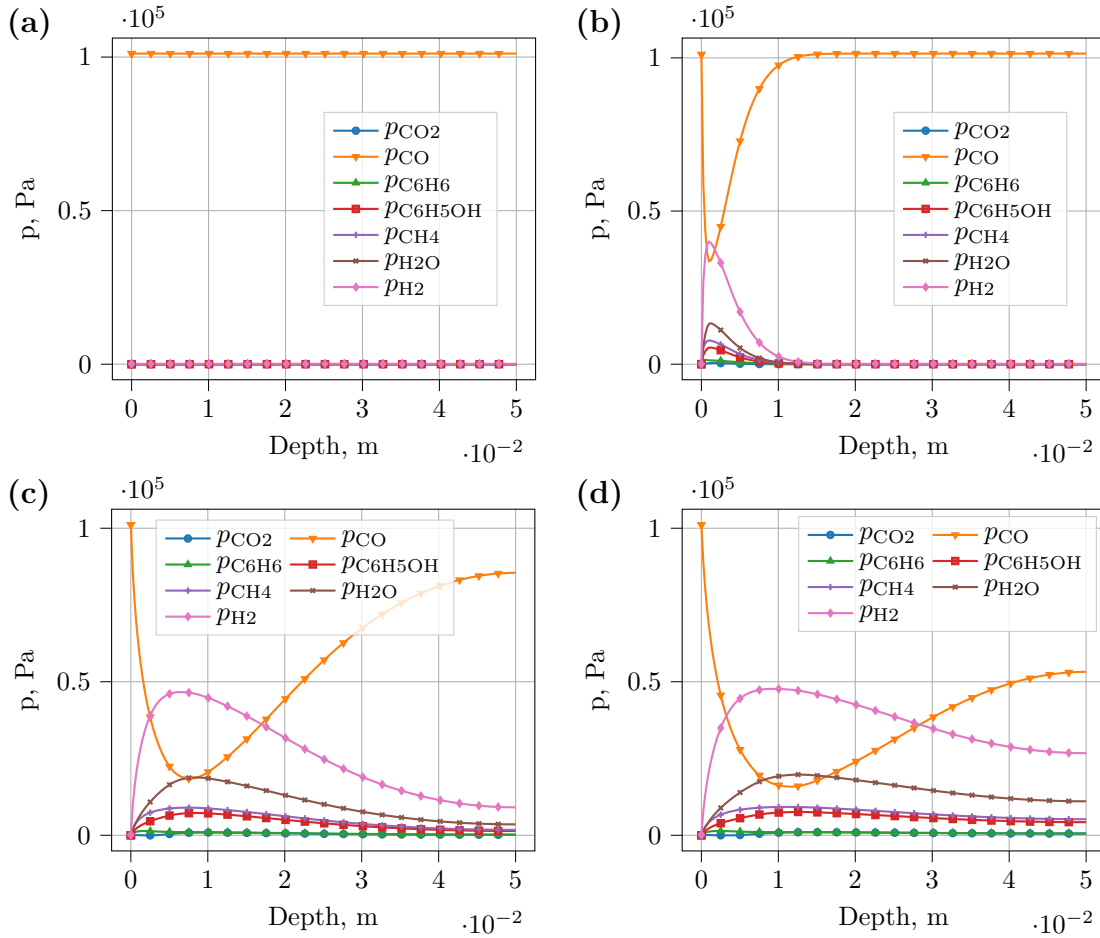


Figure 4.4: Species partial pressures for the Ablation Workshop test case at (a) $t = 0$ s, (b) $t = 1.0$ s, (c) $t = 30$ s, and (d) $t = 60$ s.

where “ t ” stands for *threshold* value. Results appear to be again in good agreement with PATO. A discrepancy of less than 5% is observed for the 98% virgin which is attributed again to the difference in gas density and momentum law.

Finally, the species partial pressures in the material are shown in Fig. 4.4. Because of the finite-rate chemistry model of Argo, each species is tracked individually by the solver. The production of pyrolysis gas can be clearly observed from the graph. Initially, only CO is present, which is our initial assumption. Then, very rapidly after the increase of surface temperature, species are produced close to the surface (Fig. 4.4(a)). The partial pressure of CO decreases and pyrolysis gases are produced abundantly, with H_2 , H_2O and CH_4 the most dominant species. Partial pressure for heavier species (C_6H_5OH and C_6H_6) are smaller and only a little amount of CO_2 is observed. From Fig. 4.4(c) and Fig. 4.4(d), we see that the production peak shifts as the pyrolysis front is moving, which is in agreement with what was already observed in Fig. 4.3(b). The peaks progressively spread as the species diffuse toward the top surface and the bottom wall of the material.

However, it is important to mention that the Dirichlet condition on the gas composition assumed at the surface may be questionable physically as the surface species composition should evolve according to the external flow and the species production coming from the in-depth material pyrolysis. A better assumption could be to consider a Neumann condition on the

species mass fractions, although in this case, we did not achieve to obtain a solution. Hence, discrepancies in the results could also stem from the differences in the top boundary condition applied.

4.3.2 Pyrolyzing material subjected to arc-jet conditions

In this section we perform a two-dimensional axisymmetric simulation of the ablation of a material subjected to aerothermal conditions close to arc-jet experiments. Boundary conditions are applied at the inlet and the outlet of a channel in which the sample is located, as showed in Fig. 4.5, and the comparison with material response code is not feasible. Within the unified approach, it is more difficult to control the heat flux at the surface of the material, which depends on the enthalpy of the flow provided at the inlet, whereas in material-response codes the heat flux can be directly provided at the surface of the material. The purpose of the test case of this section is therefore not to provide a code-to-code comparison but rather to show the capability of the code to simulate an ablation problem with pyrolysis and ablation of the material.

The computational domain and the mesh is represented in Fig. 4.5 and was generated using the software Gmsh (Geuzaine and Remacle [2009]). The sample considered in this section is described in van Eekelen et al. [2014]. It is designed to have minimal shape changes when the material is receding and is such that it features almost iso-flux contours, and is thus referred to as *iso-Q*. The sample is cylindrical with an elliptic front surface that minimizes the increase of heat flux at the shoulder compared to classical sphere-cylinder shaped samples used in arc-jet experiments (Milos and Chen [2010]). The shape of the initial geometry of the sample is sketched in Fig. 4.6(a), and is similar to the geometry used in Weng and Martin [2014] and Lachaud et al. [2015] to perform numerical simulations of a pyrolysis gas flow within a porous material. As seen in Fig. 4.6(b), the mesh at the surface of the material is refined in order to track the recession, and a high species production is expected to occur close to the surface due to pyrolysis and ablation. The hybrid mesh from Fig. 4.5 is composed of 2781 triangular elements and 396 quadrangles. The quadrangles are located around the surface (Fig. 4.6(b)) and follow the surface curvature. The height of the quadrangles progressively decreases toward the initial position of the surface. With a minimum cell width of 2.7×10^{-5} m at the interface and a second-order interpolation, we fix the value of the numerical interface thickness ϵ_ϕ to 5×10^{-4} m, such that it is not too large compared to the size of the sample, and such that it respects the empirical condition from Eq. (2.41) at the initial time and after the surface has moved (where the mesh size is larger). The influence of ϵ_ϕ was investigated in Schrooyen [2015] and is not studied here.

Boundary and initial conditions

The boundary conditions applied, sketched in Fig. 4.5, are the ones typically encountered in an arc-jet testing (Milos and Chen [2009]) and which correspond to the thermal transfer encountered during a hypersonic reentry, that are the post-shock conditions, i.e., a high enthalpy condition at the inlet. The contact between the support (holder) and the sample is assumed to be perfect.

The initial freestream temperature is equal to 3000 K and inside the material, it is equal to 298 K. Because of the increase in computational time, only a mixture of air with five species plus carbon monoxide for pyrolysis and ablation product is assumed, instead of including the simplified phenol mixture for the pyrolysis gas as in the previous section. The initial gas composition is the equilibrium composition at the initial temperature, hence, within the material, there is neither carbon monoxide nor dissociated species. For the heterogeneous reactions, the following oxidation reactions are considered:



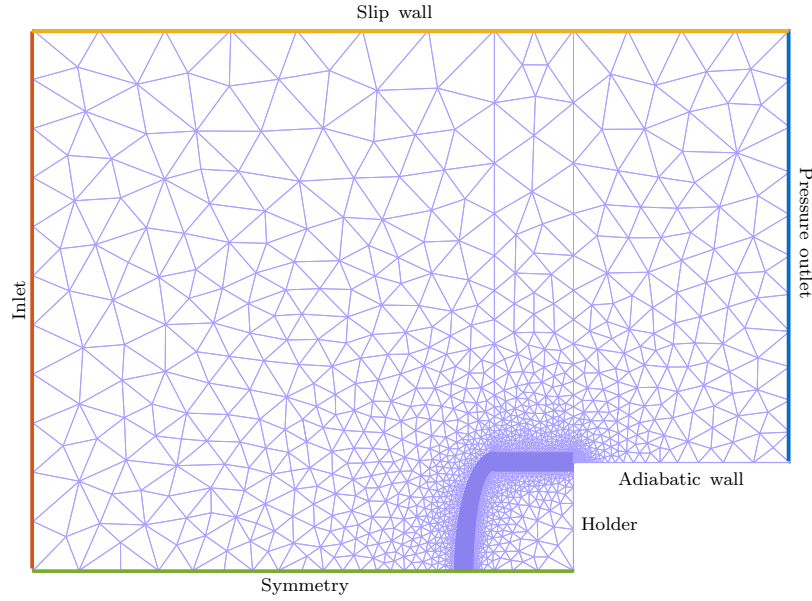


Figure 4.5: Full computational domain and mesh for the iso-Q test case and boundary conditions.

Table 4.2: Arrhenius coefficients for oxidation reactions.

	Reacting species	
	O	O ₂
\mathcal{A}_0	3.22	5.73
\mathcal{E}/\mathcal{R}	0	9.65e3
m	0.5	0.5

for which reverse reactions do not need to be considered as the equilibrium constant is very small (Park et al. [2001]). The presence of CN was also observed experimentally (Helber [2016]), however reactions with atomic nitrogen are not taken into account. The forward reactions rates $k_f^{i,C(s)}$ for the two previous reactions follow an Arrhenius law with a temperature-dependent pre-exponential factor

$$k_f^{i,C(s)} = \mathcal{A}_0 T^m \exp\left(\frac{-\mathcal{E}}{\mathcal{R}\langle T \rangle}\right), \quad (4.27)$$

where i denotes the limiting species and the Arrhenius coefficients are taken from Park et al. [2001] and given in Table 4.2. No catalytic recombination at the surface of the material is considered. The heat flux coming from the radiation of the gas particles (see Eq. (2.20)) is neglected.

Results

In Fig. 4.7 the averaged solid densities within the sample are analyzed. The solid densities along the stagnation line are shown in Fig. 4.7(a) for the total density (which represents the sum of the fibers, the char, and the two resin components) and the resin density (which represents the sum of the two resin components). The recession of the surface can be computed from the char front line and is equal to $770 \mu\text{m}$, which leads to an average recession rate of $19.25 \mu\text{m s}^{-1}$. The resin content after simulating 40 s of the experiment is shown in Fig. 4.7(b) for the whole

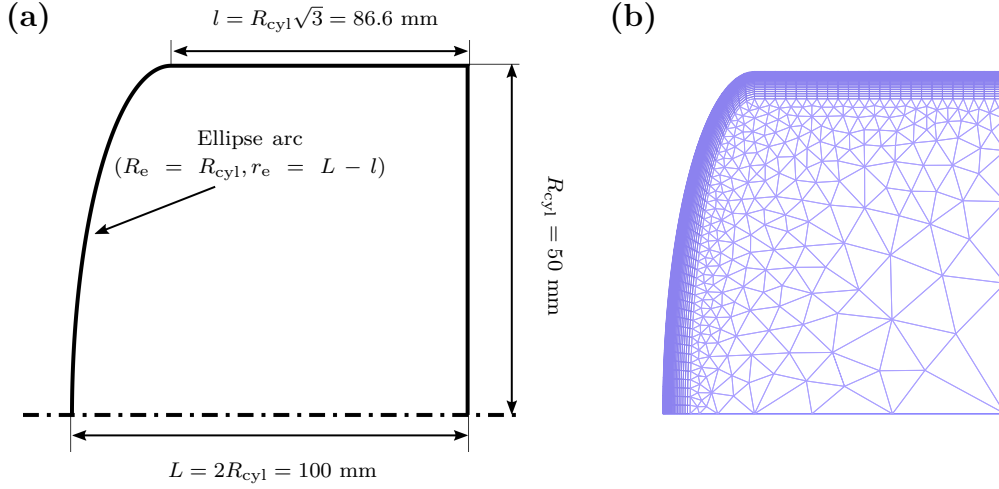


Figure 4.6: (a) Dimensions of the iso-Q elliptic sample. (b) Zoom on the hybrid mesh of the sample. Quadrilateral elements are used in a thin layer below the surface.

sample. The recession is almost uniform on the part of the sample that faces the flow field. The recession is the highest at the shoulder and then decreases progressively toward the holder. Close to the holder, we can observe that the resin has not decomposed because of the cold wall condition imposed at the holder.

In Fig. 4.7(b), we also represented the pyrolyzing front (dotted line) in the case where the holder is considered to be adiabatic. This impacts mainly the resin content close to the intersection between the material, the holder and the fluid where the temperature increases. Resin is not attached anymore to the holder at the corner because it has pyrolyzed. One can observe a bump, which is due to the injection of gas species in the boundary layer. The recession is therefore higher at the corner: the fluid exits at a higher velocity than close to the wall, thus the boundary layer is thicker and the blockage effect more important, thus reducing the temperature at that location and less resin is decomposed.

The temperature field is showed in Fig. 4.9. The temperature increases almost uniformly below the surface of the material at the front. The temperature is also higher at the shoulder, which corroborates the higher recession around this position. On the side, toward the holder, lower temperatures are observed and is even transported outside the surface of the material. In Coheur [2016], it was observed that mass was ejected on the sides of the sample. Those gases from the inner part of the material can carry flow at a lower enthalpy than the boundary layer gases, and is ejected toward the boundary layer, hence reducing the temperature.

Fig. 4.10 shows the pressure distribution. High pressures ($1.062 \times 10^4 \text{ Pa}$) are observed close to the pyrolyzing surface and in particular close to the shoulder region, whereas lower pressures are observed at the side of the sample. The partial pressure of the gas species along the stagnation line are also showed in Fig. 4.8 at four different time instants. In Fig. 4.8(a), pyrolysis and ablation reactions occur at the surface and the production of CO is important. One can note the presence of CO within the material, which is due to the initial presence of O and O₂ that comes from the oxidation reaction of the char during the first steps of the simulation, but the amount of carbon removed is negligible. Then, CO progressively diffuses below the surface (Fig. 4.8(b) and Fig. 4.8(c)) increasing its partial pressure. The partial pressures of the other species decrease abruptly, but one can still observe a large amount of N₂ and a few traces of NO. Thus, it seems that the initial amount of N₂ is trapped inside the material due to the surrounding pyrolyzing front. As it will be seen in the next paragraph, the flow below the pyrolyzing front is slowly

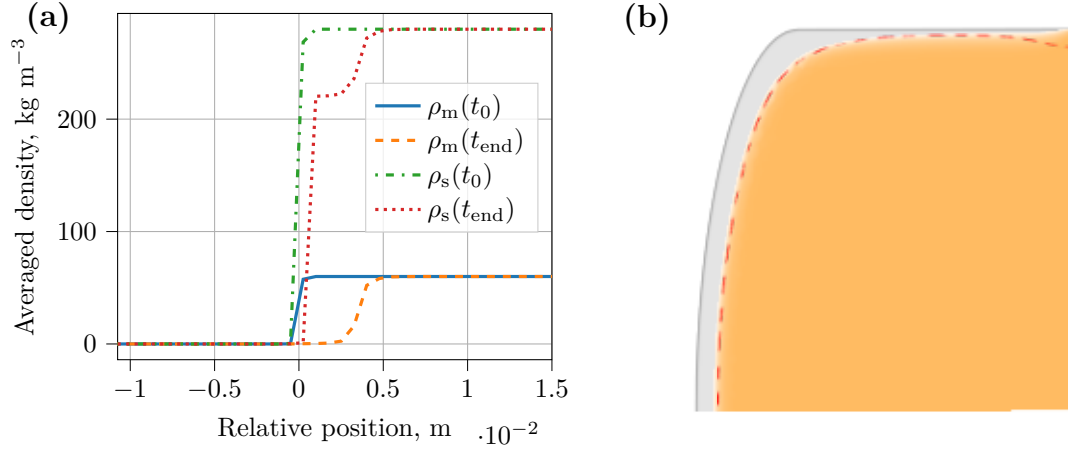


Figure 4.7: (a) Comparison of the initial averaged densities and the averaged densities after 40 s for the total solid and resin matrix densities, and (b) illustration of the overall shape change for the iso-Q sample. In (b) the grey zone represents the char zone and the yellow zone represents the virgin zone. The dotted line represents the pyrolyzing front when the holder is assumed to be adiabatic.

moving and therefore species convection may be weak. Finally, at $t = 40$ s, the partial pressures are almost constant within the material. One can also observe the shift of the production front with time. The higher partial pressure of O close to the surface compared to O_2 , whereas the temperature decreases and should lead to a higher partial pressure of O_2 , is attributed to the oxidation reaction that is slower for O than for O_2 , for which reaction the activation energy is equal to zero.

Streamlines are shown in Fig. 4.11 within and around the sample at $t = 40$ s. Particle trajectories are computed during 10 s using the velocity field at $t = 40$ s using Gmsh on a 20×50 uniform grid representing the particles located around and within the porous material. One can first observe the recirculation region on the top of the material, in which the particles ejected from the side of the material are trapped. On the other side, the particles that are ejected at the front and at the shoulder pass above the recirculation region and are convected toward the outlet. Second, one can observe that the pyrolysis reactions eject the species towards the surface of the material but also inside the material. This pyrolyzing front is characterized by the absence of streamlines below the surface and follows the curvature of the surface and the high pressure distribution at the front of the material (Fig. 4.10). Deeper below the surface, we can note that the particles do not exit the material after the 10 s during which the flow trajectory was simulated for the computation of the streamlines, showing that the particles are moving slowly within the material.

4.4 Conclusion

The unified approach with the VANS equations for solving pyrolysis-ablation problems was verified against a state-of-the-art material response code for the temperature, virgin and char fronts, and mass blowing. The finite-rate chemistry implemented in Argo allows us to track individually the pyrolysis species produced, which was assumed to be a simplified pyrolysis gas made up of seven species. However, a Dirichlet condition was imposed on the species mass fraction at the surface, and more physical boundary conditions should be investigated.

We performed a numerical simulation of an iso-Q sample subjected to high enthalpy condi-

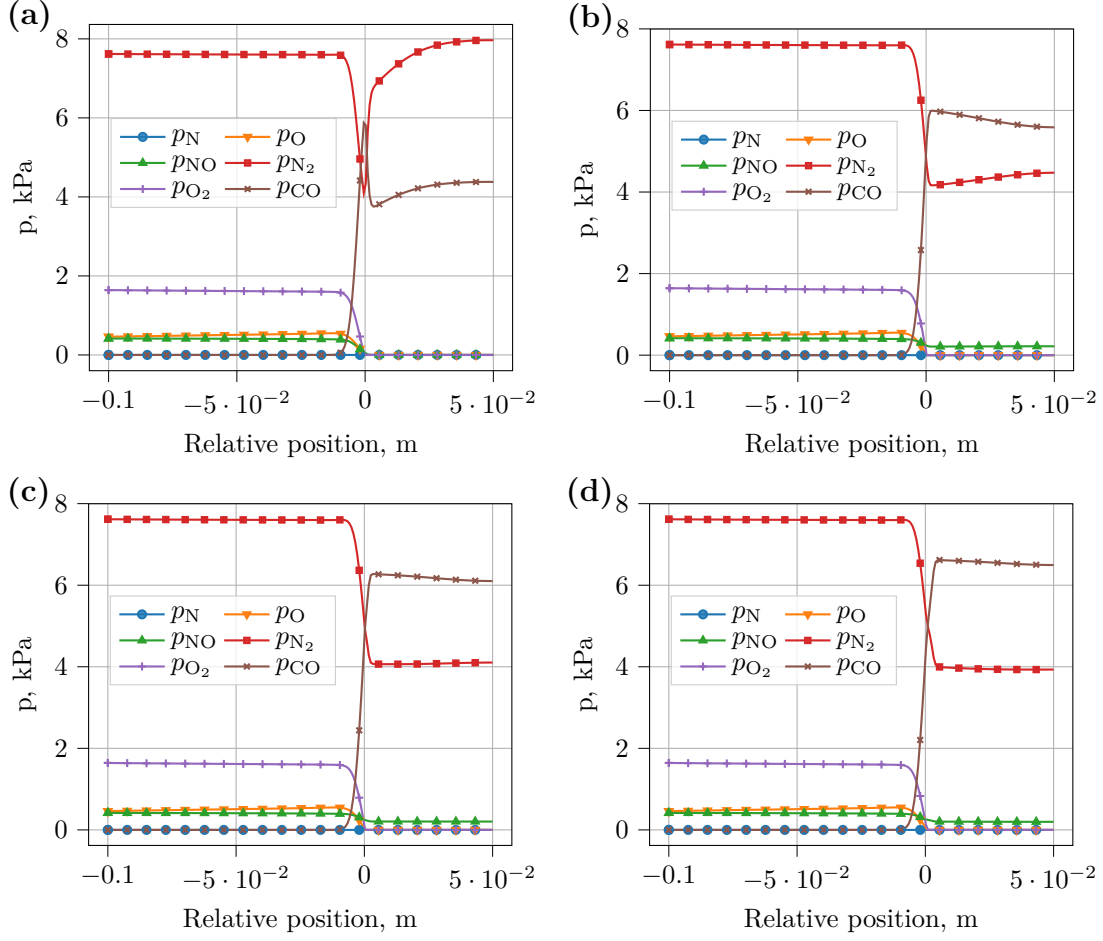


Figure 4.8: Partial pressures for the six-species mixture along the stagnation line at (a) $t = 0.01$ s, (b) $t = 10$ s, (c) $t = 20$ s, and (d) $t = 40$ s.

tions. We showed the recession of the resin and the ablation of the fibers and the char, and a surface recession of $770 \mu\text{m}$ was observed. Close to the surface of the sample, gas species are blown off and are injected into the boundary layer. Below the surface, a stagnation line is observed where the pyrolyzing front is located, and below the pyrolyzing front, the flow is creeping inside the porous material, and the species present within the materials are trapped. Close to the holder, we observed a re-circulation region on the side of the material due to the ejection of ablation material, and where pressure and temperature are lower.

As a future work, a more sophisticated pyrolysis gas than the single CO species should be considered for the iso-Q test case. However, a finite-rate chemistry model as the one implemented within Argo requires to track all the species present in the mixture, hence increasing the computational cost of the simulations. Next, it would be valuable to perform a code-to-code comparison for the iso-Q test case, first in order to corroborate the observations that are made in the present work, and second to emphasize the advantages and disadvantages of the unified approach compared to material-response codes. In order to do that, the heat flux at the surface of the sample have to be similar and, for instance, the surface heat flux could be extracted from Argo and imposed as a boundary condition to material-response codes.




Figure 4.9: Temperature field (K) of the iso-Q test case after 40 s. The temperature ranges from 300 K within the sample to 3000 K at the inlet and in most of the flow field environment. The solid line represents the initial surface of the material.




Figure 4.10: Pressure field (Pa) of the iso-Q test case after 40 s. The solid line represents the initial surface of the material.




Figure 4.11: Streamlines around and within the material.

Part II

Uncertainty Quantification: Theory and Methods

Uncertainty quantification methods for computational models

This chapter is an introduction to the basic concepts and mathematical tools used for uncertainty quantification in computational models. We first provide notions of probability theory and set the mathematical framework for studying the forward and inverse problems in a stochastic framework. Next, we focus on inverse problems and in particular on the Bayesian approach for inverse problems that allows a probability density function for the input parameters to be obtained by means of Bayes' theorem. We introduce an important class of methods based on the sampling of the probability density function, the so-called Markov chain Monte Carlo (MCMC) methods. Then, we present methods for uncertainty propagation for computational models with a focus on non-intrusive methods. We present the Monte Carlo sampling method and finally the method of polynomial chaos, which allows a computationally expensive model involved in a stochastic forward problem to be replaced with a cheaper surrogate model.

5.1 Introduction

Predictions are usually subject to numerous sources of uncertainty. These uncertainties are diverse and can be due to the presence of variability in the parameters of the model (*parameter uncertainty*) either stemming from inherent variability of the system or resulting from a limited accuracy on experimental data, or inaccuracies in the mathematical model (*model structure error*) that is used to represent the physical process; see for instance [Smith \[2014\]](#), [Kennedy and O'Hagan \[2001\]](#) or [Soize \[2017\]](#) for a comprehensive discussion on the classification of uncertainties. These uncertainties in turn introduce uncertainty in the predictions, which needs to be quantified in order to yield useful and reliable information regarding the physical system. The field of uncertainty quantification (UQ) proposes a theoretical framework for studying the uncertainty in a given problem, algorithms to solve these problems and methods to quantify the predicted outcome uncertainty.

There are two categories of problems that can be identified for uncertainty quantification. The first category is related to the propagation of uncertainty in forward problems. The *forward problem* concerns the prediction of the outcome of a physical system described by a mathematical model given the value of its parameters, possibly by means of a computer simulation. The role of uncertainty quantification in forward problems then consists in characterizing in a statistical framework the uncertainty in the output of the model based on the input uncertainties.

The second category of problem investigates the quantification of uncertainty in inverse problems. The *inverse problem* consists in inferring, or deducing, the model that represents

the physical system and the value of its parameters based on the outcome of measurements, or observations, of the physical system that are the results of noise and randomness in the physical process. Inverse problems have long been studied in both the deterministic framework, that is without providing an uncertainty quantified model, and in the stochastic framework. In the deterministic framework, the objective is to find the model and the value of the model parameters that best describe the observations, and to do so, one generally seeks to minimize a misfit function between the predicted outputs and the observations by means of optimization algorithms. However, it is well-known that inferring parameters from indirect and noisy measurements can lead to ill-posed problems, for which regularization techniques are usually applied to approximate the ill-posed problem by a well-posed problem (Engl et al. [1996]). Moreover, deterministic inverse problems usually do not account for the various uncertainties in the parameters stemming from measurement errors, model inadequacy, or parameter variability.

In the stochastic framework, the plausibility of the resulting model and parameter values is represented by means of probability distributions. In particular, the Bayesian approach (see for instance Tarantola [2005]; Stuart [2010]) allows Bayes' theorem to be used to account for parameter uncertainties and attributes a PDF to the parameter estimates resulting from the inference. The Bayesian approach consists in assigning a prior PDF to the parameters by using the current knowledge of the parameters and to update this PDF by incorporating the physical model and the new information (observations) with their uncertainty through a likelihood function. The resulting PDF is the so-called posterior PDF. An introduction to the Bayesian approach applied to engineering problems can be found in Smith [2014] or Kaipio and Somersalo [2005].

To sample from the posterior, the MCMC method is widely used (Gilks et al. [1996]; Robert and Casella [1999]; Smith [2014]; Brooks et al. [2011]). It is a general sampling method that aims at generating random samples from a target probability distribution which, in the context of Bayesian inverse problems, is the Bayesian posterior distribution. For instance, the random-walk Metropolis–Hastings (RWMH) algorithm generates samples from the target distribution by proposing a sample that is drawn from a simpler distribution, called the *proposal distribution* (generally a Gaussian distribution), based on the value of the previous iteration, hence the term *random-walk*, which is then accepted or rejected. The closer this proposal distribution is to the target distribution, the faster the algorithm will generate samples that are actually from the target distribution.

An appropriate choice of the proposal distribution is crucial to achieve an efficient sampling of the target distribution. Finding an optimal or even a suitable proposal distribution might be tedious, in particular when the problem is high dimensional and when trial-and-error tuning of the hyperparameters of the proposal distribution becomes unfeasible. Moreover, a proper choice of the proposal distribution is essential to obtain samples from the posterior distribution in a reasonable amount of time and a limited number of function evaluations. Hence, several improvements of the RWMH exist such as the adaptive random-walk Metropolis–Hastings (AMH) sampling algorithm that adapts the covariance of the Gaussian proposal based on previous iterations to improve the efficiency of the RWMH algorithm (Haario et al. [2001]). However, a complex dependence structure that results in a target distribution that is far from the Gaussian distribution can make the random walk inefficient, especially when the dimension increases. This leads to the use of sampling algorithms for which the choice of the proposal distribution is not based on the Gaussian distribution, as for instance algorithms that use the gradient of the target distribution to propose the new sample values (Neal [2010]; Roberts [2002]; Soize [2008]). The evaluations of the gradients increase the computational time compared to other methods, but it can counterbalance the slow exploration of the posterior, or it can converge when other methods fail. However, even for gradient-informed algorithms, they still require some tuning that can be challenging depending on the structure of the target distribution and the number of parameters to be inferred (Girolami and Calderhead [2011]).

The objective of this chapter is to review essential tools for uncertainty characterization and

propagation in engineering problems. For uncertainty characterization, we focus on Bayesian inference methods and we present general Metropolis–Hastings algorithms as well as advanced sampling algorithms that can be used to solve Bayesian inverse problems. We address the propagation of uncertainty through computational models by means of non-intrusive methods, which require no modification of the computer code, and we describe the Monte Carlo and the polynomial chaos methods.

This chapter is organized as follows. First, in Section 5.2, we start with a brief review of the probability theory. Then in Section 5.3, we introduce the forward and inverse problems. We describe inverse problems in both the deterministic and the stochastic framework, and for the latter, we focus on the Bayesian approach. Section 5.4 is concerned with the description of MCMC methods that are used to generate samples from a probability distribution and we present several algorithms. In Section 5.4.1, the general class of Metropolis–Hastings algorithms is presented and the two subsequent sections are dedicated to the description of the RWMH and the AMH. We then focus on two gradient-informed sampling methods, namely the Hamiltonian Monte Carlo method (Section 5.4.4) and a method based on the solution of an Itô stochastic differential equation (Section 5.4.5). Finally, the propagation of uncertainty in forward problems is described in Section 5.5, with a focus on non-intrusive methods for uncertainty propagation, and in particular the polynomial chaos method.

5.2 Fundamentals of probability theory and statistics

In this section, we summarize concepts from probability theory and statistics. Further details can be found in many statistical textbooks, see for instance Billingsley [2012] or Hogg et al. [2005] for mathematical aspects and a general description of probability theory. For an introduction in the context of uncertainty quantification, see for instance Smith [2014], or Arnst and Ponthot [2014] for a condensed overview. For the concepts related to stochastic processes and stochastic differential equations, see Soize [2017].

5.2.1 Probability space

In order to have a sensible measure of the probability of a real-world situation, or outcome of a random experiment, a probability space needs to be provided. It is defined by the probability triplet $(\Omega, \mathfrak{B}, \mathbb{P})$. In this triplet, Ω represents the sample space or outcome space, that is the totality of the possible outcomes ω of the experiment. Second, \mathfrak{B} is called the event space and contains subsets of the sample space (all the outcomes of the experiment that we would like to consider). An element B of \mathfrak{B} is called an event. The event space \mathfrak{B} is defined as a σ -field, or σ -algebra, that satisfies the following conditions:

1. $\emptyset \in \mathfrak{B}$, (\mathfrak{B} is not empty).
2. If $B \in \mathfrak{B}$, then $B^c \in \mathfrak{B}$, (B is closed under complements).
3. If the sequence of sets $\{B_n, 1 \leq n \leq \nu\}$ is in \mathfrak{B} , then $\bigcup_{n=1}^{\nu} B_i \in \mathfrak{B}$, (\mathfrak{B} is closed under countable unions).

From conditions 1 and 2, we see that $\emptyset^c = \Omega \in \mathfrak{B}$ and the σ -field always contains \emptyset and Ω . Third, $\mathbb{P}(\cdot)$ is the probability measure, that is a real-valued function defined on (Ω, \mathfrak{B}) which satisfies the following three conditions:

1. $\mathbb{P}(B) \geq 0$, for all $B \in \mathfrak{B}$.
2. $\mathbb{P}(\Omega) = 1$.
3. If $\{B_n, 1 \leq n \leq \nu\}$ is a sequence of sets in \mathfrak{B} and $B_m \cap B_n = \emptyset$ for all $m \neq n$, then $\mathbb{P}(\bigcup_{n=1}^{\nu} B_n) = \sum_{n=1}^{\nu} \mathbb{P}(B_n)$.

The probability measure tells us how the probability is distributed over the event space \mathfrak{B} and in this sense we speak of a probability distribution. From the three conditions above, the following noticeable properties can be demonstrated. For each event $B_1, B_2 \in \mathfrak{B}$,

$$\mathbb{P}(B_1 \cup B_2) = \mathbb{P}(B_1) + \mathbb{P}(B_2) - \mathbb{P}(B_1 \cap B_2), \quad (5.1)$$

$$\mathbb{P}(B_1) = 1 - \mathbb{P}(B_1^c), \quad \mathbb{P}(\emptyset) = 0, \quad 0 \leq \mathbb{P}(B_1) \leq 1. \quad (5.2)$$

5.2.2 Conditional probability and Bayes' theorem

Conditional probability

Let \mathbb{P} be the probability measure on (Ω, \mathfrak{B}) and let consider that an event $B_1 \in \mathfrak{B}$ has occurred with $\mathbb{P}(B_1) > 0$. A new probability measure on the sample space B_1 can be expressed for the probability that an event $B_2 \in \mathfrak{B}$ given the event B_1 occurs, defined as

$$\mathbb{P}(B_2|B_1) = \frac{\mathbb{P}(B_1 \cap B_2)}{\mathbb{P}(B_1)}, \quad (5.3)$$

where $B_1 \cap B_2$ denotes the elements of B_2 that are also elements of B_1 . Such a probability measure is called the *conditional probability function* and is denoted by $\mathbb{P}(\cdot|B_1)$.

Bayes' theorem

Let $B \in \mathfrak{B}$ and suppose that $\mathbb{P}(B) > 0$. Consider further a partition of \mathfrak{B} , that is, a set of k mutually exclusive and exhaustive events B_1, \dots, B_k such that $\mathbb{P}(B_i) > 0$, $i = 1, \dots, k$. It can be demonstrated that

$$\mathbb{P}(B) = \sum_{i=1}^k \mathbb{P}(B_i) \mathbb{P}(B|B_i), \quad (5.4)$$

which is known as the *law of total probability*. From the definition of conditional probability and using the law of total probability, we have

$$\mathbb{P}(B_j|B) = \frac{\mathbb{P}(B \cap B_j)}{\mathbb{P}(B)} = \frac{\mathbb{P}(B_j) \mathbb{P}(B|B_j)}{\sum_{i=1}^k \mathbb{P}(B_i) \mathbb{P}(B|B_i)}. \quad (5.5)$$

which is the so-called *Bayes' theorem*. It allows us to calculate the conditional probability of B_j given B from the probabilities of the B_i 's and the conditional probabilities of B given B_i , $i = 1, \dots, k$.

5.2.3 Random vectors and multivariate distributions

Random variables

The sample space Ω is an abstract space for which it may be tedious to describe all the outcomes ω ; instead it is desirable to represent the elements of Ω by real numbers. For a given probability space $(\Omega, \mathfrak{B}, \mathbb{P})$, an \mathbb{R} -valued random variable X is a measurable function from Ω into \mathbb{R} . Thus, X is a mapping from (Ω, \mathfrak{B}) onto $(\mathbb{R}, \mathfrak{D})$ where the sample space of X is the set of real numbers and \mathfrak{D} is a σ -field on \mathbb{R} and is such that $\forall D \in \mathfrak{D}$, $X^{-1}(D) \in \mathfrak{B}$. The probability distribution \mathbb{P}_X of X is defined such that for any element $D \in \mathfrak{D}$, $\mathbb{P}_X(D) = \mathbb{P}(X^{-1}(D)) = \mathbb{P}(\{\omega \in \Omega : X(\omega) \in D\})$, which satisfies the probability measure conditions mentioned previously.

A random variable can be discrete or continuous, depending on whether it takes values from a finite or countable infinite subset of \mathbb{R} or takes values continuously in \mathbb{R} . In the latter, the probability distribution $\mathbb{P}_X(\cdot)$ admits a probability density function, that is a function π_X from \mathbb{R} into $\mathbb{R}^+ = [0, \infty[$ such that $\mathbb{P}_X(D) = \int_D \pi_X(x) dx$ for any element D of \mathfrak{D} .

Various statistical descriptors can be computed from the probability density function π_X . The mean value, or expectation, of X is defined as

$$\begin{aligned}\mathbb{E}[X] &= \int_{\Omega} X(\omega) \mathbb{P}(d\omega) \\ &= \int_{\mathbb{R}} x \mathbb{P}_X(dx) \\ &= \int_{\mathbb{R}} x \pi_X(x) dx \equiv \mu_X,\end{aligned}\tag{5.6}$$

and the variance of X is defined as

$$\mathbb{V}[X] = \int_{\Omega} (X(\omega) - \mu_X)^2 \mathbb{P}(d\omega) \tag{5.7}$$

$$= \int_{\mathbb{R}} (x - \mu_X)^2 \mathbb{P}_X(dx) \tag{5.8}$$

$$= \int_{\mathbb{R}} (x - \mu_X)^2 \pi_X(x) dx \equiv \sigma_X^2, \tag{5.9}$$

assuming that these integrals are bounded. The quantity σ_X is called the *standard deviation* of X .

Random vectors

An \mathbb{R}^d -valued random vector defined on a probability space $(\Omega, \mathfrak{B}, \mathbb{P})$ is the measurable mapping $\omega \rightarrow \mathbf{X}(\omega) = \mathbf{X}$ from (Ω, \mathfrak{B}) into $(\mathbb{R}^d, \mathfrak{D})$ where \mathfrak{D} is a σ -field on \mathbb{R}^d . Its probability distribution is denoted by $\mathbb{P}_{\mathbf{X}}$. A continuous random vector admits a *joint probability density function* $\pi_{\mathbf{X}}$ such that $\mathbb{P}_{\mathbf{X}}(D) = \int_D \pi_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}$ for any meaningful subset D of \mathfrak{D} . In the following, we will denote a random vector using the notation $\mathbf{X} = (X_1, \dots, X_d)$.

Provided the probability distribution $\mathbb{P}_{\mathbf{X}}(\cdot)$, statistical descriptors for a random vector can be computed. The mean value of \mathbf{X} is defined as

$$\mathbb{E}[\mathbf{X}] = (\mathbb{E}[X_1], \mathbb{E}[X_2], \dots, \mathbb{E}[X_d]) \equiv \boldsymbol{\mu}_{\mathbf{X}}. \tag{5.10}$$

The marginal distributions of the random variables X_i , $i = 1, \dots, d$ can be computed from the knowledge of the joint probability distribution. The marginal PDF of the random variable X_i is defined by the $(d-1)$ -fold integral

$$\pi_{X_i}(x_i) = \int_{\mathbb{R}} \dots \int_{\mathbb{R}} \pi_{X_1, \dots, X_d}(x_1, \dots, x_d) dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_d. \tag{5.11}$$

Consider the matrix $(\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}})^T (\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}})$ whose elements are random variables. The mathematical expectation $\mathbb{E}[(\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}})(\mathbf{X} - \boldsymbol{\mu}_{\mathbf{X}})^T] \equiv \Sigma_{\mathbf{X}}$ is called the *covariance matrix* whose elements are computed, for $i \neq j$, as

$$\begin{aligned}(\Sigma_{\mathbf{X}})_{ij} &= \mathbb{E}[(X_i - \mu_{X_i})(X_j - \mu_{X_j})] \\ &= \int_{\mathbb{R}} (x_i - \mu_{X_i})(x_j - \mu_{X_j}) \pi_{X_i, X_j}(x_i, x_j) dx_i dx_j\end{aligned}\tag{5.12}$$

assuming that these integrals are bounded and where π_{X_i, X_j} is the bivariate probability distribution defined by the $(d-2)$ -fold integral of the joint PDF over all variables except X_i and X_j .

The element ij of the covariance matrix is called the covariance of X_i and X_j and the scaled quantity

$$\rho_{X_i X_j} = \frac{\mathbb{E} \left[(X_i - \mu_{X_i}) (X_j - \mu_{X_j}) \right]}{\sigma_{X_i} \sigma_{X_j}} \quad (5.13)$$

is the *correlation coefficient*. The diagonal elements of the covariance matrix $\Sigma_{\mathbf{X}}$ are the variances of X_i (Eq. (5.9)). Finally, we note that the covariance matrix is a symmetric positive semi-definite (PSD) matrix, that is, $\mathbf{a} \Sigma_{\mathbf{X}} \mathbf{a}^T \geq 0$ for all vectors $\mathbf{a} \in \mathbb{R}^d$.

The *joint conditional PDF* of the random vector $(X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_d)$ given X_i is defined as

$$\pi_{X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_d | X_i}(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_d | x_i) = \frac{\pi_{X_1, \dots, X_d}(x_1, \dots, x_d)}{\pi_{X_i}(x_i)}. \quad (5.14)$$

provided that $\pi_{X_i}(x_i) > 0$. The joint conditional PDF of an $(d-1)$ -dimensional random vector is easily extended to the joint conditional PDF of an $(d-k)$ -dimensional random vector. Let \mathbf{X} be an \mathbb{R}^k -valued random vector and \mathbf{Y} an $\mathbb{R}^{(d-k)}$ -valued random vector defined on the same probability space $(\Omega, \mathfrak{B}, \mathbb{P})$ that admit the joint PDF $\pi_{\mathbf{X}, \mathbf{Y}}$. Suppose that the marginal distribution $\pi_{\mathbf{Y}}(\mathbf{y}) > 0$. The joint conditional distribution of \mathbf{X} given \mathbf{Y} can be obtained from

$$\pi_{\mathbf{X} | \mathbf{Y}}(\mathbf{x} | \mathbf{y}) = \frac{\pi_{\mathbf{X}, \mathbf{Y}}(\mathbf{x}, \mathbf{y})}{\pi_{\mathbf{Y}}(\mathbf{y})}. \quad (5.15)$$

5.2.4 Stochastic processes

A stochastic process, or random process, defined on a probability space $(\Omega, \mathfrak{B}, \mathbb{P})$ indexed by an infinite subset \mathbb{T} of \mathbb{R} is a uncountable family

$$\mathbf{X} = \{\mathbf{X}(t), t \in \mathbb{T}\} = \{\mathbf{X}^t, t \in \mathbb{T}\} \quad (5.16)$$

of \mathbb{R}^d -valued random vectors $\mathbf{X}(t)$ defined on a probability space $(\Omega, \mathfrak{B}, \mathbb{P})$. Taking \mathbb{T} an interval of real numbers yields a continuous process, whereas taking \mathbb{T} a subset of integers yields a discrete random process.

The realization $x(t) = \mathbf{X}^t(\omega) = \mathbf{X}(t, \omega)$ is a function of t that is called the *sample path* or *trajectory* associated with the outcome $\omega \in \Omega$.

5.2.5 Markov chains

An \mathbb{R}^d -valued Markov chain is a discrete random process, that is an ordered set $\{\mathbf{X}^n, n \in \mathbb{N}\}$ of random vectors $\mathbf{X}^0, \mathbf{X}^1, \mathbf{X}^2, \dots$ defined on $(\Omega, \mathfrak{B}, \mathbb{P})$ with values in \mathbb{R}^d which satisfies the Markov property. Denoting \mathbf{x}^n the current realization of the random variable \mathbf{X}^n at the *iteration*, or time, n , and \mathbf{x}^{n+1} the next state, the Markov property reads as:

$$\mathbb{P}_{\mathbf{X}}(\mathbf{X}^{n+1} \in D | \mathbf{X}^0 = \mathbf{x}^0, \dots, \mathbf{X}^n = \mathbf{x}^n) = \mathbb{P}_{\mathbf{X}}(\mathbf{X}^{n+1} \in D | \mathbf{X}^n = \mathbf{x}^n). \quad (5.17)$$

that is, every new sample \mathbf{X}^{n+1} is drawn from a probability distribution that depends only on the current state \mathbf{X}^n and not on the previous states $\mathbf{X}^{n-1}, \dots, \mathbf{X}^0$. The probability $\mathbb{P}_{\mathbf{X}}(\mathbf{X}^{n+1} \in D | \mathbf{X}^n = \mathbf{x}^n)$ of going from one state to the other is called the *transition probability*, or the *Markov kernel*. Thus, if the initial state is known, the construction of the Markov chain $\{\mathbf{X}^n, n \in \mathbb{N}\}$ depends only on the transition probability. A Markov chain is said to be *homogeneous* if the transition probability is homogeneous, which is written, for \mathbf{x}^j , for $D \in \mathfrak{D}$, and for $0 \leq j \leq n$, as

$$\mathbb{P}_{\mathbf{X}}(\mathbf{X}^{n+1} \in D | \mathbf{X}^j = \mathbf{x}^j) = \mathbb{P}_{\mathbf{X}}(\mathbf{X}^{n+1-j} \in D | \mathbf{X}^j = \mathbf{x}^j). \quad (5.18)$$

The conditional density of the transition, the so-called the *transition kernel*, K reads as

$$\mathbb{P}_{\mathbf{X}}(\mathbf{X}^{n+1} \in D | \mathbf{X}^n = \mathbf{x}^n) = \int_D K(\mathbf{x}^n, \mathbf{x}') d\mathbf{x}' \quad (5.19)$$

and is such that $\forall \mathbf{x} \in \mathbb{R}^d$, $K(\mathbf{x}, \cdot)$ is a probability measure, and $\forall D \in \mathfrak{D}$, $K(\cdot, D)$ is measurable from \mathbb{R}^d into $[0, 1]$. An *invariant distribution* for the transition kernel K is a probability distribution π that satisfies

$$\pi_{\mathbf{X}}(D) = \int_{\mathbb{R}^d} K(x, D) \pi_{\mathbf{X}}(dx), \quad \forall D \in \mathfrak{D} \quad (5.20)$$

For homogeneous Markov chains, the invariant distribution is also referred to as *stationary*, since $\mathbf{X}^n \sim \pi_{\mathbf{X}}$ implies that $\mathbf{X}^{n+1} \sim \pi_{\mathbf{X}}$. The existence and uniqueness of a stationary distribution will be particularly interesting for MCMC algorithms. A sufficient condition for π to be the stationary distribution associated with the transition kernel K is if it satisfies the *detailed balance condition*:

$$K(y, x) \pi_{\mathbf{X}}(y) = K(x, y) \pi_{\mathbf{X}}(x), \quad (5.21)$$

for every (x, y) . Furthermore, for MCMC algorithms, the convergence in distribution of the Markov chain towards the stationary distribution $\pi_{\mathbf{X}}$ will be of interest, and this usually requires for the chain to be *irreducible* (all states “communicate”) and *aperiodic* (there is no cycle in the chain). Hence, it can be proven that ultimately, samples will forget their initial state \mathbf{X}^0 and will eventually converge to the unique *stationary distribution*. Some of the transition kernels for MCMC algorithms will be presented in Section 5.4.1. For a more comprehensive discussion on Markov chains, see Gilks et al. [1996] or Robert and Casella [2004].

5.2.6 Distinction between random and stochastic differential equations

Random differential equation

Random differential equations are differential equations in which the random effects are manifested in parameters, initial or boundary conditions, or forcing conditions that are regular (e.g. continuous) with respect to space and time. As an example, let us consider the following one-dimensional ordinary differential equation (ODE) that describes the evolution of a random process X defined on a probability space $(\Omega, \mathfrak{B}, \mathbb{P})$:

$$\frac{dx}{dt}(t) = a_1 x(t) + a_2(t), \quad (5.22)$$

with the initial condition $x(0) = X^0 = X(0, \omega)$, where $a_1 = A_1(\omega)$ is the realization of a random variable A_1 and $a_2(t) = A_2(t, \omega)$ is the realization of a random process B for an outcome $\omega \in \Omega$. Hence, for a given realization of ω , the random differential equation can be analyzed and solved by means of standard resolution techniques if $A_2(t, \omega)$ is taken to be smooth.

For instance, considering the Arrhenius parameters to be random variables, the pyrolysis equations described in Section 3.3.2 are random differential equations. For a given realization of the random parameters, they can be solved by means of numerical integration techniques or by computing the solution presented in Section 3.3.3.

Stochastic differential equation and Itô stochastic differential equation

As opposed to random differential equations, stochastic differential equations (SDEs) describe phenomena that are forced by an irregular process. For instance, the random motion of a particle in a liquid (Brownian motion) can be described by the so-called Wiener process, a continuous stochastic process representing the random effect of thermal fluctuations of the liquid

molecules on the particle. An SDE of particular interest is the Itô stochastic differential equation (ISDE), which solution can be used for defining a transition kernel in MCMC algorithms (see Section 5.4.5). It is defined as the following SDE:

$$d\mathbf{X}(t) = \mathbf{b}(\mathbf{X}(t), t)dt + a(\mathbf{X}(t), t)d\mathbf{W}(t), \quad t \in]0, t_1], \quad (5.23)$$

with the initial condition $\mathbf{X}(0) = \mathbf{X}^0$, in which $\{\mathbf{X}(t), t \in \mathbb{T}\}$ is a stochastic process indexed by $\mathbb{T} = [0, t_1]$. The function $(\mathbf{x}, t) \rightarrow \mathbf{b}(\mathbf{x}, t)$ is a function from $\mathbb{R}^d \times \mathbb{T}$ into \mathbb{R}^d and $(\mathbf{x}, t) \rightarrow a(\mathbf{x}, t)$ is a function from $\mathbb{R}^d \times \mathbb{T}$ into $\mathbb{R}^{d \times d}$. The \mathbb{R}^d -valued stochastic process $\{\mathbf{W}(t), t \in \mathbb{R}^+\}$, with $\mathbf{W}(0) = \mathbf{0}$, is the so-called Wiener process, that is, a process such that for all $0 \leq s < t < u < v$, the increments $\mathbf{W}(t) - \mathbf{W}(s)$ and $\mathbf{W}(v) - \mathbf{W}(u)$ are statistically independent; and for all $0 \leq s < t < +\infty$, the increment $\Delta \mathbf{W}_{st} = \mathbf{W}(t) - \mathbf{W}(s)$ is a \mathbb{R}^d -valued Gaussian random variable with mean zero and covariance matrix $C_{\Delta \mathbf{W}_{st}} = (t - s)I_d$. Hence, the irregular nature of $\mathbf{X}(t)$ are due to the fluctuations of the Wiener process \mathbf{W} .

The solution of the ISDE can be written as

$$\mathbf{X}(t) = \int_0^t \mathbf{b}(\mathbf{X}(t), t)dt + \int_0^t a(\mathbf{X}(t), t)d\mathbf{W}(t), \quad (5.24)$$

where the second integral on the right-hand side is not interpreted in terms of a classical integral but as an *Itô stochastic integral*. Because of the nature of the Wiener process, \mathbf{W} is nowhere differentiable and the continuous trajectories of \mathbf{W} are not of bounded variation on any bounded interval (Soize [2017]). The Itô stochastic integral allows to extend the definition of integrals in the case where the integrand and the integrator are stochastic processes. More details on stochastic calculus and ISDE can be found in Soize [2017].

5.3 Forward and inverse problems

5.3.1 Forward problem

The forward problem consists in predicting the solution of a physical system given a computational model, which describes the system, and its parameters.

In this work, we will mainly consider systems that can be described by PDEs defined in a spatial domain $D \subset \mathbb{R}^l$, $l = 0, 1, 2, 3$, and a time domain $[0, T]$ with $T \geq 0$, as follows:

$$\begin{cases} \frac{\partial \mathbf{u}}{\partial t}(\mathbf{x}, t, \mathbf{m}) = \mathcal{L}(\mathbf{u}), & D \times]0, T] \times \mathbb{R}^{n_u}, \\ B(\mathbf{u}) = 0, & \partial D \times [0, T] \times \mathbb{R}^{n_u}, \\ \mathbf{u} = \mathbf{u}_0, & D \times \{t = 0\} \times \mathbb{R}^{n_u}, \end{cases} \quad (5.25)$$

where \mathcal{L} is a (nonlinear) differential operator, $B(\mathbf{u})$ represents the values of \mathbf{u} at the boundaries ∂D of the domain D , \mathbf{u}_0 the initial condition and \mathbf{m} the inputs, or the parameters, of the system. The solution is the quantity $\mathbf{u}(\mathbf{x}, t, \mathbf{m}) : \bar{D} \times [0, T] \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_d}$, where n_d is the dimension of \mathbf{u} .

It happens that the solution to the governing PDEs for \mathbf{u} cannot be observed directly through the experiment, for instance when the parameters are tuning parameters of the model, or when direct observation is not feasible. Observations of the physical process, denoted by \mathbf{y} , are linked to the state of the system \mathbf{u} through the relation

$$\mathbf{y} = \boldsymbol{\eta}(\mathbf{x}, t, \mathbf{m}) = \mathcal{L}_2(\mathbf{u}(\mathbf{x}, t, \mathbf{m}), \mathbf{m}), \quad (5.26)$$

where \mathcal{L}_2 is a (nonlinear) observation operator that may also depend on the parameters. The function $\boldsymbol{\eta}$ represents the computer model, that is, all the algorithms and the equations used to represent the observations.

When it comes to comparing the predicted outcome of the computer model with real-world experiments, there are inherent deviations from the physical process due to the inadequacy of the model and observation errors from the measurement instruments. Hence, a more reliable description of a real-world experiment can be represented by the relation (Kennedy and O'Hagan [2001])

$$y_i^{\text{obs}} = \zeta(\mathbf{x}_i) + \varepsilon_i = \rho\eta(\mathbf{x}_i, \mathbf{m}) + \delta(\mathbf{x}_i) + \varepsilon_i, \quad (5.27)$$

where y_i^{obs} denotes the observed data, $\zeta(\cdot)$ denotes the true process, ε_i denotes the observation error for the i -th observation, and $\delta(\cdot)$ is the model inadequacy (or the model structure error) that is independent of the computer model output $\eta(\mathbf{x}_i, \mathbf{m})$ that depends on the parameters and the variables inputs, denoted by $\mathbf{x} = (\mathbf{x}, t)$.

Assuming that there is no model error ($\delta = 0$), the observed data are obtained from the computer model and the observation errors as

$$y_i^{\text{obs}} = \eta(\mathbf{x}_i, \mathbf{m}) + \varepsilon_i. \quad (5.28)$$

Assuming further that there are no observation errors, the model would correctly replicate the experimental data. In practice however, observation errors are inherent to the measurement process and a mismatch between the model output and the data is generally observed.

5.3.2 Inverse problem

There are two distinct groups of inputs in the computer model η . The \mathbf{x}_i are called the variable inputs whose value might change during the calibration process and that we assume here to be known. The quantity \mathbf{m} represents the calibration parameters, the ones that we wish to learn about. Finding the value of the calibration parameters \mathbf{m} in the absence of model error (Eq. (5.28)) is sometimes referred to as the classical inverse problem. The designation “classical” inverse problem is used to differentiate with the “extended” inverse problem (Sun and Sun [2015]). In the extended inverse problem, the purpose is to find the model and its parameters that replicate the data in the case where the model itself is unknown and there exist several plausible alternative models. More details about the extended inverse problem can be found in Burnham and Anderson [2002]; MacKay [2003] or Sun and Sun [2015].

When the objective is to quantify the model error δ separately from the observation error ε , the statistical model from Eq. (5.27) can be used. However, in practice, the modeling and the measurement errors can be collectively quantified by the quantity ε_i along with the use of the statistical model in Eq. (5.28) (Smith [2014]). Including the modeling errors in the term ε_i will be discussed further in Chapter 7.

Inverse problems can be studied both in the deterministic and the Bayesian frameworks. In the deterministic framework, the theory is concerned with how to define and find the best approximate solution, and make sure that this solution is unique and stable. The principal limitation of this approach is that it only identifies a single “best” model and yields point estimates of the underlying parameter values, without providing a meaningful description of the uncertainty among competing models and their parameters (Galagali and Marzouk [2015]).

In the Bayesian framework, the parameters to estimate are regarded rather as a random vector \mathbf{M} with realization $\mathbf{m} = \mathbf{M}(\omega)$ and associated probability density function. The experimental observations are merely the consequence of a particular realization of the random vector. The theory of Bayesian inverse problems uses the solution of the computer model, the experimental observations and the *a priori* knowledge of the model parameters in order to obtain the probability distribution of the random vector representing the parameters

Deterministic inverse problem

In the deterministic approach, one assumes that the parameters are fixed but unknown. There exists a parameter value \mathbf{m}_0 which is a deterministic quantity that describes best the experi-

mental data, thus $\eta(\mathbf{x}_i, \mathbf{m}_0)$ is also deterministic. However, because of the observations errors, one can only have an estimation of the best-fit parameter value. Indeed, considering that the measurement errors ε_i are the realizations of the random vector $\mathbf{E} \rightarrow \mathbf{E}(\omega)$, the observations are also random variables $\mathbf{Y}^{\text{obs}} \rightarrow \mathbf{Y}^{\text{obs}}(\omega)$, $\omega \in \Omega$, that are related through the statistical model

$$Y_i^{\text{obs}} = \eta(\mathbf{x}_i, \mathbf{m}_0) + E_i. \quad (5.29)$$

Deterministic inverse problems usually seek to minimize a misfit function and optimization algorithms are used to find the solution that best reproduces the data (Nocedal and Wright [1999]). The misfit function is a measure of the mismatch between the model data and the observed data. Denoting $\boldsymbol{\eta}(\mathbf{m}) = (\eta(\mathbf{x}_1, \mathbf{m}), \dots, \eta(\mathbf{x}_n, \mathbf{m}))$ an \mathbb{R}^n -valued vector that represents the output of the computer model for a given value of the parameters, the misfit function can be written as

$$S(\mathbf{m}) = \|\boldsymbol{\eta}(\mathbf{m}) - \mathbf{y}^{\text{obs}}\|^2, \quad (5.30)$$

where the operator $\|\cdot\|$ denotes a given norm. The deterministic inverse problem consists in finding the best solution that minimizes this misfit function:

$$\hat{\mathbf{m}} = \arg \min_{\mathbf{m} \in \mathcal{P}_{\text{ad}}} S(\mathbf{m}), \quad (5.31)$$

where $\hat{\mathbf{m}}$ is the parameter values for which the model should fit the observed data as close as possible. For instance, the ordinary least-squares (OLS) estimate uses the Euclidean norm (or L^2 norm)

$$\hat{\mathbf{m}}_{\text{OLS}} = \arg \min_{\mathbf{m} \in \mathcal{P}_{\text{ad}}} \sum_{i=1}^n (\eta(\mathbf{x}_i, \mathbf{m}) - y_i^{\text{obs}})^2, \quad (5.32)$$

which minimizes the distance between all the data and the computer model outputs with equal importance.

Because \mathbf{Y}^{obs} is a random vector, the value of $\hat{\mathbf{m}}_{\text{OLS}}$ obtained from Eq. (5.32) may be different for each repetition of the experiment. The OLS estimate from Eq. (5.32) is a particular realization of the OLS estimator \mathbf{M}_{OLS} , which is a random vector. Providing a large number N of realizations of the experiments and corresponding estimates $\hat{\mathbf{m}}_{\text{OLS},k}$, $1 \leq k \leq N$, one could be able to compute the error-free best-fit parameter from

$$\mathbb{E}[\mathbf{M}_{\text{OLS}}] = \mathbf{m}_0 \approx \frac{1}{N} \sum_{k=1}^N \hat{\mathbf{m}}_{\text{OLS},k}, \quad (5.33)$$

which comes from the law of large numbers (LLN). This interpretation of the optimal parameter in terms of frequency of occurrence is what is referred to as the *frequentist perspective*. Thus, the computation of the mean and the variance of \mathbf{M}_{OLS} , or computing confidence intervals, provide information on the true parameter. From a frequentist perspective, a $(1 - \alpha) \times 100$ % confidence intervals has the following interpretation: in repeated procedures, $(1 - \alpha) \times 100$ % of realized intervals would include the optimal parameter. For linear computer models with statistically independent Gaussian errors, the best-fit parameter actually corresponds to the expectation of \mathbf{M}_{OLS} .

Another estimator is the maximum likelihood estimator, that is the value that maximizes the likelihood function $L(\mathbf{m}|\mathbf{y}^{\text{obs}})$. The likelihood function is denoted in a similar way to as a conditional probability but is not, as \mathbf{m} is not a random vector. Then, the optimal parameter can be estimated by the maximum likelihood estimate (MLE)

$$\hat{\mathbf{m}}_{\text{MLE}} = \arg \max_{\mathbf{m} \in \mathcal{P}_{\text{ad}}} L(\mathbf{m}|\mathbf{y}^{\text{obs}}), \quad (5.34)$$

in which case the optimal parameter value is interpreted as the parameter value that makes the observed output the most likely.

Because parameters are fixed, but unknown, results of the deterministic inverse problem cannot be directly applied to obtain parameter PDFs that can be propagated through models to quantify uncertainty. In the Bayesian framework, parameters are considered to be random variables with associated densities, thus lending themselves naturally to uncertainty quantification.

Bayesian inverse problem

Whereas the deterministic inverse problem is concerned with finding a single value of the parameters that best fits the observations, the perspective in the Bayesian inverse problem is rather to figure out *what is the state of information about the parameters*.

In the Bayesian approach, what differs significantly from the frequentist perspective is that the parameters are also considered to be a random vector. Therefore, from a Bayesian perspective, there is no fundamental distinction between the observables and the parameters of a statistical model: all are considered as random quantities. Let $\mathbf{M} = (M_1, \dots, M_{n_u})$ be the random vector of parameters with values in \mathbb{R}^{n_u} and $\mathbf{Y}^{\text{obs}} = (Y_1^{\text{obs}}, \dots, Y_{n_{\text{obs}}}^{\text{obs}})$ the random vector with values in $\mathbb{R}^{n_{\text{obs}}}$ defined on the same probability space $(\Omega, \mathfrak{B}, \mathbb{P})$. In the context of parameter inference, the Bayesian approach is based on Bayes' theorem Eq. (5.5) and using conditional distribution yields the following relation:

$$\pi_{\mathbf{M}|\mathbf{Y}^{\text{obs}}}(\mathbf{m}|\mathbf{y}^{\text{obs}}) = \frac{\pi_{\mathbf{Y}^{\text{obs}}|\mathbf{M}}(\mathbf{y}^{\text{obs}}|\mathbf{m})\pi_{\mathbf{M}}^0(\mathbf{m})}{\int_{\mathbb{R}^{n_u}} \pi_{\mathbf{Y}^{\text{obs}}|\mathbf{M}}(\mathbf{y}^{\text{obs}}|\mathbf{m})\pi_{\mathbf{M}}^0(\mathbf{x})d\mathbf{x}}, \quad (5.35)$$

where $\pi_{\mathbf{M}}^0(\mathbf{m})$ is the prior probability density function, $\pi_{\mathbf{M}|\mathbf{Y}^{\text{obs}}}(\mathbf{m}|\mathbf{y}^{\text{obs}})$ is the posterior probability density function, and $\pi_{\mathbf{Y}^{\text{obs}}|\mathbf{M}}(\mathbf{y}^{\text{obs}}|\mathbf{m})$ is the likelihood function. The denominator represents the marginal distribution for the observations and is a normalizing constant. In the following, the subscript \mathbf{Y}^{obs} and \mathbf{M} will be dropped off for the sake of clarity.

Tarantola [2005] derives the relation from Eq. (5.35) by using the general notion of the “combination of states of information” rather than using conditional densities. The likelihood represents the information on the physical correlations between \mathbf{y}^{obs} and \mathbf{m} , as obtained from the physical law. It gives a measure of how good the model with parameters \mathbf{m} is in explaining the observations \mathbf{y}^{obs} . The prior probability density represents both information obtained on the data and on the model parameters. These two states of information combine to produce the a posteriori state of information.

The prior density should compile all the information we know about the parameters before any computation. It can stem from physical consideration (e.g., a mass can not be negative), or from expert knowledge about the parameters. When only little or no information is available a priori, it is possible to specify noninformative priors with the desire that they play a minimal role on the resulting posterior. Reference Gelman et al. [2014] defines weakly informative priors that are allowed to contain enough information to keep the posterior within reasonable bounds while still playing a minimal role in the inference process and being proper, in the sense that they integrate to one. A posterior built with an improper prior can still be proper and have a finite integral, which could be the case for instance when the likelihood function dominates the prior (Gelman et al. [2014]). For instance, we can express that the variable must be positive or not exceed a given limit by providing an equiprobable density function whose support is the parameter set. It is improper in the sense that its value does not integrate to 1 and this will require to check numerically that the resulting posterior has a finite integral.

The posterior density is what we know about the parameters after observing the data and is what we want to know. The solution of the inverse problem is the posterior probability distribution function $\pi(\mathbf{m}|\mathbf{y}^{\text{obs}})$ and this principle makes the Bayesian approach quite different from the deterministic approach.

The most general solution of the Bayesian inverse problem provides a probability density over the model parameter set, and its computation requires the evaluation of a multidimensional integral. It is only when the probability density is very simple (for instance, when it has only one maximum) and very low dimensional that analytical techniques can be used to characterize it. For more general probability densities, one needs to perform an extensive exploration of the parameter set. This can be performed numerically by means of quadrature rules or random (stochastic) explorations. However, when the dimension increases, quadrature rules suffer from the curse of dimensionality and the number of model evaluations becomes intractable. Therefore, random explorations are needed, which also offer two additional advantages: first, they provide samples directly drawn from the posterior distribution that can be used to estimate statistics of the posterior distribution (by Monte Carlo integration), and second, these samples can be used for propagation through numerical solvers for uncertainty quantification. Random explorations can solve many complex problems and a type of methods that explore efficiently the posterior distribution are the MCMC methods. The problem of exploring the posterior PDF by means of MCMC methods is the topic of the next section.

5.4 Markov chain Monte Carlo methods

The principle underlying the MCMC methods is the construction of a Markov chain whose stationary distribution is the probability distribution we wish to sample from. In the context of Bayesian inverse problems, this probability distribution will be the posterior distribution $\pi(\mathbf{m}|\mathbf{y}^{\text{obs}})$. Implementing MCMC algorithms is the subject of many textbooks (Gilks et al. [1996]; Brooks et al. [2011]; Smith [2014]; Gelman et al. [2014]) and we review here the general Metropolis–Hastings algorithm but also advanced versions of sampling algorithms. The algorithms are described in the case where the distribution is the Bayesian posterior $\pi(\mathbf{m}|\mathbf{y}^{\text{obs}})$ but they can be applied to other probability distribution in non-Bayesian frameworks.

5.4.1 The Metropolis–Hastings algorithm

Let $\mathbf{M} = \{\mathbf{M}^n, n \in \mathbb{N}\}$ be a Markov chain with realizations denoted by $\mathbf{m}^n = \mathbf{M}^n(\omega)$ for $\omega \in \Omega$ defined on a probability space $(\Omega, \mathfrak{B}, \mathbb{P})$ and with values in \mathbb{R}^{n_u} . The principle for constructing a Markov chain whose stationary distribution is the distribution of interest (e.g., the posterior distribution in the Bayesian framework) is relatively simple. The concept was first proposed by Metropolis et al. [1953] and later generalized by Hastings [1970]. The Metropolis–Hastings algorithm consists in sampling a candidate value \mathbf{m}^* from a simpler distribution $J(\cdot|\mathbf{m}^n)$ called the proposal distribution that may depend on the current state \mathbf{m}^n , with n the current iteration. The candidate value is accepted with probability $\alpha = \min(1, r)$ where r is the acceptance ratio

$$r = \frac{\pi(\mathbf{m}^*|\mathbf{y}^{\text{obs}}) J(\mathbf{m}^n|\mathbf{m}^*)}{\pi(\mathbf{m}^n|\mathbf{y}^{\text{obs}}) J(\mathbf{m}^*|\mathbf{m}^n)}. \quad (5.36)$$

If the candidate point is accepted, the new state is $\mathbf{m}^{n+1} = \mathbf{m}^*$, and if it is rejected, the state of the chain remains at the same value, i.e. $\mathbf{m}^{n+1} = \mathbf{m}^n$. The transition kernel hence obtained from the Metropolis–Hastings algorithm satisfies the detailed balance equation from Eq. (5.21) with π the stationary distribution, and the chain is ergodic, aperiodic and irreducible (see details in Robert and Casella [2004]). Therefore, once a sample from the stationary distribution is obtained, the next samples will also be from the stationary distribution and the chain is said to have *converged*. The collected samples $\{\mathbf{m}^n\}$ fully characterize the posterior distribution. The Metropolis–Hastings algorithm for the sampling of a multivariate distribution is summarized in Algorithm 1.

In Eq. (5.36), the posterior density can be replaced by the product of the likelihood with the

prior distribution from Eq. (5.35),

$$r = \frac{\pi(\mathbf{y}^{\text{obs}}|\mathbf{m}^*)\pi^0(\mathbf{m}^*)J(\mathbf{m}^*|\mathbf{m}^n)}{\pi(\mathbf{y}^{\text{obs}}|\mathbf{m}^n)\pi^0(\mathbf{m}^n)J(\mathbf{m}^n|\mathbf{m}^*)}. \quad (5.37)$$

By forming the ratio of the posterior densities, the evaluation of the normalization constant is thus avoided.

Algorithm 1: Metropolis–Hastings algorithm.

Result: Parameter values $\mathbf{m}^n \forall n \in [0, n_{\text{iterations}}]$ sampled from the posterior $\pi(\cdot|\mathbf{y}^{\text{obs}})$.
 Initialise \mathbf{m}^0 ;
while $n < n_{\text{iterations}}$ **do**
 Draw new sample $\mathbf{m}^* \sim J(\mathbf{m}^*|\mathbf{m}^n)$;
 Compute acceptance probability $\alpha = \min(1, r)$ with r from Eq. (5.36) and sample
 $u \sim U(0, 1)$;
 if $u \leq \alpha$ **then**
 $\mathbf{m}^{n+1} = \mathbf{m}^*$;
 else
 $\mathbf{m}^{n+1} = \mathbf{m}^n$;
 end
end

5.4.2 Random-walk Metropolis–Hastings algorithm

In the Metropolis algorithm, the proposal distribution is taken to be symmetric, i.e., $J(\mathbf{m}^n|\mathbf{m}^*) = J(\mathbf{m}^*|\mathbf{m}^n)$. When the proposal distribution is of the form $J(\mathbf{m}^*|\mathbf{m}^n) = J(|\mathbf{m}^* - \mathbf{m}^n|)$, the algorithm is called the random-walk Metropolis–Hastings (RWMH). A popular choice for the proposal distribution is the multivariate Gaussian distribution centered on \mathbf{m}^n with covariance matrix Σ of dimensions $n_u \times n_u$, thus,

$$J(\mathbf{m}^*|\mathbf{m}^n) = \frac{1}{\sqrt{(2\pi)^{n_u} \det(\Sigma)}} \exp\left(-\frac{1}{2}(\mathbf{m}^* - \mathbf{m}^n)^T \Sigma^{-1}(\mathbf{m}^* - \mathbf{m}^n)\right), \quad (5.38)$$

$$\text{with } \Sigma = \begin{pmatrix} \sigma_1^2 & \rho_{12}\sigma_1\sigma_2 & \cdots & \rho_{1n_u}\sigma_1\sigma_{n_u} \\ \rho_{21}\sigma_2\sigma_1 & \sigma_2^2 & \rho_{2n_u}\sigma_2\sigma_{n_u} & \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{n_u1}\sigma_{n_u}\sigma_1 & \rho_{n_u2}\sigma_{n_u}\sigma_2 & \cdots & \sigma_{n_u}^2 \end{pmatrix}, \quad (5.39)$$

where σ is the standard deviation of the i -th variable and ρ_{ij} the correlation coefficient between the i -th and the j -th variables. It is worth noting that in the case where we assume a uniform prior distribution, the acceptance ratio further simplifies into

$$r = \frac{\pi(\mathbf{y}^{\text{obs}}|\mathbf{m}^*)}{\pi(\mathbf{y}^{\text{obs}}|\mathbf{m}^n)}, \quad (5.40)$$

and we only need to sample from J for determining the proposed value of the sample \mathbf{m}^* and compute the value of the posterior PDF at \mathbf{m}^* to obtain r .

The choice of the proposal distribution J strongly influences the rate of convergence of the Markov chain towards the posterior distribution. The probability of having large jumps decreases with the distance to the current state and depends on the proposal covariance matrix, which must be chosen carefully. On the one hand, a covariance matrix that is too large, in the sense that the principal dimensions (eigenvectors) have a large magnitude (eigenvalues), will result in

an acceptance ratio that is low and the proposed samples will be mainly rejected. On the other hand, a small covariance matrix will result in small jumps between the successive sample values and a high acceptance rate; thus a very slow exploration of the posterior distribution will be observed, that we refer to as a slow mixing of the Markov chain. In both cases, this results in successive samples that are highly correlated, which is not desirable for Monte Carlo estimations.

There are many other simulation algorithms that are based on the generic Metropolis–Hastings algorithm, such as the slice sampler, the Gibbs sampler, that are not reviewed here. See [Gilks et al. \[1996\]](#) or [Robert and Casella \[1999\]](#) for a general overview of simulation methods based on Markov chains.

5.4.3 Adaptive Metropolis–Hastings algorithm

In order to avoid trial-and-error tuning of the covariance matrix (Eq. (5.39)) and improve the mixing of the Markov chain, methods that learn from past sample values have been proposed and are referred to as *adaptive algorithms*. As a consequence of being dependent on previous iterates, adaptive algorithms are by essence non-Markovian and the challenge in developing such algorithms requires to show that the chain has the correct ergodic properties (converges toward the right distribution and is stationary). For the adaptive random-walk Metropolis–Hastings (AMH) algorithm, [Haario et al. \[2001\]](#) showed that by ensuring that the target density is bounded from above and has a bounded support, the AMH process does indeed have the correct ergodicity properties. The latter assumption was further relaxed by [Atchadé and Rosenthal \[2005\]](#) and [Andrieu and Moulines \[2006\]](#). [Atchadé and Rosenthal \[2005\]](#) showed that adaptive algorithms achieve the same mixing time as an optimally scaled algorithm without requiring all the preliminary efforts to manually tune the scaling parameter of the proposal. This feature is important when little prior information is available about the scales and correlations of the parameters. A slightly different version is proposed by [Roberts and Rosenthal \[2009\]](#) where the proposal is a mixture distribution. [Rosenthal \[2010\]](#) has summarized the most important advancements in AMH and [Andrieu and Thoms \[2008\]](#) have reviewed several AMH algorithms with examples.

We summarize below the adaptive procedure initially proposed by [Haario et al. \[2001\]](#) for the random-walk Metropolis–Hastings algorithm. Starting from the RWMH, the AMH algorithm consists in taking the proposal covariance matrix as the (scaled) covariance matrix of the previous samples. Let denote by \hat{C}_0 the initial covariance matrix of the Gaussian proposal, n_0 the iteration at which the adaptive procedure starts and Δn_{adapt} the adaptation period at which the covariance matrix is updated. The first adaptation is based on the covariance estimation from the n_0 first iterations. Then, the covariance matrix is computed using a recursive formula at each iteration in order to avoid saving all the sample values. Considering that a new value \mathbf{m}^{n+1} has just been sampled from the posterior distribution from, using the Gaussian proposal distribution with current covariance matrix \hat{C}_n , the new covariance matrix estimate \hat{C}_{n+1} is computed using the new sample and the previous estimates as follows¹:

$$\hat{C}'_{n+1} = \frac{n+1}{n+2} \hat{C}'_n + \frac{1}{n+2} ((n+1)\bar{\mathbf{m}}_n \bar{\mathbf{m}}_n^T - (n+2)\bar{\mathbf{m}}_{n+1} \bar{\mathbf{m}}_{n+1}^T + \mathbf{m}_{n+1} \mathbf{m}_{n+1}^T), \quad (5.41)$$

$$\bar{\mathbf{m}}_{n+1} = \frac{1}{n+2} ((n+1)\bar{\mathbf{m}}_n + \mathbf{m}^{n+1}), \quad (5.42)$$

$$\hat{C}_{n+1} = s_d \hat{C}'_{n+1} + s_d \epsilon I_{n_u}, \quad (5.43)$$

where s_d is an additional scaling parameter that depends only on the dimension n_u of the parameter space, $\epsilon > 0$ is a correcting parameter and I_{n_u} is the n_u -dimensional identity matrix. The parameter s_d is used to improve the acceptance rate and [Gelman et al. \[1996\]](#) showed that the

¹[Haario et al. \[2001\]](#) use n to index the current sample value while we use the index $n+1$, resulting in slight differences in Eqs. 5.41–5.43 with no influence.

optimal scaling for a Gaussian distribution using a random-walk algorithm was $s_d = 2.38^2/n_u$. The correcting parameter $\epsilon > 0$ is used to ensure that the covariance estimate will not become singular and that the algorithm has the correct ergodic properties in theory, but can be set to zero in practice.

The AMH algorithm above has a proposal distribution that is Gaussian and the sampling of the target distribution can be slow when the latter has a complex dependence structure or is far from Gaussian. For instance, the covariance matrix can fail to adapt to the global scales of the distribution as the scales can be dependent of the current state. Therefore, several other improvements of the random-walk Metropolis–Hastings algorithms have been proposed in the literature, such as the delayed rejection Metropolis–Hastings, which scales the covariance matrix by a given factor in the case of rejections in order to increase the acceptance rate, its adaptive version the delayed-rejection adaptive Metropolis–Hastings, and many others that can be found, for instance, in [Gilks et al. \[1996\]](#); [Smith \[2014\]](#); [Robert and Casella \[2004\]](#). However, will see in the two next sections methods that seek to improve the mixing of the Markov chain by constructing proposal functions that take advantage of the local gradients of the target distribution and for which the proposal distribution is not Gaussian.

5.4.4 Hamiltonian Monte Carlo

The Hamiltonian Monte Carlo (HMC) method uses the solution of Hamilton’s equations as new states in the Markov chain Monte Carlo algorithm ([Duane et al. \[1987\]](#); [Neal \[2010\]](#); [Gelman et al. \[2014\]](#)). The classical Hamiltonian dynamics characterizes the evolution of a system described by its position \mathbf{q} and momentum \mathbf{p} whose total energy corresponds to the Hamiltonian $\mathcal{H}(\mathbf{q}, \mathbf{p})$. The time evolution of the system is uniquely defined by Hamilton’s equations

$$\frac{dq_i}{dt} = \frac{\partial \mathcal{H}(\mathbf{q}, \mathbf{p})}{\partial p_i}, \quad (5.44)$$

$$\frac{dp_i}{dt} = -\frac{\partial \mathcal{H}(\mathbf{q}, \mathbf{p})}{\partial q_i}. \quad (5.45)$$

The Hamiltonian is usually defined as the sum of a potential energy and the total kinetic energy of the system $\mathcal{H}(\mathbf{q}, \mathbf{p}) = \mathcal{U}(\mathbf{q}) + \mathcal{K}(\mathbf{p})$. The potential energy function $\mathcal{U}(\mathbf{q})$ in the Hamiltonian can be related to the probability distribution we wish to sample from. In the Bayesian framework, this will be minus the log of the posterior and the position variable \mathbf{q} becomes the variable of interest \mathbf{m} . The momentum variable \mathbf{p} is used for the Hamiltonian dynamics to work. The potential and kinetic energies of the Hamiltonian thus write

$$\mathcal{U}(\mathbf{m}) = -\log(\pi(\mathbf{m}|\mathbf{y}^{\text{obs}})), \quad (5.46)$$

$$\mathcal{K}(\mathbf{p}) = \frac{1}{2} \mathbf{p}^T M^{-1} \mathbf{p}. \quad (5.47)$$

where M is a symmetric, positive-definite matrix (or “mass matrix”). Using this, the derivatives in Eq. (5.44) are $\frac{\partial \mathcal{K}(\mathbf{p})}{\partial \mathbf{p}} = M^{-1} \mathbf{p}$, and the derivatives of the potential energy must be computed either numerically or analytically.

The Hamiltonian dynamics possesses several interesting properties that are necessary for constructing MCMC updates, namely reversibility and symplecticness¹. The numerical solution of Hamilton’s equations is therefore best carried out using a discretization scheme that respects those properties. The most common scheme is the leapfrog method which leads to the following

¹In particular, a consequence of the symplectic property is that the volume in the momentum-position phase space is preserved.

system of equations:

$$p_i \left(t + \frac{\tau}{2} \right) = p_i(t) - \frac{\tau}{2} \frac{\partial \mathcal{U}(\mathbf{q}(t))}{\partial q_i}, \quad (5.48)$$

$$q_i(t + \tau) = q_i(t) + \tau \frac{\partial \mathcal{K}(\mathbf{p}(t + \frac{\tau}{2}))}{\partial p_i}, \quad (5.49)$$

$$p_i(t + \tau) = p_i \left(t + \frac{\tau}{2} \right) - \frac{\tau}{2} \frac{\partial \mathcal{U}(\mathbf{q}(t + \tau))}{\partial q_i}, \quad (5.50)$$

where τ is the time step. The first equation is a half Euler step for the momentum, then one full step using the already computed momentum, then finally an other half step for the momentum with the already computed position. The integration is then performed for L leapfrog steps to reduce correlation between successive samples.

As the conservation of the Hamiltonian can only be approximated, the numerical scheme from Eqs. (5.48)–(5.50) introduces an error in the solution. Thanks to the reversibility of the Hamiltonian system, the resulting values for \mathbf{m} and \mathbf{p} can be inserted in a Metropolis–Hastings accept-reject step to correct the error in the solution introduced by the numerical integration. The acceptance ratio at iteration n writes:

$$r = \exp(-\mathcal{U}(\mathbf{m}^*) + \mathcal{U}(\mathbf{m}^n) - \mathcal{K}(\mathbf{p}^*) + \mathcal{K}(\mathbf{p}^n)). \quad (5.51)$$

By following trajectories of constant Hamiltonian values, the Hamiltonian Monte Carlo algorithm allows the new state to be quite distant from the current state, while nevertheless keeping a high acceptance ratio. This property is particularly advantageous to address the issue of the slow exploration of the posterior distribution that can be encountered when using a Gaussian proposal distribution from the random-walk dynamics. The different Hamiltonian values are finally explored by randomly drawing \mathbf{p} at the beginning of each new MCMC iteration from a centered multivariate Gaussian distribution with covariance M . The HMC algorithm is summarized in Algorithm 2.

Algorithm 2: Hamiltonian Monte Carlo algorithm (adapted from Neal [2010]).

```

Set  $\tau, L$ ;
Initialize  $\mathbf{m}^0$ ;
while  $n < n_{\text{iterations}}$  do
    Sample  $\mathbf{p}^* \sim \mathcal{N}(\mathbf{0}, M)$ ;
     $\mathbf{p}^n = \mathbf{p}^*$ ;
     $\mathbf{p}^* \leftarrow \mathbf{p}^* - \frac{\tau}{2} \frac{\partial \mathcal{U}}{\partial \mathbf{m}}(\mathbf{m}^n)$ ;
    while  $j \leq L$  do
         $\mathbf{m}^* \leftarrow \mathbf{m}^* + \tau M^{-1} \mathbf{p}^*$ ;
        if  $j \neq L$  then
             $\mathbf{p}^* \leftarrow \mathbf{p}^* - \tau \frac{\partial \mathcal{U}}{\partial \mathbf{m}}(\mathbf{m}^*)$ ;
        end
    end
     $\mathbf{p}^* \leftarrow \mathbf{p}^* - \frac{\tau}{2} \frac{\partial \mathcal{U}}{\partial \mathbf{m}}(\mathbf{m}^*)$ ;
    Compute acceptance probability  $\alpha = \min(1, r)$  with  $r$  from Eq. (5.51) and sample
     $u \sim U(0, 1)$ ;
    if  $u \leq \alpha$  then
         $\mathbf{m}^{n+1} = \mathbf{m}^*$ ;
    else
         $\mathbf{m}^{n+1} = \mathbf{m}^n$ ;
    end
end

```




Figure 5.1: Sampling from a two-dimensional Gaussian distribution with standard deviation of 0.98 using (a) the random-walk Metropolis–Hastings algorithm and (b) the Hamiltonian Monte Carlo algorithm. The starting iteration at (0,0) is marked with an asterisk and the circles represent lines of equal probability density.

To avoid the parameter values to go out of the support of the posterior, we exploit the rejection mechanism and we automatically reject such parameter values.

Illustration

The following example is reproduced from Neal [2010]. We sample from a two-dimensional Gaussian distribution with unit standard deviation and correlation equal to 0.98. For the kinetic energy, we take $M = I$. The Hamiltonian is therefore defined as

$$\mathcal{H}(\mathbf{m}, \mathbf{p}) = \frac{\mathbf{m}^T \Sigma^{-1} \mathbf{m}}{2} + \frac{\mathbf{p}^T \mathbf{p}}{2} \quad \text{with} \quad \Sigma = \begin{bmatrix} 1 & 0.98 \\ 0.98 & 1 \end{bmatrix}. \quad (5.52)$$

Thanks to the accept-reject step in the Metropolis–Hastings algorithm, the constant in the Gaussian potential function can be dropped.

The comparison of the sampling of the two-dimensional Gaussian distribution using the RWMH and HMC (implemented in `pybitup`) is shown in Fig. 5.1. In the HMC algorithm, the time step used is $\tau = 0.18$ with $L = 20$ leapfrog steps resulting in a total of 20 effective iterations. More than one leapfrog step is generally used to reduce the correlation between the samples. For the comparison with the RWMH to be fair, we performed 400 iterations and show every 20 iterations. We can observe from Fig. 5.1 that the HMC algorithm explores the support of the distribution more efficiently than the RWMH, which remains during a few iterations in the same region of low probability density.

5.4.5 Algorithm based on an Itô stochastic differential equation

The second gradient-informed MCMC algorithm described is based on the use of an ISDE (Soize [2008, 2015, 2017]) associated with a dissipative Hamiltonian and a stochastic forcing term. Soize [2015]; Arnst et al. [2017] and Arnst and Soize [2019] applied a generalized version of the ISDE that takes into account the scales and correlations between the different parameters. We denote by $\{\mathbf{Q}(t), t \geq 0\}$ and $\{\mathbf{P}(t), t \geq 0\}$ two \mathbb{R}^{n_u} -valued stochastic processes indexed by \mathbb{R}^+ and $\{\mathbf{W}(t), t \geq 0\}$ the \mathbb{R}^{n_u} -valued normalized Wiener process whose increment $\Delta \mathbf{W}_{st} = \mathbf{W}(t) - \mathbf{W}(s)$ is an \mathbb{R}^{n_u} -valued centered Gaussian random vector with covariance matrix $(t - s)I_{n_u}$.

The following stochastic differential equation describes a second-order system damped with a stochastic excitation:

$$d\mathbf{Q} = M^{-1}\mathbf{P}dt, \quad (5.53)$$

$$d\mathbf{P} = -\nabla_{\mathbf{q}}\mathcal{U}(\mathbf{Q})dt - \frac{1}{2}f_0\mathbf{P}dt + \sqrt{f_0}L_Md\mathbf{W}, \quad (5.54)$$

where $M = L_M L_M^T$ is a symmetric, positive-definite matrix with L_M the lower triangular matrix resulting from the Cholesky decomposition, f_0 is a free parameter (Arnst and Soize [2019]) that controls the amplitude of the deterministic damping and the stochastic forcing. In the limit of zero damping, the system is similar to the (deterministic) conservative Hamiltonian system describes in Section 5.4.4.

Soize [2008] showed that under some regularity conditions regarding the potential function \mathcal{U} , the stationary density of (\mathbf{Q}, \mathbf{P}) is independent of the damping parameter and is given by

$$\pi(\mathbf{q}, \mathbf{p}) \propto \exp\left(-\mathcal{U}(\mathbf{q}) - \frac{\mathbf{p}^T M^{-1} \mathbf{p}}{2}\right). \quad (5.55)$$

Once again, by letting the position variable \mathbf{q} be the variable of interest \mathbf{m} and choosing a potential function that corresponds to $\mathcal{U} = -\log(\pi(\mathbf{m}))$, the marginal stationary density for \mathbf{Q} will be exactly the distribution we wish to sample from (e.g., the Bayesian posterior $\pi(\mathbf{m}|\mathbf{y}^{\text{obs}})$).

To solve the stochastic irreversible system, symplectic integrators also show a favorable long-time behavior (Hairer et al. [2006]). They will have the advantage of being able to solve the system in the long-term behavior even for low damping values and the same leapfrog discretization as for the Hamiltonian system is applicable. In Soize and Ghanem [2016]; Guilleminot and Soize [2014] and Arnst and Soize [2019], the leapfrog discretization is used with the slight modification of Mannella [2006], which allows to have an error in the variance of the momentum variable that is independent of the damping parameter (Burrage et al. [2007]), that is referred subsequently to as the Störmer-Verlet method. The numerical solution of the stochastic irreversible (non-Hamiltonian) system is performed by applying the Störmer-Verlet scheme (Soize [2017]):

$$\mathbf{Q}^{(n+\frac{1}{2})} = \mathbf{Q}^{(n)} + \frac{\tau}{2}M^{-1}\mathbf{P}^{(n)}, \quad (5.56)$$

$$\mathbf{P}^{(n+1)} = \frac{1-b}{1+b}\mathbf{P}^{(n)} - \frac{\tau}{1+b}\nabla_{\mathbf{Q}}\mathcal{U}\left(\mathbf{Q}^{(n+\frac{1}{2})}\right) + \frac{\sqrt{f_0}}{1+b}L_M\Delta\mathbf{W}^{(n+1)}, \quad (5.57)$$

$$\mathbf{Q}^{(n+1)} = \mathbf{Q}^{(n+\frac{1}{2})} + \frac{\tau}{2}M^{-1}\mathbf{P}^{(n+1)}, \quad (5.58)$$

where $b = f_0\tau/4$. The dynamics of the ISDE is not reversible due to the dissipative behavior of the system, as opposed to the Hamiltonian dynamics, and the accept-reject step from the Metropolis–Hastings cannot be applied. In order to limit the drift induced by the numerical discretization, the time step τ will have to be chosen carefully.

The implementation of the ISDE algorithm assumes that the support of the target PDF is the entire space \mathbb{R}^{n_u} . It is not directly applicable in the case of a bounded or semi-bounded parameter set as the algorithm may reach parameter values where the potential is not defined. To avoid this, Soize [2017] has proposed to use regularization or rejection techniques. An other possibility is to use a change of variables, as in Arnst et al. [2017], that will map the original variables into a transformed parameter space where the ISDE algorithm is applicable, which is discussed in the next section.

Illustration

In Fig. 5.2, we compare the sampling strategy of the ISDE algorithm with the HMC algorithm for a one-dimensional Gaussian distribution, thus allowing us to represent the phase space for the

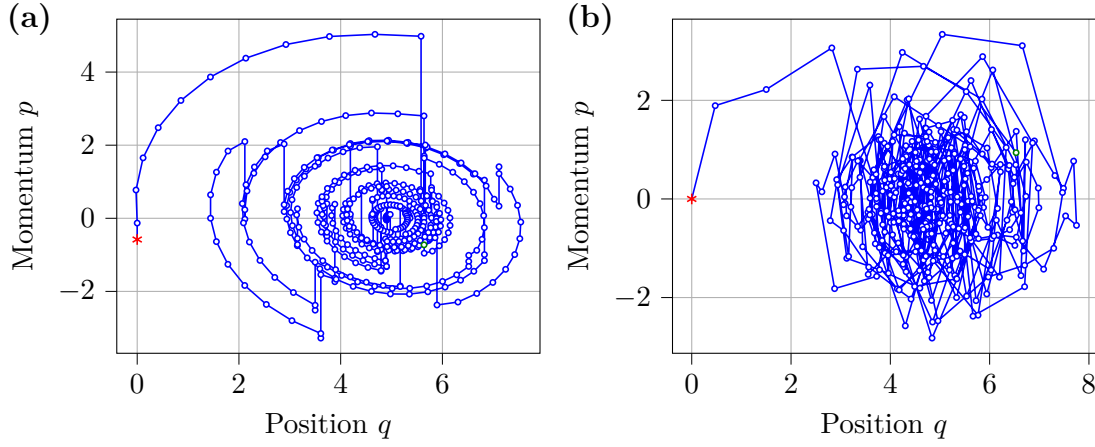


Figure 5.2: Sampling from a one-dimensional standard Gaussian distribution using (a) the Hamiltonian Monte Carlo algorithm and (b) the algorithm based on an Itô stochastic differential equation. The starting iteration at $(0,0)$ is marked with an asterisk.

position and momentum. For the HMC algorithm (Fig. 5.2(a)), the time step used is $\tau = 0.18$ with $L = 10$ leapfrog steps. For the ISDE (Fig. 5.2(b)), the time step is set to $\tau = 0.5$ and $f_0 = 4.0$. In Fig. 5.2(a), we can clearly observe the trajectory of the samples that follows the line of equal Hamiltonian values. At each new iteration, the value for the momentum p is re-sampled, resulting in the straight lines on the graph. The samples selected are the values at the end of each trajectory along a constant Hamiltonian line.

Compared to the HMC algorithm, the ISDE has the advantage of controlling the stochastic perturbation through the damping coefficient f_0 , while the Hamiltonian is based on a conservative system to which the stochasticity is added through the resampling of the momentum at each new iteration. Both methods have the advantage of avoiding the need to introduce a proposal distribution and thus avoiding the random-walk.

5.4.6 Change of variables based on the support of the posterior

A change of variables can accelerate convergence because it can simplify the structure of the posterior or it can help to decorrelate random variables. When applying a change of variables, the Jacobian rule for the transformation of the PDF under a differentiable bijection requires the multiplication of the image of the posterior PDF for the original variables by the inverse of the determinant of the Jacobian. See, for instance, in Appendix B for the computation of the acceptance ratio in the case of a change of variables for the Metropolis–Hastings algorithms.

The use of gradient-informed sampling algorithms requires the computation of the gradient of the logarithm of the posterior distribution and in turn of the gradient of the logarithm of the prior and the likelihood distributions. The use of distributions that are not C^1 (continuously differentiable) such as uniform priors might therefore be troublesome. When the support of the variable is a subset $\mathcal{S} \subset \mathbb{R}$, we therefore need to apply a change of variables from \mathcal{S} to the real line \mathbb{R} . Exception can be made for the class of Metropolis–Hastings algorithms, as the Hamiltonian Monte Carlo, which are rejection methods. The gradient of the uniform prior is set to a constant value everywhere, and the accept-reject step will automatically reject the new value if it is outside of the bounds of the prior distribution. For the ISDE algorithm, a change of variable is applied as there is no rejection step, and we make the ISDE feasible by mapping the initial parameter set into a transformed unbounded parameter set.

Thus, we introduce a general change of variables that is based on the structure of the poste-

rior. We distinguish two cases. First, the case where \mathbb{S} is bounded from above and below with $\mathcal{S} = [L_b, U_b]$, we choose a change of variables of the following form:

$$\tilde{m}_i = \tan \left\{ \left(\frac{m_i - L_b}{U_b - L_b} - \frac{1}{2} \right) \pi \right\}, \quad (5.59)$$

$$\frac{d\tilde{m}_i}{dm_i} = \left[1 + \tan^2 \left\{ \left(\frac{m_i - L_b}{U_b - L_b} - \frac{1}{2} \right) \pi \right\} \right] \frac{\pi}{U_b - L_b}. \quad (5.60)$$

Second, in the case where the support \mathcal{S} is semi-infinite, i.e. m_i is in \mathbb{R}_0^+ (or \mathbb{R}_0^-), we apply a logarithmic change of measure $m_i \mapsto \log(m_i)$ (or $\log(-m_i)$) from \mathbb{R}_0^+ (or \mathbb{R}_0^-) into \mathbb{R} .

5.4.7 Checking and monitoring convergence

We address here the two questions of how long the MCMC chain should be run and how we can diagnose a slow or a lack of convergence. Mixing and stationarity of a Markov chain are two important properties that need to be verified when deciding whether the Markov chain has converged or not. The first concept refers to the samples of the chain that must have covered all the support of the distribution. The second one refers to the case where two chains have mixed and cover the same distribution, but that neither sequence appears stable in which case the chains are not stationary (see, for instance, in [Raftery and Lewis \[1996\]](#) for an example). A first indicator of a good convergence is therefore to compare two chains that were simulated independently with possibly different starting values. After discarding an initial burn-in phase, the juxtaposition of the two chains should be indistinguishable at convergence.

Mixing can be checked by looking at moments of the distribution and comparing their values. For instance, in order to diagnose a slow mixing of the chains, between-sequence and within-sequence variances can be computed ([Gelman et al. \[2014\]](#)). First, after discarding the initial burn-in samples, each chain is split in half and the mixing is checked for the resulting half-sequences. Let denote by n_c the number of chains ran in parallel, hence the total number of sequences is $2n_c$, and by $n_h = n_{it}/2$ the length of the half-sequence chains. Let denote by $m_k^{(ij)}$ the i -th iteration from chain j of the k -th component of the vector of parameters \mathbf{m} . To assess the mixing, the between-sequence B and within-sequence W variances are computed as ([Gelman et al. \[2014\]](#))

$$B_k = \frac{n_h}{2n_c - 1} \sum_{j=1}^{2n_c} \left(\bar{m}_k^{(\cdot j)} - \bar{m}_k^{(\cdot \cdot)} \right)^2, \quad (5.61)$$

$$\text{where } \bar{m}_k^{(\cdot j)} = \frac{1}{n_h} \sum_{i=1}^{n_h} m_k^{(ij)}, \quad \bar{m}_k^{(\cdot \cdot)} = \frac{1}{2n_c} \sum_{j=1}^{2n_c} \bar{m}_k^{(\cdot j)}, \quad (5.62)$$

$$W_k = \frac{1}{2n_c} \sum_{j=1}^{2n_c} \left(s_k^{(j)} \right)^2, \quad \text{where } \left(s_k^{(j)} \right)^2 = \frac{1}{n_h - 1} \sum_{i=1}^{n_h} \left(m_k^{(ij)} - \bar{m}_k^{(\cdot j)} \right)^2 \quad (5.63)$$

Marginal posterior variances can be estimated as

$$\hat{\text{var}}^{(k),+}(m^{(k)}|\mathbf{y}^{\text{obs}}) = \frac{n_h - 1}{n_h} W^{(k)} + \frac{1}{n_h} B^{(k)}, \quad (5.64)$$

and convergence of the chain is monitored through the potential scale reduction factor

$$\hat{R} = \sqrt{\frac{\hat{\text{var}}^{(k),+}(m^{(k)}|\mathbf{y}^{\text{obs}})}{W^{(k)}}} \quad (5.65)$$

which declines to 1 as $n_h \rightarrow \infty$.

5.5 Propagation of uncertainties

Having defined the computational model and the characterization of the input uncertainty, we now address the problem of how this uncertainty propagates through the computational model to the quantity of interest. We formulate the problem as follows: we have an \mathbb{R}^n -valued random vector \mathbf{X} for which we know the probability distribution and we are interested in the \mathbb{R} -valued random variable Y which is some transformation of \mathbf{X} through the computational model as follows:

$$Y = \eta(\mathbf{X}), \quad (5.66)$$

and we want to know the probability distribution of Y . The probability distribution \mathbb{P}_Y of Y is such that for any meaningful event $D_Y \in \mathfrak{D}_Y$, $\mathbb{P}_Y(D_Y) = P_{\mathbf{X}}(\{\mathbf{x} \in \mathbb{R}^n : \eta(\mathbf{x}) \in D_Y\})$. Thus, to obtain $\mathbb{P}_Y(D_Y)$ one needs to obtain the probability $\mathbb{P}_{\mathbf{X}}$ of corresponding values of $\mathbf{x} \in \mathbb{R}^n$ that the computational model maps into values in D_Y , which subset is written as $\{\mathbf{x} \in \mathbb{R}^n : \eta(\mathbf{x}) \in D_Y\}$. In the cases where the transformation $\eta(\cdot)$ is a one-to-one mapping, the characterization of \mathbb{P}_Y can be performed analytically (see for instance in [Hogg et al. \[2005\]](#)). However, this case is seldom in real-world applications and the propagation of uncertainties is most often solved numerically.

Numerical methods for uncertainty propagation are often divided into two main categories, namely intrusive and non-intrusive methods. On the one hand, intrusive methods require the modification of the computer code. On the other hand, non-intrusive methods require little or no modification of the code and the model can be seen as a black-box. There are polynomial chaos methods that consist in replacing the existing computational model by using a simpler representation of the mapping between the uncertain input parameters and the quantity of interest (called the surrogate) when the mapping can be truncated at a sufficiently low order. This solution to the computer code is approximated by means of an orthogonal polynomials basis with unknown coefficients that must be estimated ([Xiu and Karniadakis \[2002\]](#)). See in [Le Maître and Knio \[2010\]](#) or [Xiu \[2010\]](#) for an overview of intrusive and non-intrusive methods for uncertainty propagation.

In this section, we focus on two non-intrusive methods for uncertainty propagation. First, the most intuitive and straightforward approach to uncertainty propagation using non-intrusive methods is the direct Monte Carlo sampling method. It has the advantage of being applicable to problems of any stochastic dimension and is suitable when we have a set of samples drawn from an unknown PDF, which can have correlation between its variables. Second, we focus on the non-intrusive polynomial chaos method which is widely used in uncertainty quantification studies.

5.5.1 Monte Carlo sampling method

The Monte Carlo approach involves repeated simulations with random sampling from the space of the random vector $\mathbf{X} = (X_1, \dots, X_n)$ with values in \mathbb{R}^n . An ensemble of independent and identically distributed (i.i.d.) samples $\{\mathbf{x}^{(k)}, 1 \leq k \leq \nu\}$ are generated from the probability distribution $\mathbb{P}_{\mathbf{X}}$. Then, the solution of the computational model is generated for each of the samples to obtain the corresponding ensemble of i.i.d. samples from \mathbb{P}_Y written as $\{y^{(k)} = \eta(\mathbf{x}^{(k)}), 1 \leq k \leq \nu\}$.

The statistical descriptors for a random variable introduced in Eq. (5.6) and Eq. (5.9) can then be estimated from the i.i.d. samples $y^{(k)}$ as

$$\mu_Y \approx \mu_Y^\nu = \frac{1}{\nu} \sum_{k=1}^{\nu} y^{(k)}, \quad (5.67)$$

$$\sigma_Y^2 \approx (\sigma_Y^\nu)^2 = \frac{1}{\nu} \sum_{k=1}^{\nu} \left(y^{(k)} - \mu_Y^\nu \right)^2. \quad (5.68)$$

where the superscript ν is used to indicate that the quantity of interest y is estimated from the ν i.i.d. samples. The law of large number ensures that the approximation μ_Y^ν converges toward μ_Y when the sample size ν increases and the accuracy is proportional to $1/\sqrt{\nu}$ (Robert and Casella [2004]). Hence, the convergence of the Monte Carlo method is independent of the dimensionality of the random vector, which is advantageous for problems involving a large number of random input variables.

5.5.2 Polynomial chaos method

Polynomial chaos was originally described in the work of Wiener [1938], and was later used by Ghanem and Spanos [1991] to model uncertainty in mechanical engineering problems. Since then, it has been the subject of many investigations. Xiu and Karniadakis [2002] extended the classical polynomial chaos for Gaussian processes using Hermite polynomial (Wiener-Hermite chaos) to a general set of polynomials provided by the so-called Askey scheme of polynomials (Askey and Wilson [1985]). Generalized polynomial chaos (gPC) are summarized in this chapter.

Let \mathbf{X} be a \mathbb{R}^d -valued random vector that admits the probability density $\pi_{\mathbf{X}}$. Let us assume a random scalar output quantity Y is the transformation η of the random vector \mathbf{X} , that is the mapping $Y : \mathbf{X} \in \mathbb{R}^d \rightarrow \eta(\mathbf{X}) \in \mathbb{R}$. Let $\mathbf{i} = (i_1, \dots, i_n) \in \mathbb{N}_0^n$ be a multi-index of modulus $|\mathbf{i}| = i_1 + \dots + i_n$. Then the polynomial chaos method consists in representing the model η by a spectral expansion using a predefined basis of orthonormal polynomials $\{\psi_{\mathbf{i}}(\mathbf{X}), \mathbf{i} \in \mathbb{N}_0^n\}$ from \mathbb{R}^d into \mathbb{R} (the *polynomial chaos*) such that

$$Y = \eta(\mathbf{X}) = \sum_{|\mathbf{i}|=0}^{\infty} y_{\mathbf{i}} \psi_{\mathbf{i}}(\mathbf{X}). \quad (5.69)$$

Equation (5.69) is the polynomial chaos expansion (PCE) of the output quantity and the coefficients $y_{\mathbf{i}}$ are unknown. In practice, the series from Eq. (5.69) is truncated and the model $\eta(\mathbf{X})$ is approximated by the truncated sum

$$\eta(\mathbf{X}) \approx s^p(\mathbf{X}) = \sum_{|\mathbf{i}|=0}^p y_{\mathbf{i}} \psi_{\mathbf{i}}(\mathbf{X}). \quad (5.70)$$

There are two common approaches to compute the expansion coefficients: the spectral projection and regression. In regression methods, the problem of finding the coefficients of the PCE is cast under the form of a weighted least-squares problem. The model function is replaced by its truncated PCE approximation and the expansion coefficients are obtained as the solution of the minimization problem of the residual

$$\arg \min_{\{y_{\mathbf{i}}, 0 \leq |\mathbf{i}| \leq p\}} \sum_{j=1}^m \left(\eta(\mathbf{x}_j) - \sum_{|\mathbf{i}|=0}^p y_{\mathbf{i}} \psi_{\mathbf{i}}(\mathbf{x}_j) \right)^2, \quad (5.71)$$

which requires the computation of the forward model at m minimization points which need to be selected carefully. See Le Maître and Knio [2010] for a more comprehensive discussion on the least-squares fit approach for the computation of the expansion coefficients.

Spectral projection methods take into account the orthonormality condition of the basis with respect to the probability distribution $\mathbb{P}_{\mathbf{X}}$, which implies that

$$\mathbb{E}[\psi_{\mathbf{i}} \psi_{\mathbf{j}}] = \int_{\Xi} \psi_{\mathbf{i}}(\mathbf{x}) \psi_{\mathbf{j}}(\mathbf{x}) \pi_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} = \delta_{\mathbf{ij}}, \quad \forall \mathbf{i}, \mathbf{j} \in \mathbb{N}_0^n, \quad (5.72)$$

where $\delta_{\mathbf{ij}} = 1$ for $\mathbf{i} = \mathbf{j}$ and 0 otherwise, such that the polynomial chaos coordinates write as

$$y_{\mathbf{i}} = \mathbb{E}[\psi_{\mathbf{i}} Y] = \int_{\Xi} \psi_{\mathbf{i}}(\mathbf{x}) \eta(\mathbf{x}) \pi_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}, \quad \mathbf{i} \in \mathbb{N}_0^n. \quad (5.73)$$

The evaluation of the coefficients amounts to the computation of a multidimensional integral. This integral is usually obtained numerically and several approximation methods have been proposed to estimate its value. These techniques are classified into simulations approaches, that use Monte Carlo sampling or improved sampling strategies in a similar manner as discussed in Section 5.5.1, and cubature approaches, that consist in approximating the integral by a discrete sum. Only the latter is discussed below. The PCE expansion that uses numerical integration to compute the value of its coefficient is sometimes referred to as “pseudo-spectral” (Xiu [2007]).

Cubature rules for multidimensional integration

Multidimensional cubature rules are generally constructed on the basis of 1D integration rules, or *quadratures*. The univariate integral I of a function f with respect to the measure \mathbb{P}_X is approximated by a weighted sum of function evaluations at specified points within the domain of integration

$$I = \int_{\mathbb{R}} f(x) \pi_X(x) dx \approx \sum_{i=1}^{N_l} f(x_i) w_i. \quad (5.74)$$

where the x_i are the quadrature points or nodes and the w_i the quadrature weights.

Gauss integration formulas consist in finding the nodes x_i of the quadrature rules from the zeros of the N_l -th order orthogonal polynomials with regard to the measure \mathbb{P}_X . Gauss quadratures yield an optimal degree of exactness for polynomial functions f of degree $\leq 2N_l - 1$. Hence, given that the model solution η is polynomial and for a given degree of the polynomial ψ_i , one can deduce the number of nodes that are needed to yield the exact integration of the polynomial function $\eta(x)\psi_i(x)$. For labeled distributions, the orthogonal polynomials are known and efficient formulas for computing the values of the nodes and the weights are available (Abramowitz and Stegun [1972]). For instance, in the case of a uniform measure on the interval $[-1, 1]$, the associated orthogonal polynomials are from the Gauss–Legendre family, or for the standard Gaussian measure the associated orthogonal polynomials are the Gauss–Hermite family. Xiu and Karniadakis [2002] established the numerical rate of convergence for a set of labeled distribution with their corresponding polynomials from the Askey scheme family of orthogonal polynomials and showed that the error exhibits an exponential convergence rate.

In general, the model solution η is rarely polynomial and the exact integration can not be achieved. Moreover, Gauss integration formulas have the drawback of having a node distribution that depends on the degree of the polynomial assumed. If one wishes to improve the accuracy of the PCE approximation by increasing the number of nodes, one cannot re-use the nodes computed at a lower degree. Instead, efficient quadrature points seek to reduce the number of model evaluations. Nested quadratures are interesting in the sense that they can re-use the nodes computed at a level l for the next level $l + 1$, although the convergence rate might be slower than using the quadrature rule that corresponds to the probability measure. Examples of nested quadrature rule are the Clenshaw–Curtis and Fejèr rules, or Gauss–Patterson rule.

One-dimensional quadrature rules are easily extended to n -dimensional cubature rules via tensor product construction. Considering that the parameters are independent, with possibly different distributions, one can apply the individual quadrature rules along the different integration directions. However, tensor product construction suffers from the curse of dimensionality, as the total number of cubature points resulting from the tensor product formula increases exponentially as the dimension increases. Sparse tensorization quadrature formulas, or sparse-grid quadrature rules, have been proposed and offer an appreciable reduction of the number of cubature nodes for a given degree of exactness, and most of the applications rely on the Smolyak sparse grid algorithm (Smolyak [1963]). In practice, using sparse-grid quadrature rules can result in errors on the coefficients associated with high-order basis polynomials. In the application of Chapter 8, we will use a slightly different method to construct the sparse-grid quadrature rules, the so-called sparse pseudo-spectral approximation method (Constantine et al. [2012]),

implemented in the Dakota ([Adams et al. \[2018\]](#)) software, which attempts to reduce the error on the coefficients associated to the high-order polynomials.

5.6 Summary

In this chapter, we discussed the theory for the analysis of computational models, from the characterization of the uncertainty in the input parameters to the propagation of uncertainty through the computational model and how to quantify it.

For the characterization of input uncertainty, we focused on the Bayesian approach in which input prior knowledge is updated through the likelihood function to provide the posterior probability distribution. Then, we have reported several algorithms to sample from a probability distribution, including the random-walk Metropolis–Hastings methods and its adaptive version, as well as gradient-informed sampling algorithms, namely the Hamiltonian Monte Carlo and the method based on an Itô stochastic differential equation. For the propagation of uncertainty, we discussed the Monte Carlo method, which can be applied to problems with an arbitrary number of input parameters but which features a low convergence rate, and the polynomial chaos methods.

These methods will be of interest for the quantification of uncertainty in the characterization of thermal protection material properties and the propagation of uncertainty through numerical codes for computing the material response during an atmospheric reentry.

Development of a numerical toolbox for Bayesian inference and uncertainty propagation

*This chapter provides an introduction to the software **Pybitup** developed during this thesis. The **Python** Bayesian inference toolbox for uncertainty propagation (**Pybitup**) proposes an integrated tool for performing uncertainty quantification, from Bayesian calibration to uncertainty propagation. The code is written in **Python** such that it allows to easily implement new methods and to test them directly on external or built-in applications, and intends to cover a wide range of engineering problems. After a quick review of existing software and discussing the objectives of the present software, we describe the main features of the code. It was initially tailored to the calibration of time-dependent processes defined by a mathematical model that represents the physics with unknown parameters using MCMC methods. Several sampling algorithms are implemented, including Metropolis–Hastings algorithms and gradient-informed (HMC and ISDE) algorithms. For the propagation of uncertainty, the sample values characterizing the probability density function resulting from the parameter calibration can be used and propagated through built-in or third party numerical models. Polynomial chaos methods were further implemented for uncertainty propagation. We finally provide a short example that illustrates the workflow for model parameter calibration. The new software does not intend to compete with current “industrial level” toolboxes for uncertainty quantification but is intended to be a research code in which state-of-the-art algorithms can be implemented in **Python**.*

6.1 Brief review of existing software

There exists many computational libraries for Bayesian inference and uncertainty propagation. For Bayesian inference, the BUGS (Bayesian inference Using Gibbs Sampling) program was first developed in the early 1980s (Lunn et al. [2009]) with the aim of providing a statistical toolbox for sampling from the posterior distribution using a convenient programming language. A related package of BUGS is JAGS (Just Another Gibbs Sampling) which is more UNIX-oriented. More recently, the software Stan (Carpenter et al. [2017]) written in C++, implements efficient sampling algorithms for high dimensions and has a good computational efficiency. The main simulation algorithm is the Hamiltonian Monte Carlo that features several advanced improvements. OpenTURNS is a general C++ library designed to perform uncertainty quantification and reliability analysis (Baudin et al. [2015]), and implements the Metropolis–Hastings algorithm for

Bayesian calibration. Dakota (Adams et al. [2018]) is an open-source C++ software developed by Sandia National Laboratories. It was initially designed for optimization but was then extended to uncertainty propagation, and more recently to Bayesian calibration with the use of the Queso library. The latter library is a collection of algorithms aimed for the solution of statistical inverse problems (among others), and other functionalities for uncertainty propagation (Prudencio and Schulz [2012]). Similarly, the UQToolkit (Debusschere et al. [2017]), written in C++, proposes a collection of tools for uncertainty quantification, ranging from forward propagation of uncertainty to inverse problems and sensitivity analysis. In Matlab, UQLab (Marelli and Sudret [2014]) has become a reference software for uncertainty quantification. It proposes a set of tools for uncertainty propagation, sensitivity analysis, or Bayesian inference for inversion and model calibration. For Bayesian calibration, MCMC methods are implemented among which are the RWMH, AMH and HMC. The MIT Uncertainty Quantification Library (MUQ) (Parno et al. [2014]) is a C++/Python library that allows sampling by means of MCMC techniques, among which are the RWMH, AMH, and Metropolis-adjusted Langevin (MALA), and forward uncertainty propagation using PCE methods. The description of other software for uncertainty quantification and more details about some of the software mentioned above can be found in Ghanem et al. [2017].

In the recent years, the use of Python as programming language has seen a growing interest. The development by researchers and programmers of pieces of codes tailored to their problem has become popular, and to which the development of new libraries for UQ has followed. For instance, the toolbox PyMC (Patil et al. [2010]) was developed for Bayesian statistical modeling. It implements Metropolis–Hastings algorithms and the Hamiltonian Monte Carlo and several versions of slice samplers. More advanced methods are made available in PyMC3 (Salvatier et al. [2016]), which has also several interesting features for probabilistic machine learning. The Python package HIPPylib (Villa et al. [2021]) features, for instance, the MALA algorithm for sampling from the posterior and uses Hessian information.

6.2 Purpose of the present code

The result of a Bayesian parameter calibration is not always reused for uncertainty propagation. Because of the complexity of nowadays computational models, uncertainty propagation is usually performed assuming standard distributions (e.g., uniform, Gaussian, gamma) for the input parameters and for which it is convenient to build a surrogate model. However, this input uncertainty is not necessarily representative of the parameter distribution obtained, for instance, from the calibration with ground-test experiments. Building surrogate models for computational models with non-standard or unknown input distribution (e.g., a posterior distribution from which we have MCMC samples) is still an active research domain (Jakeman et al. [2019]).

The main purpose of Pybitup is to implement a toolbox for uncertainty quantification from the calibration of model parameters to uncertainty propagation, whereby the use of the samples resulting from the Bayesian calibration can directly be used for uncertainty propagation through a computational model (not necessarily the model used for the calibration). The code is implemented in a versatile way such that the software is readily usable and is suitable to general engineering problems, and allow external users to apply the toolbox without the need to care about the algorithmic details. With the growing interest of using Python as a coding language, the development of the toolbox in Python was found to be convenient, and further allows more advanced users to have the flexibility to implement easily state-of-the-art algorithms for research purposes. The problems that are sought are the calibration of time-dependent observations to a computational model representing the complex physics of the process and the propagation of the calibrated parameters. Hence besides the application to the pyrolysis problems from Chapter 7, several applications were investigated, that are the inference of parameters for a heat capacity model and a thermal conductivity model, or the inference of parameters for models relevant to

sloshing problems and contact angle determination (Dania [2021]).

6.3 Implementation

6.3.1 Dependencies

As `Python` loops are inherently slower than their low-level counterpart, the library mainly relies on the `Numpy` and `Scipy` packages integrating `C` and `Fortran` parallel codes for efficient matrix operations. The `pandas` library provides the data manipulation tools, such as reading and writing “.csv” files for input and output fluxes and the `pickle` library provides packages for manipulating large data sets efficiently. The `json` and `jsmin` libraries are used for the input file management and to have a structured input file. Finally, `matplotlib` and `seaborn` libraries are used for post-processing the results and producing graphs.

6.3.2 Input/output

Every case must define a run file, which contains the different tasks that the code will perform, and an input file, which contains the corresponding inputs of the tasks. For instance, for Bayesian inference, we will specify that we want to sample from the posterior, and one needs to specify what are the experimental data, the model used with its parameters and the method for the inference.

The run file is a `Python` script (identified by the `.py` extension) that specifies the path to the input file, defines the physical model (except in the case of sampling from a known distribution) and lists the different tasks to perform. To define the tasks, we need to access the `Python` object in `Pybitup` that contains the method for reading the input file and depends on the problem. It is kept as simple as possible, and usually the method parameters are defined in the input file described below. Most of the tasks are defined in the `solve_problem.py` file. For instance, to perform a sampling task, a generic run file is as follows: `post_dist = pybitup.solve_problem.Sampling(input_file_name)`.

The input file is based on the `Python` JSON (JavaScript Object Notation) syntax that makes the file easy to read and write for the user. The `Python` JSON syntax is similar to a dictionary in `Python` with the keys that are quoted-strings. The value of these keys can be any primitive type in `Python`, for instance a string, a number or an array. Because `pybitup` can perform sampling, propagation, sensitivity analysis and data post-processing, the most basic JSON file looks as follows: `{ "Sampling": {}, "Propagation": {}, "SensitivityAnalysis": {}, "PostProcess": {} }`. An minimal syntax of the structure of the input for the sampling from a Bayesian posterior is sketched below:

```
"Sampling": {
    "BayesianPosterior" : {
        "Data": [],
        "Model": [],
        "Prior": {},
        "Likelihood": {}
    },
    "Algorithm": {}
},
```

The `BayesianPosterior` keywords means that we will sample from a posterior distribution computed from Bayes’ formula. The `Algorithm` section will define the algorithm and the parameters used to sample from it.

6.3.3 The model class

Except in the case of sampling from a known PDF, a numerical model will have to be evaluated for the sampling from the Bayesian posterior or for propagating uncertainty. There are no pre-implemented models within **Pybitup** and the model must be specified in a separated **Python** file following the **Model** class structure. The **Model** class contains the information about the physical model (and not the statistical model) and the class defines the structure that a model defined in **Pybitup** needs to provide. The simplest structure of the **Model** class is as follows:

```
import numpy as np
from pybitup import bayesian_inference as bi

class Model:
    """ Class defining the model function. """
    def __init__(self , x=[], param=[], name=""):

        def fun_x(self):
            # Return the function evaluation
```

The **fun_x** method is mandatory and specifies the function that returns the solution of the model given the parameter values set in **param** and the abscissa **x**. This function can be defined directly within the **Python** file or can be provided from an other program (either a **Python** module or an external solver written in another language) in which case the **fun_x** method needs to point towards the execution of that program. A **name** can be used to identify the model, that is used for instance when several models are provided.

Several algorithms use a change of variables (reparameterization) (e.g., the ISDE algorithm), or one can simply wish to reparameterize the model to simplify it, in which cases the Jacobian of the transformation needs to be defined. Several algorithms also use the gradients of the model. The **Model** class can incorporate such changes of variables or the gradients of the model, and are optional (default is **p** = **p̄**). The additional methods that need to be defined when a change of variables or the computation of gradients are required are the following, with the default values provided below:

```
def d_fx_dparam(self , *ext_parameters):
    return 1

def parametrization_forward(self , X=1, P=1):
    Y = X
    return Y

def parametrization_backward(self , Y=1, P=1):
    X = Y
    return X

def parametrization_det_jac(self , X=1):
    det_jac = 1
    return det_jac

def parametrization_inv_jac(self , X=1):
    inv_jac = X
    return inv_jac
```

6.4 Sampling

One of the main features of `Pybitup` is the sampling from a probability distribution function. The code is based on the definition of a `ProbabilityDistribution` class. For known distributions, it is defined by a PDF and its value at a given sample m can be accessed by calling the `compute_value` method. For Bayesian posteriors, the definition of the likelihood and the prior distributions need to be first specified. In this case the `compute_value` method returns directly the product of the likelihood times the prior distribution.

6.4.1 Sampling from a Bayesian posterior

When sampling from a Bayesian posterior, a dataset and a physical model need to be provided, for which two `Python` classes are dedicated, the `Data` class and the `Model` class (described previously), respectively. The `Data` class contains the abscissa and ordinates (the observations) of the problem. The data are specified in a list and several data set can be specified, e.g, `Data : [dataSet1, dataSet2, etc]`. The different data sets are then concatenated in a single one dimensional `numpy` array in memory. In that case, there must be as much `Model` delimiters in the input file as data sets, even if the model is the same to represent all data sets. In the case where the program is external, the model will write at every iteration a temporary input file (`write_tmp_input_file`) that automatically updates the inputs of the model according to the Markov chain.

The `Sampler` class takes as arguments a `ProbabilityDistribution` and a set of inputs that identifies the sampling method to use. The different sampling methods are based on the generic `MetropolisHastings` class. The `ProbabilityDistribution` can be any “labeled” distribution that is implemented, or can be a `BayesianPosterior` class.

6.4.2 MetropolisHastings class

The Metropolis–Hastings algorithms introduced in Chapter 5 are implemented and based on the general `MetropolisHastings` class:

```
class MetropolisHastings:
    def __init__(self , IO_util , nIterations , param_init , V,
               prob_distr , opt_arg):

    def run_algorithm(self):

    def compute_new_val(self):

    def compute_acceptance_ratio(self):

    def accept_reject(self):

    def compute_multivariate_normal(self):
```

The `MetropolisHastings` class is based on three main methods that are: `compute_new_val`, `compute_acceptance_ratio` and `accept_reject`. They appear sequentially in the `run_algorithm`, which is the main function to call when using the Metropolis–Hastings algorithm. The `run_algorithm` method of `MetropolisHastings` class above is the random-walk Metropolis–Hastings loop.

More advanced Metropolis–Hastings algorithms are implemented, that are: the adaptive Metropolis–Hastings, the delayed-rejection Metropolis–Hastings, and the delayed-rejection adaptive Metropolis–Hastings. In these cases, the `run_algorithm` method is overwritten in order

to include the additional steps of the methods, compared to the RWMH. For the adaptive Metropolis–Hastings, a fourth method `adapt_covariance` is defined, while for the delayed-rejection Metropolis–Hastings, the `accept_reject` step is redefined by adding the extra delayed-rejection step. The delayed-rejection adaptive algorithm is just a combination of the two previous ones.

Note on the `compute_acceptance_ratio` step

The value of the likelihood function can decrease rapidly when moving away from its maximum, leading to values in the acceptance ratio that can be order of magnitudes different. The logarithm of the probability density is usually computed in the acceptance ratio, leading to

$$\log(r) = \log(\pi(\mathbf{m}^*|\mathbf{y}^{\text{obs}})) + \log(J(\mathbf{m}^*|\mathbf{m}^n)) - \log(\pi(\mathbf{m}^n|\mathbf{y}^{\text{obs}})) - \log(J(\mathbf{m}^n|\mathbf{m}^*)) \quad (6.1)$$

or, when the target distribution is the Bayesian posterior

$$\begin{aligned} \log(r) = & \log(\pi(\mathbf{y}^{\text{obs}}|\mathbf{m}^*)) + \log(\pi^0(\mathbf{m}^*)) + \log(J(\mathbf{m}^*|\mathbf{m}^n)) \\ & - \log(\pi(\mathbf{y}^{\text{obs}}|\mathbf{m}^n)) - \log(\pi^0(\mathbf{m}^n)) - \log(J(\mathbf{m}^n|\mathbf{m}^*)) \end{aligned} \quad (6.2)$$

Thus, the log value of a distribution is usually implemented within every `ProbabilityDistribution` class and can be accessed through the `compute_log_value` method. The value of the acceptance ratio is finally computed by taking the exponential of Eq. (6.2) or Eq. (6.1).

6.4.3 GradientBasedMCMC

The implemented gradient-informed algorithms are the Hamiltonian Monte Carlo and the algorithm based on an ISDE. They are defined based on the `GradientBasedMCMC` class, that inherits from the `MetropolisHastings` class and defines two new attributes.

```
class GradientBasedMCMC(MetropolisHastings):
    def __init__(self, IO_fileID, nIterations, param_init, V,
                prob_distr, C_matrix, gradient, opt_arg):
```

The `C_matrix` attribute contains the options for the covariance matrix, and the `gradient` attribute specifies the method for computing the gradients (either analytical or numerical). The `C_matrix` attribute corresponds to the `covariance` keywords, under the `Algorithm` key, from the input file, and has the following options:

```
"Algorithm": {
    "covariance": {
        "value": "Hessian_diag",
        "starting_it": 1e1,
        "update_it": 1e2,
        "end_it": 1e3
    }
}
```

The keyword `value` refers to the initial value of the covariance matrix, that can be `Hessian` (we take the inverse of the Hessian matrix, computed by means of a finite difference formula), `Hessian_diag` (idem, but we only use the main diagonal), `Identity` (we take the identity matrix) or `from_file` (in which case the covariance must be provided in a “cov_init.csv” file within the working directory). The keywords `starting_it`, `update_it`, and `end_it` are options related to the adaptation of the covariance matrix, and refer to the initial iteration for adaptation, the period of adaption and the final iteration at which the covariance is adapted, respectively. If the value of `starting_it` is greater than the value of `end_it`, then no adaptation of the covariance matrix is performed. It is not implemented currently, but we note that the `C_matrix` attribute could be a common feature shared by all MCMC algorithms.

6.4.4 Gradient of the log-likelihood

For gradient-informed algorithms, the gradient of the logarithm of the posterior distribution (or more generally the potential function) will be required and usually, variables will be expressed in a transformed space following a change of variables. We denote the forward and inverse mappings by $\mathbf{m} \rightarrow f(\mathbf{m})$ and $\tilde{\mathbf{m}} \rightarrow g(\tilde{\mathbf{m}})$ such that $\tilde{\mathbf{m}} = f(\mathbf{m})$ and $\mathbf{m} = g(\tilde{\mathbf{m}})$, respectively. Denoting ϕ the log-likelihood function, we have

$$\begin{aligned} \nabla_{\tilde{\mathbf{m}}} \phi(\mathbf{y}(g(\tilde{\mathbf{m}}))) &= \nabla_{\mathbf{x}} \phi(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{y}(g(\tilde{\mathbf{m}}))} \nabla_{\mathbf{z}}(\mathbf{y}(\mathbf{z})) \Big|_{\mathbf{z}=g(\tilde{\mathbf{m}})} \nabla_{\tilde{\mathbf{m}}}(g(\tilde{\mathbf{m}})) \\ &= \underbrace{\nabla_{\mathbf{x}} \phi(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{y}(g(\tilde{\mathbf{m}}))}}_{\text{log-like gradient}} \underbrace{\nabla_{\mathbf{z}}(\mathbf{y}(\mathbf{z})) \Big|_{\mathbf{z}=g(\tilde{\mathbf{m}})}}_{\text{model gradient w.r.t. parameters}} \underbrace{(\nabla_{\mathbf{m}}(f(\mathbf{m})))^{-1} \Big|_{\mathbf{m}=g(\tilde{\mathbf{m}})}}_{\text{inv jac.}}. \end{aligned} \quad (6.3)$$

When computing gradient numerically (e.g., with finite differences), one can directly compute $\nabla_{\tilde{\mathbf{m}}} \phi(\mathbf{y}(g(\tilde{\mathbf{m}})))$. When analytical gradients of the model are available, the formula above will be preferred for saving computational time and for scalability with the number of parameters.

As an illustration, we consider the Gaussian likelihood function for which the potential ϕ expresses as

$$\phi(\mathbf{y}(\mathbf{m})) = -\frac{1}{2}(\mathbf{y}^{\text{obs}} - \mathbf{y}(\mathbf{m}))\Sigma^{-1}(\mathbf{y}^{\text{obs}} - \mathbf{y}(\mathbf{m})) \quad (6.4)$$

and its gradient

$$\nabla_{\mathbf{x}} \phi(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{y}(\mathbf{m})} = (\mathbf{y}^{\text{obs}} - \mathbf{y}(\mathbf{m}))\Sigma^{-1}. \quad (6.5)$$

Thus, from Eq. (6.3), gradient-informed MCMC methods require the computation of the derivatives of the model with respect to its parameters. This function can either be provided in the `model` file if the gradients are available analytically, or computed numerically. In the latter case, only finite differences are implemented, which may lead to significant approximation errors and may be computationally expensive, in particular in high dimensions or when the model is itself expensive. Future work could investigate the implementation of methods that aim at improving gradient computation, such as adjoint-based methods (Cao et al. [2003]; Bosco [2019]), automatic differentiation methods (Griewank and Walther [2008]), or other estimation methods (Jasra et al. [2021]).

6.5 Illustration: spring model

We consider the example of the spring model described in Smith [2014]. The displacement z (adimensional) as a function of time of the spring without any external excitation is provided by the equations

$$\begin{cases} \ddot{z} + C\dot{z} + Kz = 0 \\ z(0) = 2, \dot{z} = -C \end{cases} \quad (6.6)$$

with $C = c/m$, $K = k/m$, $m \geq 0$ (kg) is the mass, $c \geq 0$ (N s m⁻¹) is the damping coefficient and $k \geq 0$ (N m⁻¹) is the stiffness coefficient. For $C^2 - 4K < 0$, the solution is

$$z(t) = 2 \exp\left(-\frac{Ct}{2}\right) \cos\left(t\sqrt{K - \frac{C^2}{4}}\right) \quad (6.7)$$

For this problem, synthetic data z_i are generated for $0 \leq i \leq n$ with $n = 51$ the number of points in the time interval $[0, 5]$ sec from a random Gaussian distribution $\epsilon_i \sim N(0, \sigma_0^2)$ with $z_i = z(t_i, p_0) + \epsilon_i$ where $\mathbf{m}_0 = (K_0, C_0)$ the nominal parameter values with $K_0 = 20.5$, $C_0 = 1.5$ and $\sigma_0 = 0.1$.

6.5.1 Model file

The model, the parameters and its solution are specified in the `spring_model.py` file that we describe below.

```
import numpy as np
from pybitup import bayesian_inference as bi

class SpringModel(bi.Model):
    """ Class for the spring model """

    def __init__(self, x=[], param=[], name=""):
        # Initialize parent object ModelInference
        bi.Model.__init__(self, name=name)

    def set_param_values(self):
        # Parameters
        self.C = self.param[0]
        self.K = self.param[1]

        # Variable
        self.time = self.x

    def fun_x(self):
        # Get the parameters
        self.set_param_values()

        # Return the function evaluation
        return 2 * np.exp(-self.C * self.time / 2) * np.cos(np.sqrt(
            self.K - self.C**2 / 4) * self.time)
```

Exponential, cosine and square root functions are imported from the `numpy` package. The `bayesian_inference.py` file, where the `Model` class is defined, needs to be imported from `Pybitup`. Our model for the spring that we call here `SpringModel` is based on the general `Model` class and inherit from two (aside from `__init__`) essential methods, namely `set_param_values()` and `fun_x()`. On the last line is the solution from Eq. 6.7 of the spring model. The parameters C and K are assigned in the method `set_param_values()`. Because the values of C and K will change during the calibration as we will see later, this method is run every time the model is evaluated.

6.5.2 Input files

Now let's have a look at the general input file, that can be identified by the `.json` extension. We first have a look at the first keyword `Sampling` which contains the inputs for the sampling part. In this case, we want to sample from a distribution computed from Bayes' formula and is identified by the `BayesianPosterior` keyword. Mandatory keywords for `BayesianPosterior` are `Data`, `Model`, `Prior` and `Likelihood`. The experimental data are first defined under the `Data` keyword:

```
"Data": [
    {
        "Type": "ReadFromFile", // ReadFromFile, GenerateSynthetic
        "FileName": "spring_model_data",
        "xField": ["time"],
```



```

        "yField": ["d"],
        "sigmaField": ["std_d"],
        "n_runs": 1
    }
],

```

The data are read from files with general name `spring_model_data`. The `n_runs` keyword denotes the number of experimental data files that we have, which is only one here. Thus, the input file is named `spring_model_data_0.csv` and must appear in the case directory. Below is the structure of the dataset of the `spring_model_data_0.csv` file and for compactness, we only show here only the 10 first lines.

```

,time,d,std_d
0,0.0,1.9320285551921579,0.1
1,0.1,1.7158980306494984,0.1
2,0.2,1.0116134181643508,0.1
3,0.30000000000000004,0.43030863754766857,0.1
4,0.4,-0.1967505502959701,0.1
5,0.5,-0.8165563476714192,0.1
6,0.6000000000000001,-1.179212035121202,0.1
7,0.7000000000000001,-1.1151949721017202,0.1
8,0.8,-0.9116380936755248,0.1

```

The first column is the indices of the data. Following columns are the fields whose names appear in the input file, namely the time, the displacement of the spring `d` and the experimental standard deviation on `d`, `std_d`. The entries inside the input file and the data file must correspond. For this illustration, the data encoded in the `.csv` file are generated synthetically.

Going back to the input file, the `Model` keyword specifies all the parameters of the model, their nominal values and if the model will be parameterized during the calibration, which is not the case here.

```

"Model": [
    {
        "param_names": ["C", "K"],
        "param_values": [1.5, 20.5],
        "parametrization": "no"
    }
],

```

There are only two parameters here that are C and K . The nominal values of the model parameters will be replaced if those parameters are considered to be unknown in the calibration process.

Next, `Prior` defines the prior distribution.

```

"Prior": {
    "Distribution": "Mixture",
    "Param": {
        "C": {"initial_val": 1.5, "prior_name": "Uniform", "
            prior_param": [0.0, 100]},
        "K": {"initial_val": 20.5, "prior_name": "Uniform", "
            prior_param": [0.0, 100]}
    }
},

```

The parameters that are unknown in the calibration process appear under the `Param` keyword. We consider the two unknown parameters to be $\mathbf{m} = (C, K)$, for which an initial value must

be specified to start the calibration process. The **Mixture** distribution means that we consider the tensor product distribution of the two marginal distributions, assumed to be independent. Here, the marginal distributions are supposed to be uniform PDF in the domain $[0, 100]$.

The **Likelihood** keywords defines the likelihood function.

```
"Likelihood": {
    "function": "Gaussian",
    "distance": "L2",
    "theta": 1
}
```

It is assumed to be a **Gaussian** function with the weighted L2 distance between the model and the data in the argument of the exponential term. The keyword **theta** is an additional free parameter that appears in the argument of the likelihood, which will be described later in Chapter 7. So far, it is assumed to be equal to one and it does not influence the posterior.

The Bayesian posterior distribution will be sampled using the AMH that is specified under the **Algorithm** keyword.

```
"Algorithm": {
    "name": "AMH",
    "AMH": {
        "starting_it": 1e2,
        "updating_it": 1e1,
        "eps_v": 0.0
    },
    "n_iterations": 1e4,
    "proposal": {
        "name": "Gaussian",
        "covariance": {
            "type": "diag",
            "value": [0.0345, 0.7071]
        }
    }
}
```

The AMH adapts the covariance matrix of the proposal during the iterations. It requires additional options specified under the **AMH** keyword, namely the number of iterations at which the adaptation starts (**starting_it**), the frequency at which we update the covariance (**updating_it**) and a last parameter that controls the covariance adaptation (**eps_v**) which is set here to zero. (**n_iterations**) specifies the number of iterations (samples) that the algorithm will generate, which is here equal to 10^4 . For Metropolis–Hastings class algorithms, we need to defined the proposal covariance matrix. Here, it is a Gaussian function with a diagonal covariance matrix with values $[0.0345, 0.7071]$ on the main diagonal. The covariance matrix is a tuning parameter that is not known a priori and may require several trial-and-error tests before finding an adequate guess. Here, we choose it to be diagonal for convenience. However, thanks to the AMH algorithm, this covariance matrix will be adapted with the iterations and will improve the sampling efficiency.

6.5.3 Run file

The file that needs to be run by **Python** to perform the calibration is named **run.py**. It gathers the information from the input file, the model and it is where we select what function from **Pybitup** we want to run. In this case, we want to sample from the posterior distribution and post process the results. First, we to import the **spring_model** file as well as **Pybitup**.

```

import spring_model
import pybitup

case_name = "spring_model"
input_file_name = "{}.json".format(case_name)

# Define the model
my_spring_model = {}
my_spring_model[case_name] = spring_model.SpringModel(name=case_name)

# Sample
post_dist = pybitup.solve_problem.Sampling(input_file_name)
post_dist.sample(my_spring_model)
post_dist.__del__()

# Post process
pybitup.post_process.post_process_data(input_file_name)

```

The model in `Pybitup` is specified in a list where the models are provided. The only model here is the `SpringModel` from the `spring_model.py` file.

All the main capabilities of `Pybitup` (sampling, propagation, sensitivity analysis) are provided in the `solve_problem` file. We thus create an object `post_dist` for the sampling and we provide as input the `.json` file. Next, we use the method `.sample` to which we provide the model list `my_spring_model`. At the end of the sampling, the output file are written in an `/output` folder and we delete the object `post_dist`. The last line is for the post process that will be described later.

6.6 Uncertainty propagation

The main goal of the propagation capability in `Pybitup` is to use the samples from the Bayesian calibration and propagate them through built-in or external computer codes. The direct Monte Carlo can be easily performed in any computing languages and the results analysis only requires efficient post-processing tools. However, when the computational models implemented in the computer codes are expensive, methods that seek the construction of surrogate models are generally required. A PCE library was therefore implemented in `Pybitup`, with the goal of using the samples from the Bayesian posterior, for which the distribution can be non-standard with correlated variables. Nevertheless, it was implemented such that it can be used independently of the sampling module of `Pybitup`, and standard distributions can be assumed. Therefore, the calibration and the propagation can be performed separately.

The PCE library of `Pybitup` was developed in [Lacroix \[2020\]](#) and the capabilities were showed on an uncertainty propagation problem following from the calibration of a simplified pyrolysis model. The main features are discussed below.

The computation of the Vandermonde matrix and the basis by a Gram-Schmidt orthogonalization, which are central elements in most of data-driven PCE related algorithms, is straightforward. For computing a quadrature rule, the code takes advantage of several Numpy and Scipy functions. A method to extract a sparse embedded quadrature rule from the MCMC samples is implemented. Once computed, the polynomials and the PCE coefficients are stored in a compact homogeneous block representation within two classes, one containing the orthogonal polynomial basis and the other containing the surrogate model. The compact homogeneous block representation allows a memory efficient storage as well as the use of simple matrix operations to manipulate the polynomials. It is important to note that `Polynomial_basis` and `Expansion` share the same exponent table and dimension, but store a different coefficient matrix.

6.7 Future developments

Although the ISDE method it is not a Metropolis–Hastings-type algorithm, it inherits from the general class `MetropolisHastings` but does not need to compute the acceptance ratio and the accept-reject step. Future implementation should consider a more general class (e.g. `IterativeAlgorithms` class) that does not have the accept-reject step, on which `MetropolisHastings` and ISDE could be based. Also, all adaptive MCMC methods should share the `C_matrix` attribute, including the AMH, and should be implemented in the future.

While uncertainty propagation is needed to assess the reliability of the predictions, sensitivity analysis is essential to identify key parameters that influence the quantities of interest. As for uncertainty propagation, performing sensitivity analysis when two or more variables are dependent can be challenging, and future work should investigate methods for sensitivity analysis using the samples from the posterior.

The hIPPYlib (Villa et al. [2021]) shares similar sampling capabilities compared to `Pybitup`, and has also several complementary features that would be valuable to benefit from. It has large scale optimization algorithms, that are still lacking in `Pybitup` and which could help for obtaining the initial guess for MCMC chains, or it can perform the linearization of the Bayesian posterior, which information can be used to improve sampling efficiency. As several modules are written in `Python`, the hIPPYlib could be directly used in `Pybitup`.

The post-processing tools of `Pybitup` are limited and future developments should include more efficient methods for posterior analysis, for plotting functionality and methods for MCMC chain convergence diagnostics. Several functions are available in the `Arviz` (Kumar et al. [2019]) `Python` package, that could be used in `Pybitup`.

Part III

Quantification of uncertainty in thermal protection material decomposition

Bayesian identification of pyrolysis model parameters for thermal protection materials

Carbon/phenolic composite materials have gained renewed interest to serve as ablative TPMs and new experimental data relevant to the pyrolytic decomposition of the phenolic resin of lightweight carbon/phenolic TPMs have recently been published in the literature. In this chapter, we infer from these new experimental data uncertainty-quantified pyrolysis models using a Bayesian probabilistic approach to account for uncertainties in the model identification. We use an approximate likelihood function involving a weighted distance between the model predictions and the time-dependent experimental data. In a first step, we assume multi-component pyrolysis reactions and we infer a pyrolysis model for PICA at 366 K min^{-1} . To sample from the posterior, we investigate two gradient-informed MCMC methods, namely the HMC and the method based on an ISDE and we compare the sampling efficiency when we perform a reparameterization of the pyrolysis equations and an adaptation of the covariance matrix approximation. We first show on a two-parameter case that the covariance adaptation improves significantly the exploration of the posterior. To select the decomposition mechanisms to be represented in the pyrolysis model, we proceed by progressively increasing the complexity of the pyrolysis model until a satisfactory fit to the data is ultimately obtained. The pyrolysis model thus obtained involves six reactions and has 48 parameters. In a second step, a pyrolysis model including competitive mechanisms is assumed in order to be able to represent the data at two distinct heating rates. Because of the increased complexity of the model and the absence of a fast computational solution, gradient-informed algorithms are computationally expensive and the pyrolysis model parameters are thus inferred using the RWMH algorithm. We study the correlation between the parameters and we finally assess the predictive capabilities of such pyrolysis model at other heating rates. The preliminary results from Section 7.4 were published in the proceedings of the annual VKI Symposium (Coheur et al. [2018, 2019]). Preliminary results on the ISDE algorithm applied to a simple pyrolysis model were presented at the UNCECOMP conference (June 2019), and results on the covariance matrix adaptive procedure for gradient-informed sampling algorithms were provided in Coheur et al. [2020]. The results for the competitive pyrolysis reaction for PICA were published in Torres-Herrador et al. [2019a].

“If you do not make runs like that [millions or billions of iterations], you are simply not serious about MCMC.”

Charles J. Geyer.

7.1 Introduction

Several factors make it challenging to formulate in a probabilistic setting the identification of a pyrolysis model with Arrhenius-like reaction rates from the new experimental data. Notably, the experimental data are subject to measurement uncertainty and sample variability, and experimentalists can usually provide only a limited characterization of these sources of uncertainty: they repeat their experimental study a small number of times under nominally identical conditions and then use the differences in the observed species productions between the repeated experimental runs to deduce error bars for the experimental data. In addition, for the pyrolysis of the TPMs considered, significant modeling errors can be expected to be present: significant physical mechanisms responsible for species production in the experiment can be expected to be either unmodeled or represented imperfectly, for instance in the multi-component single-step model, or with Arrhenius-like reaction rates. For example, in the initial stage of the experiment, physical mechanisms other than pyrolysis, such as outgassing of gases like water that had been absorbed by the sample, may cause significant species production, or considering distinct heating rates may cause pyrolysis reaction curves to shift which cannot be accounted for by multi-component models. Furthermore, expressing the reaction rates with Arrhenius laws may imperfectly represent the physical behavior. Because of the lack of knowledge of the measurement uncertainty and sample variability and because of the likely presence of significant modeling errors, the model identification cannot be readily formulated as a standard problem of Bayesian statistical inference of a deterministic model augmented with a stochastic representation of experimental noise.

To account for modeling errors in model identification, there exist approaches that seek to incorporate a representation of the model discrepancy explicitly into the model, such as by incorporating additional stochastic terms into the governing equations ([Beck and Katafygiotis \[1998\]](#); [Kennedy and O’Hagan \[2001\]](#)). However, such approaches significantly modify the model: they convert a model in the form of a system of differential equations into a stochastic model in the form of a system of stochastic differential equations. It is then this stochastic model that is identified from the data. In the application that we study, the purpose of the identified pyrolysis model is to serve in material response codes for numerical simulations of heat shield decomposition during atmospheric entries. Such approaches would thus require the identified stochastic pyrolysis model, with the additional stochastic terms inherent to it, to be implemented within such numerical tools, thus turning these numerical simulations themselves into stochastic simulations.

In this work, we adopt an alternative approach whereby we keep the structure of the pyrolysis model unchanged (we do not introduce additional stochastic terms) and we seek instead to incorporate a representation of the uncertainties into the model parameters. We formulate the model identification as a Bayesian probabilistic inverse problem, in which we use an approximate likelihood function involving a weighted distance between the model predictions and the experimental data. This formulation provides a probability distribution over the model parameters, the posterior, which assigns higher probabilities to parameter values that correspond to predictions that better fit the experimental data. We introduce a free parameter in this approximate likelihood function that allows us to control how significantly the probability of parameter values is penalized in proportion to the misfit between the corresponding predictions and the experimental data. This approach allows the free parameter to be tuned to seek to introduce

into the model parameters an amount of uncertainty that can serve as an overall representation of the impact of various sources of uncertainty, including measurement uncertainty, sample variability, and modeling errors. In this work, we rely on an ad hoc, visual, criterion to tune the free parameter.

The number of model parameters of the pyrolysis model is high, and the posterior exhibits strong dependencies, so that sampling from the posterior is challenging. Notably, pre-exponential factors and activation energies involved in Arrhenius laws exhibit significant dependencies because of their complementary role in the expression of Arrhenius law. This high correlation is called the kinetic compensation effect (Koga [1994]; Galwey and Brown [1999]) and has already been observed in numerous deterministic kinetics studies, as in Rodionova and Pomerantsev [2005]; Pomerantsev et al. [2017] or Torres-Herrador et al. [2019a], as well as in probabilistic inference studies, see for instance in Najm et al. [2014] and Bruns [2015]. Najm et al. [2014] and Bruns [2015] applied adaptive random-walk Metropolis–Hastings algorithms, with a reparameterization (change of variables) of the pre-exponential factor using a logarithmic change of variables, to sample efficiently from posteriors over parameters of Arrhenius laws. We thus investigate a similar reparametrization of the pyrolysis equations in order to simplify the structure of the posterior distribution we sample from. However, our reparameterization introduces additional scaling factors that need to be provided and which are unknown. Therefore, we also investigate two gradient-informed methods, namely the Hamiltonian Monte Carlo (HMC) method (Neal [2010]; Gelman et al. [2014]) and a method based on an Itô stochastic differential equation (ISDE) (Soize [2008, 2015, 2017]). These methods, at least in the variants used in this chapter, require an estimate of the covariance matrix of the posterior, which we propose to learn adaptively from previous samples in the chain, similarly to adaptive Metropolis–Hastings algorithms.

To simulate the pyrolysis of TPMs, two types of model are used. The first one is a multi-component model that is built using the experimental data of PICA decomposition at 366 K min^{-1} . For multi-component models, a semi-analytical expression was derived in Section 3.3, which is computationally cheap, and millions of MCMC iterations can be performed to ensure that all the support of the posterior is sampled. This allows to study the efficiency of the proposed improvements using the change of variable and the sampling strategies. Multi-component models may lack of generality when used at significantly different heating rates, and therefore a second model that features a competitive mechanism is investigated. It allows to account for the decomposition of the PICA material at two distinct heating rates (10 K min^{-1} and 366 K min^{-1}), and we present a probabilistic inference of the model parameters. The main focus is to establish the relevance of accounting for competitive mechanisms in the pyrolysis model in order to reduce the model error structure when distinct heating rates are considered. Because the model with competitive mechanisms does not have an analytical solution, the system of ODEs is solved by means of numerical methods, which increases the computational time of the Bayesian inference. Only the RWMH MCMC method is applied, but the convergence of the algorithm is enhanced by providing the initial guess from a genetic algorithm estimation, thus allowing to run a low number of MCMC samples.

The chapter is organized as follows. In Section 7.2, the Bayesian probabilistic formulation, the expression for the distance in the approximate likelihood function, the choice of the weights and the change of variables are described. The methods for approximating the covariance matrix and the adaptive procedure are then presented in Section 7.3. Next in Section 7.4, numerical results are provided, including a study of computational aspects and efficiency of the employed adaptive sampling algorithms. Using the experimental data from Bessire and Minton [2017], the case for the production of a single species, namely, CH_4 is first considered in Section 7.4.1 in order to have a simple application for which we can show the proposed improvements and for which the posterior PDF shows a complex structure (correlated and non-Gaussian). In Section 7.4.2, the proposed methodology is applied to the pyrolysis of PICA at high heating rates using the

whole experimental data set, and the uncertainties on the kinetic parameters and the species production for the PICA material are quantified. Finally, in Section 7.5, the results are shown and analyzed for the pyrolysis scheme including the competitive mechanisms.

7.2 Bayesian approach to inverse problem for parameter calibration

Let $\eta(\cdot, \mathbf{m}) : \mathbb{R}^+ \rightarrow \mathbb{R}^{n_{\text{obs}}}$ be the time-dependent solution to the computational model representing the dynamic physical process, with the vector of model outputs $\eta(\cdot, \mathbf{m}) = (\eta_1(\cdot, \mathbf{m}), \eta_2(\cdot, \mathbf{m}), \dots, \eta_{n_{\text{obs}}}(\cdot, \mathbf{m}))$ with values in $\mathbb{R}^{n_{\text{obs}}}$, provided the model parameters $\mathbf{m} \in \mathbb{R}^{n_u}$. The value predicted by the model for a given observable indexed by i , $1 \leq i \leq n_{\text{obs}}$, at time $t = t_j$, $1 \leq j \leq n_t$, is denoted by

$$y_{ij} = \eta_{ij}(\mathbf{m}) = \eta_i(t_j, \mathbf{m}). \quad (7.1)$$

The different observables can be the different species produced by the pyrolysis reaction, e.g. $n_{\text{obs}} = N_g$. We refer, for instance, to Bruns [2015] for an example where the observable is the mass loss. In Section 7.4, we will use the individual species production because we expect that it provides more information on the different reaction steps and the reaction mechanisms. In Section 7.5, we will use both the mass loss and the derivative of the mass loss, depending on the experimental data sets considered.

We formalize the results of the experiments (described in Section 3.2.2) as follows. We denote by n_r the number of repetitions of the experiment. The data are acquired at given time values that depend on the acquisition frequency, which is fixed by the experimental device, and results in n_t data points for each observation (for instance, at $\beta = 366 \text{ K min}^{-1}$, $n_t = 69$, see Fig. 3.3). The different experimental observations (species, mass loss) as a function of time, or temperature, with n_{obs} the total number of experimental observations, can be specified in a third-order tensor of observations \mathbf{Y}^{obs} of dimensions $n_{\text{obs}} \times n_t \times n_r$. The corresponding k -th matrix of observations is thus denoted

$$Y_{::k}^{\text{obs}} = \begin{bmatrix} y_{11k}^{\text{obs}} & y_{12k}^{\text{obs}} & \cdots & y_{1n_t k}^{\text{obs}} \\ y_{21k}^{\text{obs}} & y_{22k}^{\text{obs}} & \cdots & y_{2n_t k}^{\text{obs}} \\ \vdots & \vdots & \ddots & \vdots \\ y_{n_{\text{obs}}1k}^{\text{obs}} & y_{n_{\text{obs}}2k}^{\text{obs}} & \cdots & y_{n_{\text{obs}}n_t k}^{\text{obs}} \end{bmatrix}, \quad (7.2)$$

with $1 \leq k \leq n_r$. In particular, in Bessire and Minton [2017], the experiments were repeated three times ($n_r = 3$) and a mean and a standard deviation as a function of temperature were estimated for each observation (see Fig. 3.3):

$$\bar{y}_{ij}^{\text{obs}} = \frac{1}{n_r} \sum_{k=1}^{n_r} y_{ijk}^{\text{obs}}, \quad (7.3)$$

$$s_{ij}^2 = \frac{1}{n_r - 1} \sum_{k=1}^{n_r} (y_{ijk}^{\text{obs}} - \bar{y}_{ij}^{\text{obs}})^2. \quad (7.4)$$

with \bar{Y}^{obs} the matrix of mean observations and S the matrix of the variances of the observations.

We adopt a Bayesian approach to inverse problems in which the solution to the inverse problem is the probability distribution of the unknown parameters \mathbf{m} given the value of the experimental observations \mathbf{Y}^{obs} with the following density:

$$\pi(\mathbf{m} | \mathbf{Y}^{\text{obs}}) \propto \exp \left(-\Phi \left(\mathbf{Y}^{\text{obs}}, \eta(\mathbf{m}) \right) \right) \pi^0(\mathbf{m}), \quad (7.5)$$

where the symbol \propto is interpreted as to mean “proportional up to a normalization factor independent of \mathbf{m} ”. The exponential factor is an approximate likelihood function and depends on a potential function Φ , or misfit function, that seeks to gauge the goodness of fit between the actual observations and the model output. The second factor $\pi^0 : \mathbb{R}^{n_u} \rightarrow \mathbb{R}^+$ is the prior and acts as a regularization that can incorporate any prior knowledge about the parameters \mathbf{m} .

7.2.1 Prior density function

The only prior knowledge we have about the parameters is that \mathcal{A}_i , \mathcal{E}_i and r_i are positive and the mass fractions F_{ij} are bounded between 0 and 1. We choose the following prior:

$$\pi^0(\mathbf{m}) = \prod_{i=1}^{N_p} 1_{\mathbb{R}^+}(\mathcal{A}_i) 1_{[L_{\mathcal{E}_i}, U_{\mathcal{E}_i}]}(\mathcal{E}_i) 1_{[1,10]}(r_i) \prod_{j \in I_g} 1_{[0,1]}(F_{ij}), \quad (7.6)$$

where 1_S is the indicator function of $S \subseteq \mathbb{R}$. The activation energy is strictly positive, but in cases where several pyrolysis reactions are considered, the values of \mathcal{E}_i are further limited to the bounded domain $[L_{\mathcal{E}_i}, U_{\mathcal{E}_i}]$ in order to avoid the different pyrolysis reactions to mix together, which would make the potential function more complex to explore. For multi-component parallel models, the maximum values for r_i are limited to 10, as higher values were not found in previous similar pyrolysis models (Torres-Herrador et al. [2019b]), and first-order reactions ($r_i = 1$), for which the analytical solution is different, are not considered here (Coheur et al. [2021b]). When competitive schemes are included, we fix $r_i = 1$, $1 \leq i \leq N_p$, which is considered to be known.

7.2.2 Choice of the approximate likelihood function

In Eq. (7.5), the goodness of fit is assessed through the approximate likelihood function $\exp(-\Phi(\mathcal{Y}^{\text{obs}}, \boldsymbol{\eta}(\mathbf{m})))$. We define the potential function using the sum over all the observations of the squares of the weighted norms of the differences between the model output and the observations, similarly to a weighted least-squares formulation in optimization,

$$\Phi(\mathcal{Y}^{\text{obs}}, \boldsymbol{\eta}(\mathbf{m})) = \frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \phi_i(\bar{\mathbf{y}}_{i:}^{\text{obs}}, \boldsymbol{\eta}_{i:}(\mathbf{m})) = \frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \|\bar{\mathbf{y}}_{i:}^{\text{obs}} - \boldsymbol{\eta}_{i:}(\mathbf{m})\|_{W_{::i}}^2, \quad (7.7)$$

where $\|\mathbf{x}\|_W = \mathbf{x}^\top W \mathbf{x}$ denotes the weighted norm of a vector $\mathbf{x} \in \mathbb{R}^{n_t}$ and $W_{::i}$ is the $\mathbb{R}^{n_t \times n_t}$ symmetric positive definite matrix of weights, formed from the i -th subarray of the third-order tensor of weights \mathcal{W} . If we assume further that the weighting matrices $W_{::i}$ are diagonal, i.e. $w_{kji} = 0, \forall k \neq j, 1 \leq k \leq n_t, 1 \leq j \leq n_t$, then we arrive at the following function for the likelihood

$$\exp(-\Phi(\mathcal{Y}^{\text{obs}}, \boldsymbol{\eta}(\mathbf{m}))) = \exp\left(-\frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \sum_{j=1}^{n_t} w_{jji} (\bar{y}_{ij}^{\text{obs}} - \eta_{ij}(\mathbf{m}))^2\right), \quad (7.8)$$

where $\bar{y}_{ij}^{\text{obs}}$ is computed from Eq. (7.3).

In classical optimization problems, the weights w_{jji} can be used to rescale the different physical quantities, or they can be used to give more weight to certain data rather than other. Equal importance is specified with $w_{jji} = 1, \forall i, j$. In the Bayesian approach, larger values of the weights result in a posterior that is more concentrated around the best-fit parameter values.

To choose the values of the w_{jji} , we will introduce two normalizations. The first normalization takes into account the observed dispersion (the s_{ij} in Eq. (7.4)) by choosing the weights w_{jji} inversely proportional to s_{ij} . The second normalization takes into account the number of data acquired in the time series n_t and is chosen proportional to $(\Delta t_j / t_f)$, where t_f is the total time

span of the experiment and Δt_j are the data acquisition time, and, for a constant acquisition time, the ratio $\Delta t_j/t_f = n_t^{-1}$. This second normalization prevents the sum in the exponential from becoming very high if the number of data acquired is high, as in this case Δt_j tends to 0. The presence of this factor proportional to the number of terms in the sum yields a posterior that is less dependent on the data acquisition frequency. These normalizations are further discussed in Appendix C.1.

We set $w_{ij} = \vartheta s_{ij}^{-2}(\Delta t_j/t_f)$ with s_{ij} computed from Eq. (7.4) and ϑ a free (tuning) parameter that must be chosen. This choice finally leads to the following expression for the likelihood function:

$$\exp\left(-\Phi\left(\mathcal{Y}^{\text{obs}}, \boldsymbol{\eta}(\boldsymbol{m})\right)\right) = \exp\left(-\frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \sum_{j=1}^{n_t} \vartheta \frac{\Delta t_j}{t_f} \left(\frac{\bar{y}_{ij}^{\text{obs}} - \eta_{ij}(\boldsymbol{m})}{s_{ij}}\right)^2\right). \quad (7.9)$$

The free parameter ϑ controls all the weights at the same time and will influence the width of the resulting posterior. Later in this work, we will rely on an ad hoc, visual, criterion to tune the free parameter, so that we introduce in the model parameters an amount of uncertainty that represents an overall impact of measurement uncertainty, sample variability, and modeling errors.

7.2.3 Change of variables

In Section 5.4.6, we discussed the advantage of using a change of variables on the convergence and mixing of Markov chains. We presented a first change of variables that transforms the problem based on the support of the posterior and was used to make the ISDE feasible by mapping the initial parameter set into a transformed unbounded parameter space.

We propose a second change of variables that is based on the structure of the pyrolysis equations (3.9)–(3.15), that seeks to decouple the strong kinetic compensation effect characteristic of the Arrhenius equations. A proposed change of variables based on the structure of the pyrolysis model was introduced in Coheur et al. [2019]. It is following the reparametrization for \mathcal{A}_i and \mathcal{E}_i proposed in Rodionova and Pomerantsev [2005] and Pomerantsev et al. [2017] that was applied for deterministic parameter estimation and that we extend to the pyrolysis of TPMs by introducing two additional scaling factors for r_i and F_{ij} . The transformation is such that

$$\tilde{\mathcal{A}}_i = \ln(\mathcal{A}_i \bar{t}_i) - \mathcal{E}_i/(\mathcal{R} \bar{T}_i), \quad (7.10)$$

$$\tilde{\mathcal{E}}_i = \mathcal{E}_i/\bar{\mathcal{E}}_i, \quad \tilde{r}_i = r_i/\bar{r}_i, \quad \tilde{F}_{ij} = F_{ij}/\bar{F}_{ij}, \quad (7.11)$$

where we introduced the scaling factors $\bar{\boldsymbol{m}} = \{\bar{t}_i, \bar{\mathcal{E}}_i, \bar{r}_i, \bar{F}_{ij}\}$ and the reference temperature \bar{T}_i . Considering this change of variables, reaction rates are now expressed as

$$k_i = \frac{1}{\bar{t}_i} \exp\left(\tilde{\mathcal{A}}_i + \tilde{\mathcal{E}}_i \tilde{T}_i(t)\right), \quad (7.12)$$

where $\tilde{T}_i = \bar{\mathcal{E}}_i/\mathcal{R}(1/T + 1/\bar{T}_i)$ is called the local reciprocal temperature (Pomerantsev et al. [2017]). Another normalization common in combustion literature was also discussed by Urzay et al. [2012] for its use in uncertainty quantification, but this normalization does not consider the logarithmic dependence of the pair \mathcal{A}_i and \mathcal{E}_i , or the translation used in Eq. (7.10).

The effect of the reparameterization in Eqs. (7.10)–(7.11) on the sampling of the posterior was investigated in Coheur et al. [2019]. On the right-hand side of Eq. (7.10), the logarithm in the first term $\log(\mathcal{A}_i \bar{t}_i)$ seeks to suppress the non-linear dependence between \mathcal{A}_i and \mathcal{E}_i . The second term on the right-hand side then attempts to remove the linear dependence between $\log(\mathcal{A}_i \bar{t}_i)$ and \mathcal{E}_i that remains. Because most of the dependence between \mathcal{A}_i and \mathcal{E}_i has thus been removed and \mathcal{A}_i and \mathcal{E}_i are now of comparable orders of magnitude, the jumps in the MCMC algorithms

will be more efficient, in particular for RWMH-type algorithms in which the proposal covariance matrix will be easier to tune.

An appropriate choice of the scaling factors $\overline{\mathbf{m}}$ is important and we strive to tune the reparametrization so that it has the desired effect of rescaling the posterior density around parameters with unit values. We will show in Section 7.4.1 on a single pyrolysis process that this reparametrization improves significantly the convergence and mixing of the Markov chain for the RWMH and HMC algorithms and facilitates the process of estimating the covariance matrix for the proposal function with very limited tuning. However when the dimension increases, the interactions between the parameters and the number of scaling factors to tune make it difficult to find their appropriate value.

7.3 Adaptive-gradient informed sampling algorithms

The covariance matrix can be seen as a representation of the scales and correlations of the problem. The physical parameters have their own scale where they can take their values and can differ significantly one from another, such as for instance \mathcal{A}_i or \mathcal{E}_i compared to F_{ij} or r_i , for which the reparameterization can already help to reduce scales and correlations. Moreover, the experimental data can also provide more information on some parameters than others, resulting in more uncertainty and thus more variability in those parameters that will differ in scale. It is crucial for the mixing and convergence of the Markov chain toward the posterior distribution to have a good estimation of the covariance matrix.

How the approximation of the covariance matrix can help in the mixing of the MCMC chains depends on the type of algorithm considered. For the RWMH, the covariance matrix appears in the proposal distribution from Eq. (5.38) and is sought to enable an efficient sampling from the posterior distribution. The covariance of the proposal affects directly the size of the parameter jumps between two states in the Markov chains that we want to be large (low) for the parameters with high (low) uncertainty and large (low) scales while keeping a reasonable acceptance probability. A similar behavior is of interest for gradient-informed algorithms such as the HMC and the ISDE. From Eq. (5.47) and Eq. (5.55), the momentum variable is a Gaussian random vector whose covariance matrix is the symmetric positive-definite matrix \mathbf{M} ; thus, the linear transformation $d\mathbf{Q}/dt = \mathbf{M}^{-1}\mathbf{P}$ is a Gaussian random vector with covariance matrix $\mathbf{M}^{-1}\mathbf{M}\mathbf{M}^{-T} = \mathbf{M}^{-T} = \mathbf{M}^{-1}$. Hence, the matrix \mathbf{M}^{-1} can also be interpreted in terms of the scales and correlations between two jumps in the position variable. For both the ISDE and HMC, we will relate the mass matrix to the symmetric positive-definite matrix $\mathbf{C} = \mathbf{M}^{-1}$ for which an efficient tuning can also be desirable to improve the sampling efficiency.

For the tuning of the proposal in Metropolis–Hastings algorithms, [Girolami and Calderhead \[2011\]](#) use a mechanism that learns the local structure of the log-posterior distribution at the current state of the Markov chain for the Metropolis-adjusted Langevin and the HMC algorithms. [Law \[2014\]](#) uses proposals that include a local estimation of the value of the Hessian of the sign-reversed logarithm of the posterior, which is representative of the local scales and correlations. However, this position-dependent approach requires modifications to the differential equations of the dynamical system of the proposal mechanism ([Girolami and Calderhead \[2011\]](#); [Arnst et al. \[2017\]](#)). For the ISDE algorithm, [Arnst et al. \[2017\]](#) and [Arnst and Soize \[2019\]](#) use a fixed covariance approximation $\hat{\mathbf{C}}$ based on the value of the Hessian of the sign-reversed logarithm of the posterior at the maximum likelihood estimate. This estimate of the covariance matrix can sometimes lack accuracy, for instance, if the Hessian is estimated at a parameter value that is far from the maximum likelihood value (because, for instance, we do not know it) and where the approximation of the covariance may result in a matrix that is not positive-definite. To avoid this, the value of $\hat{\mathbf{C}}$ can be set to the identity matrix, $\hat{\mathbf{C}} = \mathbf{I}$ (we do not provide any information about scales and correlations), or we can also use the diagonal of the Hessian estimate that can be easily made positive-definite by taking the absolute value of the diagonal elements (it represents

only the local scales and not the correlations). Depending on the choice of the covariance matrix estimation, the value for the timestep τ will have to be chosen carefully.

The tuning of the covariance matrix is not the only way to improve the sampling efficiency. The adaptation of the parameters of the discretization scheme can also be performed (Hoffman and Gelman [2014]; Marshall and Roberts [2011]), as it is done for instance with the no-U-turn sampler (Hoffman and Gelman [2014]) for the HMC algorithm that refines the step size τ and the number of leapfrog steps L at each iteration, or both the timestep and the covariance matrix for the Metropolis-adjusted Langevin algorithm in Marshall and Roberts [2011], but this is not investigated here.

7.3.1 Adaptive procedure for the estimation of the covariance matrix

In order to have a more global estimation of the scales and correlations, which is also robust independently of the covariance matrix estimate, we propose to use an adaptive procedure in order to refine the initial covariance estimation. Adaptation of the covariance matrix for gradient-informed algorithms was already used in Atchadé [2006] and Marshall and Roberts [2011] for the Metropolis-adjusted Langevin algorithm. Their adaptation consists in a normalized truncated covariance estimation multiplied by a scaling factor that is monitored using a criterion based on the acceptance probability. Similarly, we use here an adaptive procedure for the covariance matrix that we apply to the ISDE and HMC algorithms. The covariance matrix is estimated from past sampled values, as in the adaptive Metropolis–Hastings from Haario et al. [2001]; Atchadé and Rosenthal [2005]; Andrieu and Thoms [2008]; Roberts and Rosenthal [2009] and Rosenthal [2010]. The adaptation starts after an initial burn-in phase and the whole sample history is used in the covariance computation. The adaptation updates the proposal covariance matrix based on previous iterates at given increments Δn . The covariance matrix is set initially to \hat{C}_0 during the first n_0 iterations before being updated. We denote by n_0 the iteration at which we start the adaptive procedure and by Δn_{adapt} the period at which the covariance is adapted. The first adaptation is based on the covariance estimation of the n_0 iterations. Then we apply a recursive formula for estimating the covariance as a function of each iteration to avoid saving all the sample values (Haario et al. [2001]). Considering that we have just sampled a new value \mathbf{m}^{n+1} from the posterior distribution with the covariance estimate \hat{C}_n , we compute the new covariance matrix estimate \hat{C}_{n+1} using the $(n+2)$ from Eqs. (5.41)–(5.43), where s_d is an additional scaling parameter that depends only on the dimension n_u of the parameter space, $\epsilon > 0$ is a correcting parameter and I_{n_u} is the n_u -dimensional identity matrix. The parameter s_d is used to improve the acceptance rate in Metropolis–Hastings algorithms and it was showed in Gelman et al. [1996] that the optimal scaling for a Gaussian distribution was $s_d = 2.38^2/n_u$. We also keep this parameter as it was also observed that the posterior distribution was sampled more efficiently using it. The correcting parameter $\epsilon > 0$ is used to ensure that the covariance matrix estimate will not become singular, but it can be set to zero in practice (Haario et al. [2001]).

Since the update of the covariance matrix depends on previous iterations, adaptive algorithms are by essence non-Markovian and the challenge in developing such algorithms requires to show that the chain has the correct ergodic properties, i.e. that it converges toward the right distribution and is stationary (Haario et al. [2001]; Atchadé and Rosenthal [2005]; Andrieu and Moulines [2006]). Such a proof for the HMC and the ISDE is left as an interesting direction for future work¹.

Instead, we discard the samples drawn during the adaptation phase. Once the adaptation phase has provided a satisfactory estimation of the covariance matrix, we keep the latter fixed and the subsequent iterations will be used for posterior estimations. The adaptive phase can be more

¹This would require the *diminishing adaptation* and the *bounded convergence* conditions to be demonstrated, see examples in Roberts and Rosenthal [2009], and in Marshall and Roberts [2011] for an application to a gradient-informed algorithm, namely the Metropolis-adjusted Langevin algorithm.

or less long, depending on the initial approximation of the covariance matrix, the parameters of the algorithms and the time for the MCMC chain to learn from the whole posterior density. We denote by n_f the last iteration at which the covariance matrix is updated and before which the samples are discarded for posterior estimations. The ISDE algorithm with an initial adaptive phase is presented in Algorithm 3 and it would be similar for the HMC algorithm. The adaptive version of the gradient-informed sampling algorithms are implemented within `pybitup`.

Algorithm 3: ISDE algorithm with an initial adaptive phase of the covariance matrix.

```

Set  $\tau, f_0$ ;
Initialize  $\mathbf{m}^0, \hat{C} = \hat{C}_0$ ;
while  $n < n_{\text{iterations}}$  do
    Sample  $\Delta \mathbf{W}^{(n+1)} \sim \mathcal{N}(0, \tau)$ ;
    Solve Eqs. (5.56)–(5.58) for  $\mathbf{Q}^{(n+1)}, \mathbf{P}^{(n+1)}$ ;
     $\mathbf{m}^{(n+1)} = \mathbf{Q}^{(n+1)}$ ;
    Update covariance estimate from Eqs. (5.41)–(5.43);
    if  $n \geq n_0$  and  $n \leq n_f$  and  $\text{mod}(n, \Delta n_{\text{adapt}}) = 0$  then
        | Update  $\hat{C} = \hat{C}_{n+1}$ ;
    end
end

```

7.4 Application: parameter inference of multi-component pyrolysis models

In this first application section, the probabilistic setting is used for the identification of a multi-component pyrolysis model using experimental data at a single heating rate (366 K min^{-1}). We investigate the effects of the proposed improvements on the calibration of the pyrolysis model parameters and we study the influence of the parameter ϑ . First, in Section 7.4.1, the change of variables, the change of scales and correlations, the different sampling algorithms as well as the adaptive phase are applied to a two-parameter pyrolysis model in order to compare their efficiency, along with a discussion of the challenge of incorporating modeling errors. Then, in Section 7.4.2 we perform the parameter calibration and uncertainty quantification on a pyrolysis model based on the experiments of Bessire and Minton [2017] and we will discuss the challenge in the selection of the pyrolysis model.

7.4.1 Uncertainty analysis for the production of CH_4

We analyze the production of gaseous CH_4 that can be mainly characterized by a single peak in the production curve (Fig. 3.3 top-right). In order to feature the strong kinetic compensation effect that results in a posterior distribution with large scales and correlations and a complex dependence structure, we consider two uncertain parameters that are the pre-exponential factor and the activation energy, denoted by \mathcal{A} and \mathcal{E} , respectively. Furthermore, this two-parameter model allows for the comparison of the results of the sampling methods with the direct evaluation of the posterior distribution from formula Eq. (7.5) with prior from Eq. (7.6) and likelihood function from Eq. (7.9).

Prior to the analysis, we perform a regularization of the initial estimation of the standard deviations s_j with $1 \leq j \leq n_t$ (to simplify the notation, we omit the single i index because we consider only CH_4) which appear in the weights w_j^{-1} . Some values of the standard deviations obtained from the experiments, denoted by $s_j^{(e)}$, have a zero value where the production rate is zero, for instance, at the beginning or at the end of the pyrolysis process, and the resulting

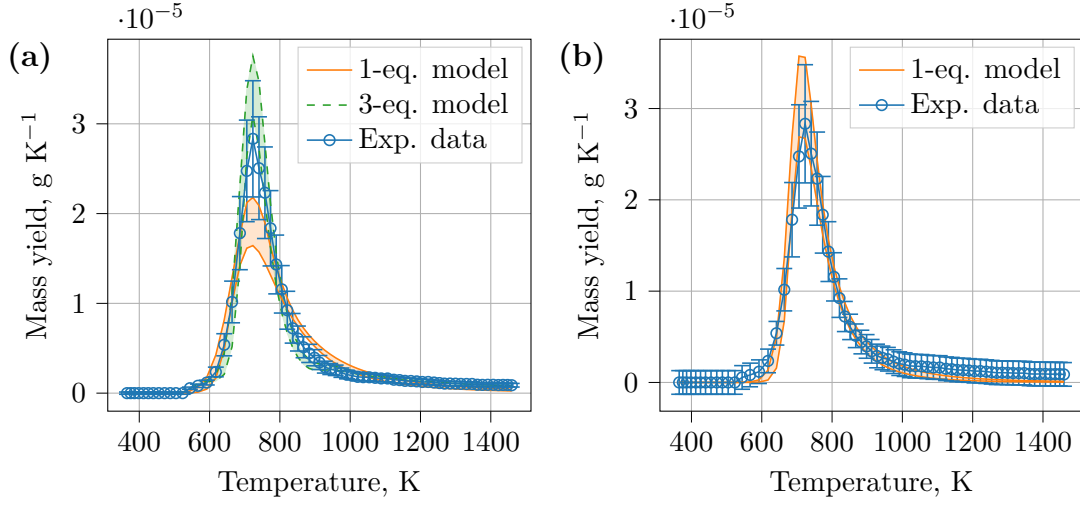


Figure 7.1: Influence of the weights on the calibration of the production rate of CH_4 for the one-equation model with uncertain parameters \mathcal{A} and \mathcal{E} . The shaded areas represent the region where 95% of the model evaluations lie and the error bars represent the values of the inverse of the weights. If no relaxation of the weights is included (a), we need to increase the number of equations in the model to better capture the maximum production peak. With the weights relaxed (b), a one-equation model is enough.

weights are not defined properly. We thus apply the following regularization:

$$s_j = \begin{cases} s_{\min} & \text{if } s_j^{(e)} = 0, \\ s_j^{(e)} & \text{else,} \end{cases} \quad \text{with } 1 \leq j \leq n_t, \quad (7.13)$$

where $s_{\min} = \min_j s_j^{(e)}$ for $1 \leq j \leq n_t$ with strictly positive values of $s_j^{(e)}$.

Further modifications of the weights

We assume first that the production rate for CH_4 can be reasonably reproduced using a one-equation model that simulates the single production peak. The initial parameter guess is obtained by manually tuning the parameters and visual inspection of the simulated production rate. The two other parameters of the pyrolysis model, namely γ and r , are kept fixed to 0.0011 and 4.96, respectively. From the experiments, we have $\Delta t/t_f = n_t^{-1}$, and we set $\vartheta = n_t$. The calibration is then performed using 10^4 iterations with the ISDE algorithm.

The evaluations of the one-equation model at the MCMC sample values are shown in Fig. 7.1(a). At every temperature, the 2.5–97.5 percentile range (shaded area in Fig. 7.1(a)) is computed using all the model evaluations. It can be seen that the resulting range misses the maximum of production rate and favors a good fitting of the extremities of the curve, which is to us unwanted. This is due to a local maximum in the parameter posterior density that has a higher density than the one at the parameter values that provide the good fitting around the maximum production peak. The presence of this spurious local maximum in the posterior appears because of the high values of the weights at the extremities of the curves (even if those values were already regularized in Eq. (7.13), the values of the s_j are still low). Any discrepancy of the simulated curve compared to the experimental data at these locations will have a significant influence on the posterior compared to the other data close to the peak, to which lower values of weights are assigned.

Table 7.1: Nominal parameter values for the three-equation pyrolysis model for the production of CH_4 . The two uncertain parameters are the ones for which the values are inferred.

R	F	$\log(\mathcal{A})$ (s^{-1})	\mathcal{E} (kJ mol^{-1})	n
1	0.000369	27.83	164.55	14.37
2	0.003789	Uncertain	Uncertain	3.34
3	0.003593	27.83	284.722	67.37

In order to suppress this undesirable outcome, a first possibility is to relax the weights at the extremities of the curves, that is, at low and high temperatures. Let $s_{\max} = \max_j s_j^{(e)}$ for $1 \leq j \leq n_t$. In order to reduce the influence of the small values of the weights that are away from the peak of maximum production, we assign

$$s_j = \begin{cases} cs_{\max} & \text{if } j \text{ s.t. } s_j^{(e)} \leq cs_{\max}, \\ s_j^{(e)} & \text{else,} \end{cases} \quad (7.14)$$

where c is a constant that needs to be fixed. The constant c is a trade-off between providing flexibility to the model by increasing the model error where it is expected to be present, and constraining the model to fit these values. The experimental data with the relaxed weights with $c = 0.2$ and the resulting model evaluations for the 10^4 samples using the ISDE algorithm are shown in Fig. 7.1(b). The higher values of the $\sigma_j^{(e)}$ can be observed at both extremities of the curve. As expected, the maximum production peak is now captured by the one-equation model.

A second possibility is to improve the overall fitting by using a more flexible model for the production of CH_4 . From Fig. 7.1(a), a first increase in the production rate of CH_4 can be observed around 575 K and a residual production can be observed above 1000 K. Thus, two pyrolysis equations are added to the previous one-equation model to improve the fitting in these temperature ranges. A manual tuning and a few pre-simulations are performed in order to find the additional parameters introduced by the three-equation model, which are then fixed to the values that are summarized in Table 7.1. The two parameters \mathcal{A}_2 and \mathcal{E}_2 (we will drop the indexes in the following) control the kinetics of the second equation that is mainly responsible for the maximum production peak and are therefore considered to be uncertain. Finally, the ISDE algorithm is run for 10^4 iterations in order to infer \mathcal{A} and \mathcal{E} and the resulting evaluations of the three-equation model are shown in Fig. 7.1(a). It can be seen that the three-equation model fits reasonably the production of CH_4 over the whole temperature range, which allows us to use the initial weights from Fig. 7.1(a).

The choice of the model and the value of the weights have thus an important influence on the result of the calibration of the model parameters. Because we mainly want to capture the main production peak, we will both favor the choice of relaxing the weights and providing flexibility to the model, as it will be done in the example of Section 7.4.2, where the model will feature more parameters.

Influence of ϑ as a parameter for controlling the posterior uncertainty

Let now consider the three-equation pyrolysis model with the two uncertain parameters \mathcal{A} and \mathcal{E} . The result of the propagated samples obtained with $\vartheta = n_t$ was shown in Fig. 7.1(a) and the corresponding computation of the posterior distribution is shown in Fig. 7.2. The posterior from Eq. (7.6) is evaluated on a uniform 200×200 logarithmic grid in order to improve the accuracy in the region of high probability density in the original parameter set. As expected, we can observe a strong correlation between \mathcal{A} and \mathcal{E} emphasized by the large values they take and a non-Gaussian shape of the posterior. Because of the difficulty of visualizing the whole




Figure 7.2: Bivariate posterior densities for the parameters \mathcal{A} and \mathcal{E} in the original parameter space (left) and logarithmic scales (right). Fourteen density contours are shown with intensity represented on the colorbars on the right of each plot.

posterior in the original parameter set (Fig. 7.2(a)), the posterior is also shown on a log-log scale in Fig. 7.2(b).

The results for $\vartheta = 2n_t$, $\vartheta = n_t$ and $\vartheta = 0.8n_t$ are shown in Fig. 7.3. As it can be seen on Fig. 7.3(b), the decrease of ϑ results in a posterior that is wider and thus results in more uncertainty in the parameter values and reduces the kinetic compensation effect between \mathcal{A} and \mathcal{E} . The 95% ranges for the model evaluations are therefore larger as seen in Fig. 7.3(a) when ϑ decreases. The increase in the uncertainty when ϑ decreases is due to the fact that ϑ appears in the numerator of the exponential term of the posterior.

It is worth mentioning that decreasing too much the value of ϑ may trigger solutions in the posterior that were previously far from the current maximum density, and render the sampling of this part of the posterior possible. If unwanted, such solutions can be removed, for instance, by incorporating more restrictive bounds on the parameters defined in the prior.

Comparison of the sampling strategies

The two gradient-informed algorithms, namely ISDE and HMC, are compared with $\vartheta = n_t$ and again the three-equation pyrolysis model with the same two uncertain parameters \mathcal{A} and \mathcal{E} . For the sake of the comparison, we also perform the inference using the RWMH and with the change of variables. For the reproducibility of the path of the Markov chains, we control the initial state of the random generator by setting the seed equal to zero.

Regarding the sampling methods using the change of variables based on the structure of the pyrolysis equations (CVR), the challenge is to find the adequate scaling factors as their value will strongly influence the mixing of the resulting chain. The scaling factors appearing in Eqs. (7.10)–(7.11) are $\bar{t} = 1$ s and $\bar{\mathcal{E}} = 1.9406 \times 10^5$ J mol⁻¹ and the reference temperature is $\bar{T} = 725$ K. The reference time \bar{t} is used to have the logarithm of a dimensionless quantity, the reference activation energy is obtained from the estimation of the maximum likelihood, and the reference temperature is obtained as the argument of the maximum yield of CH₄. The free parameters in the HMC and ISDE also need to be set. The time step τ should neither be too small, in which case the sampling is inefficient, nor too large in which case the numerical scheme is not stable. They are chosen by a few trial-and-error simulations. For both HMC and ISDE algorithms, the time step is set to $\tau = 0.1$ s. For HMC, the leapfrog step $L = 20$, and the damping parameter in the ISDE is set to $f_0 = 4$. Similarly, the covariance matrices for HMC and RWMH with

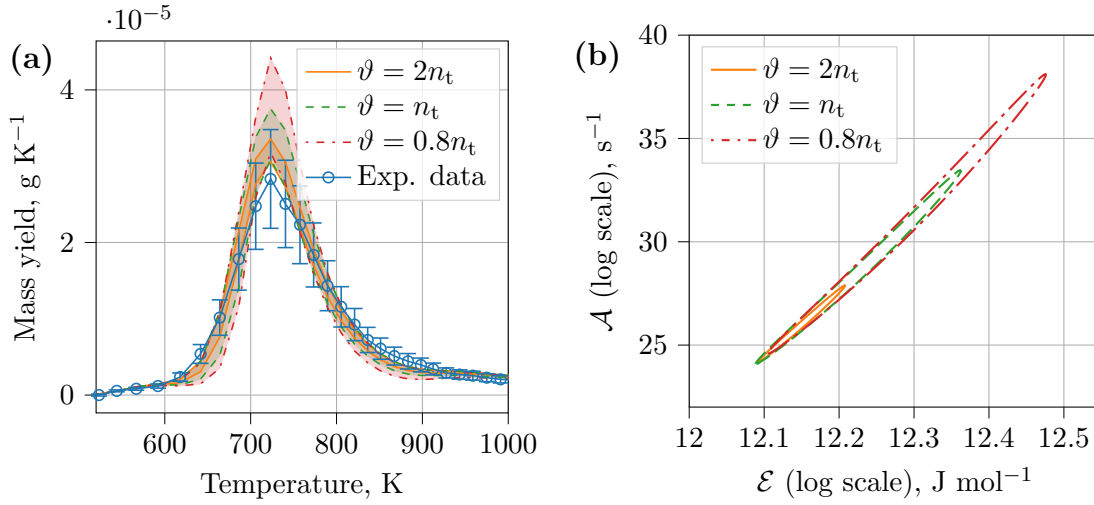


Figure 7.3: Influence of the parameter ϑ on the result of the inference for the three-equation model with two uncertain parameters \mathcal{A} and \mathcal{E} . The production rate of CH_4 (a) and the 95% posterior density probability (b) for three different values of ϑ are shown. In (a), the shaded areas represent regions where 95% of the model evaluations lie and the error bars represent the values of the inverse of the weights. In (b), the inner part of the curves represents a region of parameter values within which 95% of the probability density lies.

the CVR are set by a few trial-and-error steps. For HMC with CVR, the covariance matrix is the identity matrix, and for the RWMH it is a diagonal matrix with diagonal elements equal to 10^{-3} . The sampling algorithms are run for 1.1×10^5 iterations.

In order to assess the efficiency of the different algorithms, the ergodic mean and variance for the marginal posterior distributions as a function of the number of iterations n are computed for \mathcal{A} and \mathcal{E} as

$$\mu_i(n) = \frac{1}{n} \sum_{k=1}^n m_i^k, \quad (7.15)$$

$$\text{var}_i(n) = \frac{1}{n} \sum_{k=1}^n (m_i^k - \mu_i(n))^2. \quad (7.16)$$

The translated mean and normalized variance can be computed as $\tilde{\mu}_i(n) = \mu_i(n) - \bar{\mu}_i$ and $\tilde{\text{var}}_i(n) = \text{var}_i(n)/\bar{\text{var}}_i$ such that their values tend to zero and one, respectively. Good approximations of the reference mean $\bar{\mu}_i$ and reference variance $\bar{\text{var}}_i$ can be obtained from the direct evaluation of the posterior distribution, which in the case shown in Fig. 7.2 are equal to (for \mathcal{A} and \mathcal{E} in logarithmic values) $\bar{\mu}_{\mathcal{A}} = 28.53$, $\bar{\text{var}}_{\mathcal{A}} = 12.22$, $\bar{\mu}_{\mathcal{E}} = 3.74$ and $\bar{\text{var}}_{\mathcal{E}} = 0.0032$.

The comparisons of $\tilde{\mu}_i$ and $\tilde{\text{var}}_i$ are shown in Fig. 7.4 for $\log(\mathcal{A})$. The graphs for the parameter \mathcal{E} are almost identical due to the correlation between the two parameters and are shown in Appendix C.3. On the x-axis, we represented the number of times the model is evaluated instead of the number of MCMC iterations n . Both are the same for RWMH and ISDE, but the HMC algorithm performs L leapfrog steps before selecting the new sample, and thus L times more model evaluations are performed. The number of MCMC iterations for the HMC algorithm is therefore equal to 5×10^3 and the estimated mean and variance is constant during L model evaluations following each new MCMC iteration. We discarded the 10^4 samples from the burn-in phase (500 for HMC-type algorithms), which corresponds to the adaptive phase in the ISDE-DA and HMC-DA.

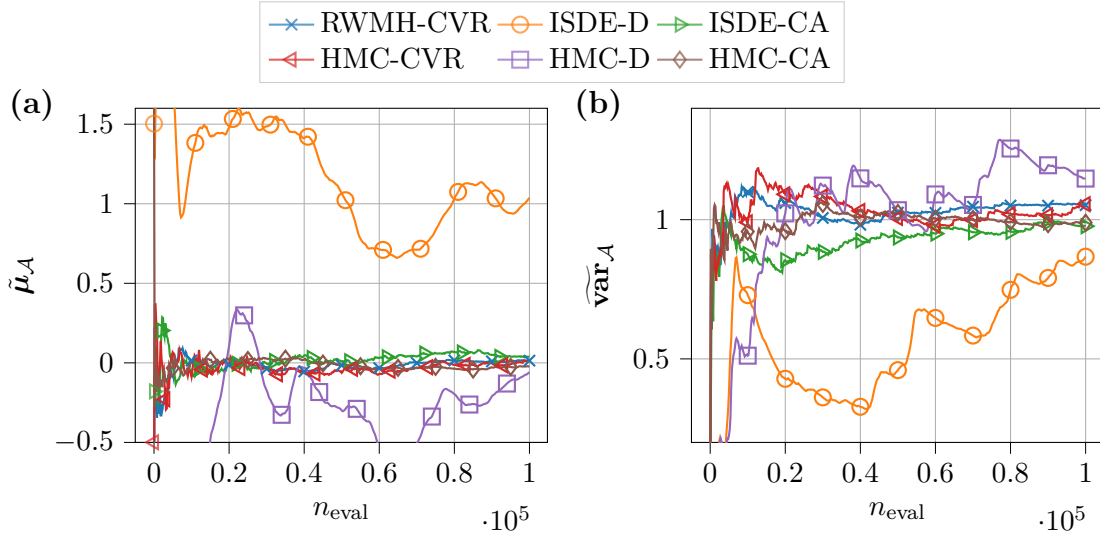


Figure 7.4: Convergence of the mean (a) and variance (b) for \mathcal{A} (in log) as a function of the number of function evaluations (after the burn-in phase). Random-Walk Metropolis–Hastings (RWMH); RWMH with the change of variables based on the structure of the pyrolysis equations (RWMH-CVR); ISDE with the covariance matrix estimated from the diagonal of the inverse of the Hessian and change of variables based on the support of the posterior (ISDE-D); ISDE-D with the covariance matrix resulting from the adaptation process (ISDE-CA); HMC with the change of variables based on the structure of the pyrolysis equations (HMC-CVR); HMC with the covariance matrix estimated from the diagonal of the inverse of the Hessian and change of variables based on the support of the posterior (HMC-D); HMC-D with the covariance matrix resulting from the adaptation process (HMC-CA).

First, from Fig. 7.4, it can be noted that the change of variables from Eqs. (7.10)–(7.11), the one that is well-suited for the pyrolysis model, applied to the random-walk Metropolis–Hastings (RWMH-CVR) shows a good convergence for the mean and the variance. These good convergence properties were not observed for the classical RWMH without change of variables, even with a careful tuning of the covariance proposal. The change of variables is also applied to the HMC algorithm (HMC-CVR) which also shows very good convergence properties.

When the change of variables based on the support of the posterior is applied using the diagonal covariance matrix estimation, both ISDE (ISDE-D) and HMC (HMC-D) are far from convergence. However, they both show good performances when the covariance matrix estimation results from the adaptation phase (ISDE-CA and HMC-CA), which is performed during the initial burn-in phase. During this initial phase, the algorithm is run for 10^3 iterations (50 for HMC) before the covariance matrix starts to update using the recurrence formulas from Eqs. (5.41)–(5.43). The covariance is then updated every 100 (5 for HMC) iterations until the 10^4 -th (500-th for HMC) iteration, which corresponds to the end of the burn-in phase. This shows that the adaptive HMC and ISDE are robust even if the estimation for the initial covariance matrix was representative of the local scales and not the correlation. Adaptation may also succeed with a less informative covariance matrix, e.g. the identity matrix, but a larger adaptive phase would be required.

The 2×10^4 first sample values from the MCMC chain for \mathcal{E} using the ISDE algorithm are shown in Fig. 7.5, where the effect of the adaptation phase on the mixing of the chain can be observed. Similar graphs are obtained for \mathcal{A} and using HMC. The first part on the left are the samples obtained during the 10^3 first iterations of the ISDE using the fixed diagonal covariance




Figure 7.5: Iterations of the Markov chain for ε (shown every 10 iterations for clarity) obtained using the ISDE algorithm. The first (green) dotted vertical line represents the iteration at which the adaptation starts. The second (red) dotted line represents the last iteration of the adaptation. Above 10^4 iterations, the covariance is fixed to the estimated value and the MCMC chain continues.

matrix, obtained from the diagonal of the inverse of the Hessian computed at the initial parameter values. From iterations 10^3 to 10^4 , the covariance is estimated from the recurrence formula and updated in the algorithm every 10^2 iterations. The effect of the adaptation is clearly observed when compared to the first part of the chain. After 10^4 , the covariance matrix is kept fixed and the chain is run for 10^4 more iterations.

In Fig. 7.6, we compare the effect of the number of iterations during which the adaptation is performed on the convergence of the mean and variance. The total number of iterations performed is 1.1×10^5 (5.5×10^3 for HMC) for all cases. Figure 7.6 shows the results when the adaptation is stopped after the first 5×10^3 (2.5×10^2 for HMC), 1×10^4 (5×10^2) and 2×10^4 (1×10^3) iterations. Once again, the mean and variance are estimated after discarding the samples from these adaptation phases. Figure 7.6 also shows the result when the adaptation is not stopped and is run during all the 1.1×10^5 (5.5×10^3 for HMC) iterations, in which case we use all the sample history to compute the mean and variance.

It can be observed that the results converge toward zero for the mean and one for the variance in all cases. For HMC, no significant difference can be noticed, except a slight difference when the adaptation is performed until the end ($n_{\text{end}} = 5.5 \times 10^3$). For ISDE, the best result is for $n_{\text{end}} = 10^4$, which motivates our choice in the previous section for stopping the adaptation at this value. The result when the adaptation is performed during all the iterations shows that the algorithm may also work with a continuous adaptation of the covariance matrix, although we do not have the theoretical support.

Similarly, the effect of the period at which the covariance matrix is updated in the algorithm (Δn_{adapt}) on the convergence rate of the mean and variance is also investigated. When the covariance matrix is effectively updated in the algorithm, it requires the computation of its inverse for computing L_M in Eq. (5.54) for the ISDE algorithm, or for sampling \mathbf{p} from the multivariate normal in HMC. This may take some time when the dimensionality increases if the covariance matrix is adapted at every iteration. Conversely, a long adaptation period Δn_{adapt} may require a longer adaptive phase for having the benefit of adapting the covariance matrix from samples. Results were obtained for $\Delta n_{\text{adapt}} = 10$, $\Delta n_{\text{adapt}} = 100$ and $\Delta n_{\text{adapt}} = 1000$ and showed that $\Delta n_{\text{adapt}} = 100$ was a good compromise between the accuracy of the estimation of the mean and variance and the computational time. Having a short adaptation period did not




Figure 7.6: Convergence of the mean (a) and variance (b) for $\log(\mathcal{A})$ as a function of the number of function evaluations. The number of iterations in the adaptation phase (n_{end}) is compared for the HMC and ISDE algorithms.

show a significant improvement in the convergence rate, and a long adaptation period did not provide bad convergence results either.

Parametric study of f_0

Finally, we provide a numerical analysis for the influence of the damping factor f_0 on the convergence of the mean and variance for the parameter \mathcal{A} . The parameter f_0 introduces dissipation in the stochastic equation and will influence the convergence speed for the estimation of the mean and variance. We consider again the case with parameters from Table 7.1 with $\vartheta = n_t$, the covariance matrix results from the adaptation of the first 10^4 iterations and is fixed for the 10^5 iterations.

Results for the estimation of the mean and variance of \mathcal{A} are shown in Fig. 7.7. For low values of f_0 ($f_0 = 0.01$ and $f_0 = 0.1$) we observe an erratic convergence, in particular for the variance. In this case, because the stochastic term is proportional to the square root of f_0 , it is predominant compared to the damping term. The system is closer to an Hamiltonian system and the algorithm generates highly correlated samples with rapid local variations that can be observed on the curves. For values of the order of unity (from $f_0 = 0.5$ to $f_0 = 10$) the convergence is faster, with the closest results to zero mean and unit variance achieved for $f_0 = 0.5$ and $f_0 = 1.5$. For larger values of f_0 ($f_0 = 50$), the damping factor is predominant and the convergence is again slower, but the local oscillations are not observed. Therefore, it seems that values of f_0 ranging from 0.5 to 10 can all achieve a good result for the convergence of the mean and variance, while values outside of this range should be avoided.

7.4.2 Bayesian parameter inference for PICA devolatilization at a high heating rate

In this last section, we consider the whole experimental data set of Bessire and Minton [2017] at the heating rate $\beta = 360 \text{ K min}^{-1}$ for the pyrolysis of PICA (Fig. 3.3) and we perform

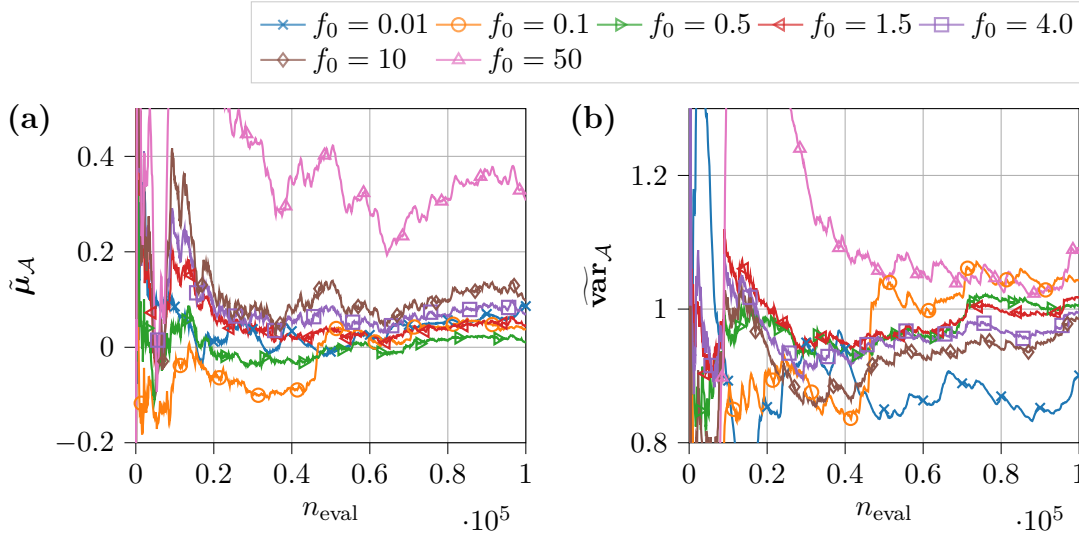


Figure 7.7: Convergence of the mean (a) and variance (b) for $\log(\mathcal{A})$ using the ISDE algorithm as a function of the number of function evaluations with different values for the damping factor f_0 .

the Bayesian calibration of the parameters. The data for the 1-propanol and 2-propanol are discarded, as they are scarce and might be subjected to high experimental errors, and therefore $n_{\text{obs}} = 12$.





Iterative procedure for the construction of the pyrolysis model

There are two main issues that need to be solved: the first issue is that the structure of the pyrolysis model is unknown; the second issue is that we do not have an initial estimate of the parameters for a pyrolysis model once its structure is given.

In [Torres-Herrador et al. \[2019b\]](#), a six-equation ($N_p = 6$) pyrolysis model was identified by visual inspection of the number of maximum species production peaks. A model fitting procedure was then performed using a multi-objective genetic algorithm to infer the values of the 35 parameters. However, by performing the Bayesian parameter inference from the model structure identified in [Torres-Herrador et al. \[2019b\]](#), similar observations as the ones from [Fig. 7.1\(a\)](#) were obtained, that is, some production peaks were not appropriately simulated due to the influence of the weights. It does not mean that the model is wrong, but that some misfits in the model that are not accounted for in the weights degrade the overall fitting of the production peaks.

Because of the complexity of the pyrolysis model structure, identifying which values of the weights are degrading the goodness of fit is not obvious. This leads us to the need of rebuilding the pyrolysis model structure in an iterative procedure, in which models of increasing complexity are successively inferred. We start by inferring a one-equation model using a single experimental curve, as it was done for CH_4 in [Section 7.4.1](#). Calibrating this model is simpler as it is defined by fewer parameters and the initial guess can be easily obtained, for instance, by running the sampling algorithms for a few iterations to converge toward the region of higher probability. Once we have a satisfactory solution for the parameters \mathcal{A}_1 , \mathcal{E}_1 , m_1 and γ_{11} , we add one experimental curve in the set of calibration data (we add one species j). If a new reaction maximum is observed in the additional curve, we increase the number of equations by one and we need to calibrate the new set of parameters \mathcal{A}_2 , \mathcal{E}_2 , m_2 and γ_{2j} together with the previous set of parameters for which we have an initial guess. If a reaction already exists at the new peak, then we just need

to calibrate the new mass fractions γ_{2j} . This procedure is repeated until the model includes all the species from the data set. As already mentioned in Section 7.4.1, issues may arise when the inverses of the weights approach zero where the production rate is zero. Therefore, we apply once again the regularization of the weights from Eq. (7.13) and the relaxation of the weights at the extremities from Eq. (7.14) with $c = 0.6$. The following six-equation pyrolysis model is found:



Posterior analysis

Using the model from Eqs. (7.17)–(7.22), the number of parameters to infer is equal to 48. The prior is the product distribution from the marginal priors. The marginal priors for pre-exponential factors \mathcal{A}_i are semi-infinite uniform priors. The activation energies \mathcal{E}_i are bounded in given ranges in order to separate the different peaks of production. Evidence from some of our previous simulations shows that the solution can converge slowly or can sample unwanted regions in the posterior if the activation energies are not limited within a given range. The values of the reaction orders are also limited in the range $]1, 10]$ and all mass fractions range between 0 and 1.

The ISDE algorithm is run for 10^6 iterations with $\tau = 0.002$ s and $f_0 = 4.0$. The initial covariance matrix is estimated using the diagonal matrix of the inverse of the Hessian computed at the initial guess value. Then, the covariance is updated every 100 iterations from $n_0 = 10^4$ to $n_f = 10^5$. Due to the higher dimensionality, the algorithm needs more iterations to learn the local scales and correlations and the adaptation phase is therefore longer. The solution to the pyrolysis model and its derivatives with respect to the parameters are all computed analytically (Coheur et al. [2021b]).

The results for $\vartheta = n_t$ are shown in Figs. 7.8–7.10. In Fig. 7.8, the Markov chains are provided for the kinetic parameters of the reaction R_1 and the corresponding mass fraction of CH_4 . The 10^5 first iterations from the adaptation phase are discarded. The other chains representing the samples drawn from the 48-dimensional posterior feature similar behavior of mixing but are not shown here.

In Fig. 7.9 the bivariate and marginal posteriors for the reaction R_1 are shown. In the bivariate posterior for \mathcal{A}_1 (in log) and \mathcal{E}_1 , the correlation between the two parameters can be observed. This correlation is reduced compared to the previous section because of the restricted range on \mathcal{E} , which is necessary to limit the potential interactions between the different reactions. The region where \mathcal{A} varies is therefore reduced. Regarding the marginal parameter distributions, they feature a single mode distribution.

The production rate curves at the sample values are shown in Fig. 7.10. We represent the regions where 95 % of the model evaluations lie with the mean value in between. From this figure, we can directly assess the goodness of fit of the model, and we can conclude that the model structure chosen provides an overall fit that is good. One can notice that this model

underestimates a bit the production of trimethyl phenol, xylene and toluene and overestimates slightly the maximum of CO_2 . In general, the model represents well the observations within the range [500, 1000] K. Above 1000 K, the production almost reaches a plateau for CO , CO_2 , H_2O , and H_2 , which is difficult to account for using the multi-component approach from Eq. (3.9).

Discussion of the choice of the value of ϑ

Selecting the value of ϑ based on the overall goodness of fit not obvious as there is significant model inadequacy at low and high temperatures. However, we can compare the fitting by visual inspection based on the results in the range from [500, 1000] K, which is fairly good, and we can compare the overlap of the computed intervals with the experimental intervals. From that point of view, the computed intervals from Fig. 7.10 might still be underestimating the uncertainty ranges and the posterior distribution be over-confident.

In order to increase the uncertainty in the simulated curves and be more in agreement with the experimental ranges, the inference is also performed with $\vartheta = 0.7n_t$. This value was obtained after a few trials, and the resulting simulated curves are shown in Fig. 7.11. The increase in the uncertainty range can be clearly noticed, which is due to the posterior distribution that is less concentrated, showing that the parameter ϑ allows to control the overall uncertainty range on the predictions. The maximum a posteriori values for the whole set of parameters are provided in Table C.1 in the Supplementary Material.

Yet, some observations are still outside of the computed uncertainty ranges, particularly below 500 K and above 1000 K. This issue highlights limitations of the approach of seeking to incorporate the effect of modeling errors as an uncertainty in the model parameters. If unmodeled or poorly represented physical behavior results in model predictions that are not sufficiently close to the experimental data, then such model inadequacy cannot be compensated by introducing uncertainty in the model parameters.

7.5 Application: parameter inference of competitive pyrolysis models

In this second application section, we consider a competitive pyrolysis model for the PICA TPM that allows us to simulate the decomposition of the material at two distinct heating rates, namely 10 K min^{-1} and 366 K min^{-1} using the results from two experimental data set taken from Wong et al. [2016] and Bessire and Minton [2017], respectively (see Section 3.2). Competitive pyrolysis models were described in Section 3.4 and their introduction was motivated by the observation that the mass loss curve and the pyrolysis production peak shift towards lower temperatures as the heating rate increases, which cannot be accounted for by using multi-component models.

We perform the Bayesian parameter inference and the results are presented, and the interaction between the parameters of the different branches of the competitive mechanism are analyzed. The extrapolation capability of the calibrated model is assessed by propagating and comparing with experimental results at other heating rates. Results from this section were published in the co-authored manuscript presented in Torres-Herrador et al. [2019a].

7.5.1 Competitive reaction scheme for PICA

The observation that at high heating rates the decomposition curve shifts towards lower temperatures suggests to isolate into two different pathways the contribution of the low heating and the high heating rates. A kinetic model for PICA consistent with this observation and the experimental data at 10 and 366 K min^{-1} is shown in Fig. 7.12 (Torres-Herrador et al. [2019a]). The superscripts (g) and (s) are used to clarify which components are gaseous or solids, respectively. In addition, the superscript “*” highlights intermediate solids that are produced without




Figure 7.8: MCMC chains for selected parameters after a burn-in of the firsts 10^5 samples.

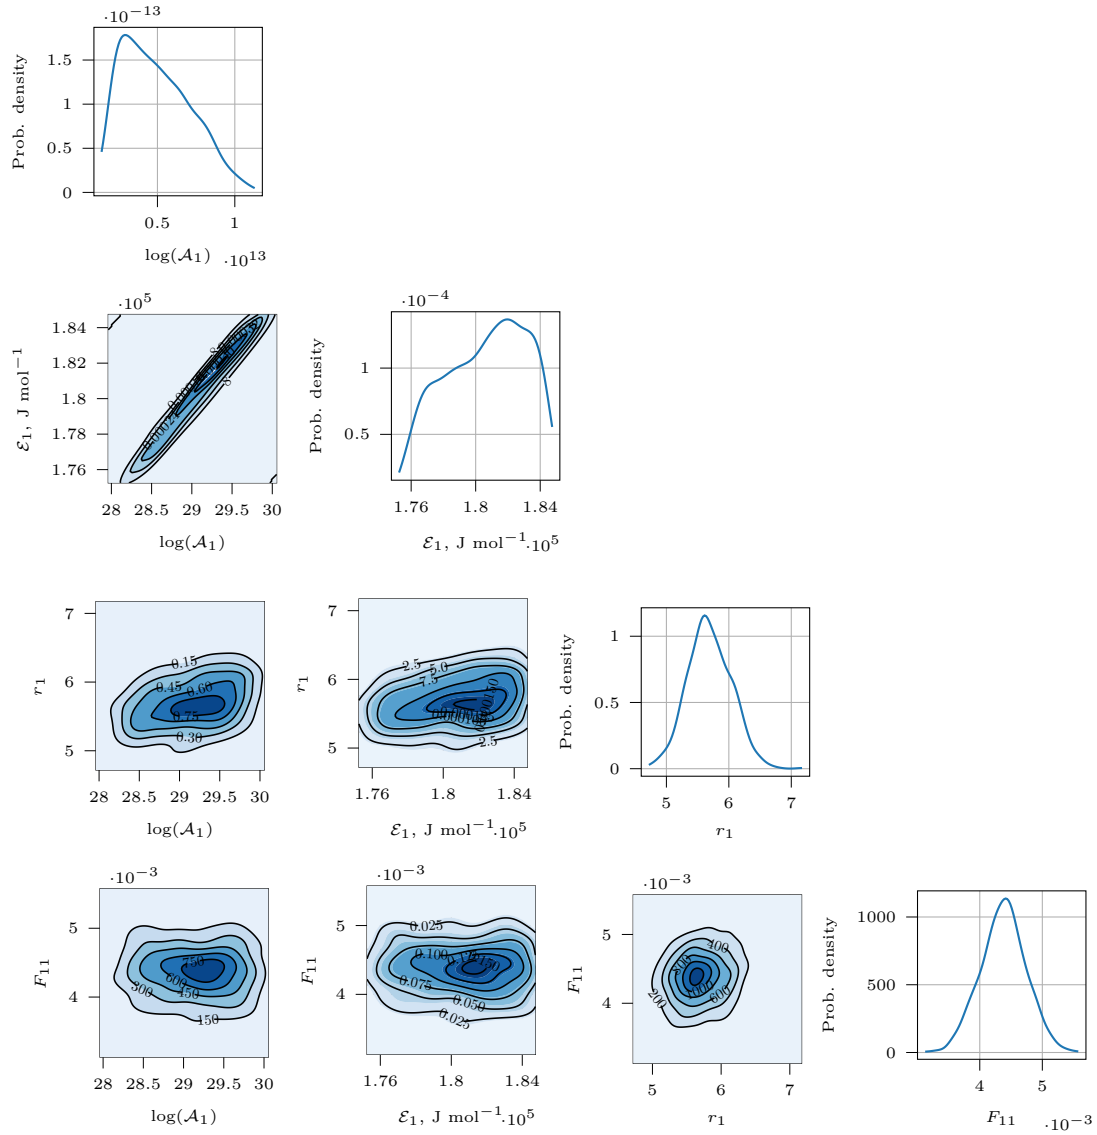


Figure 7.9: Marginal and bivariate posterior distributions for selected parameters. \mathcal{A}_1 is represented on a log-scale (unit of \mathcal{A}_1 is in s^{-1}). The probability densities have the inverse units of their variable(s).

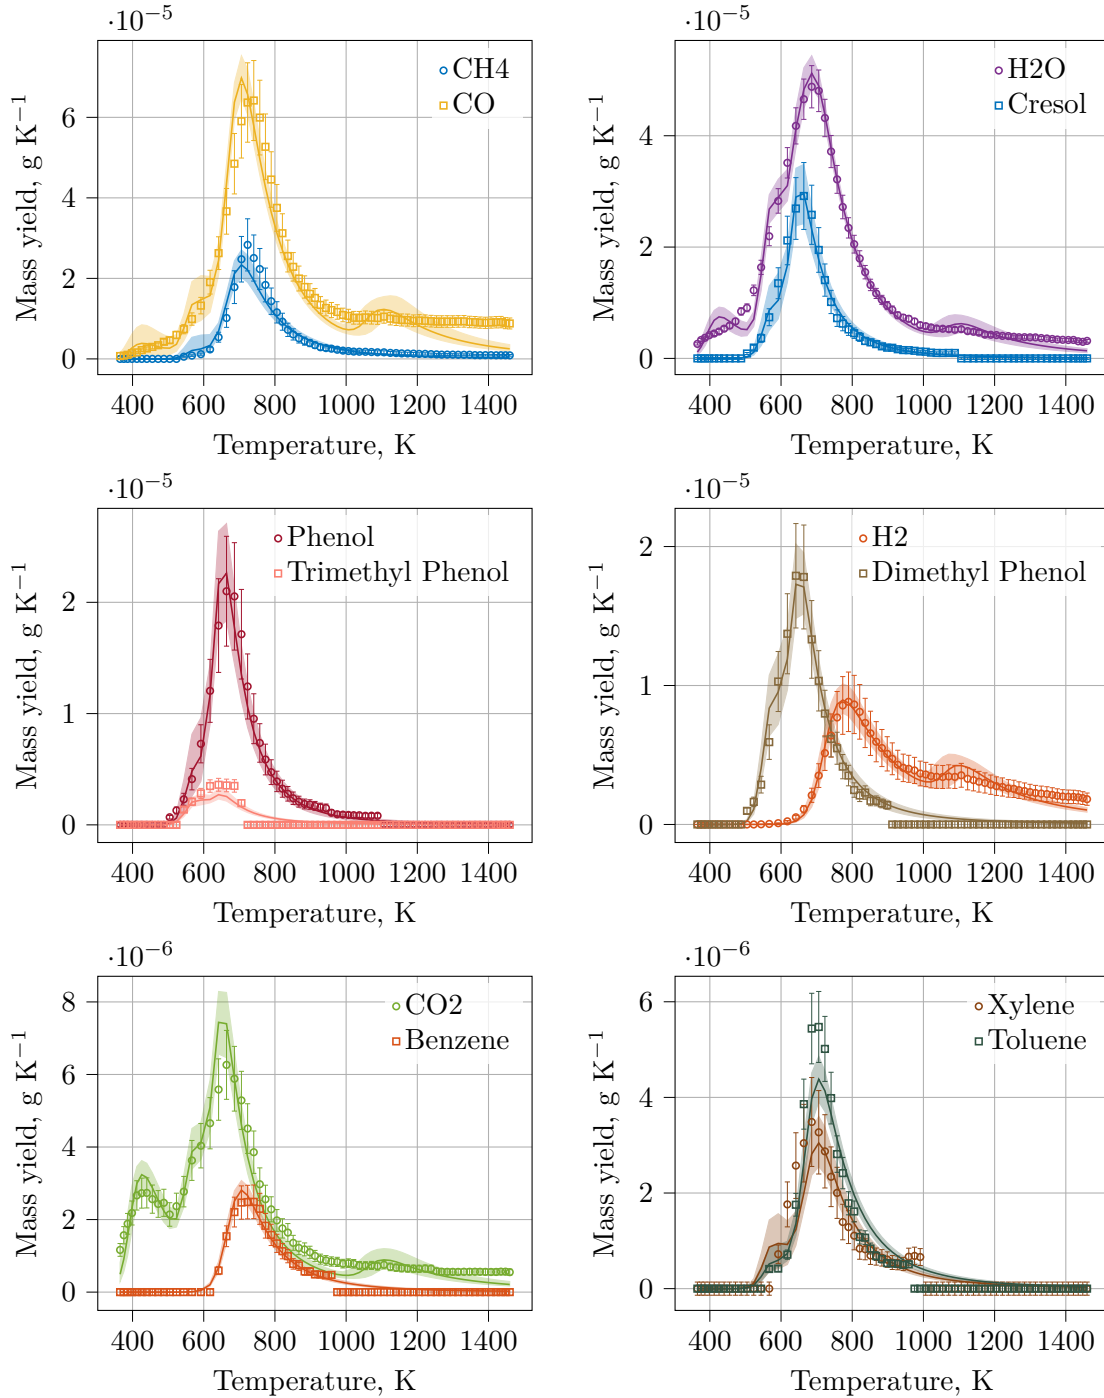


Figure 7.10: Mass yields as a function of temperature for each species obtained after the propagation of the MCMC samples with $\vartheta = n_t$. The symbols with the error bars represent the experimental results from Bessire and Minton [2017]. The continuous lines are the results of the propagation for the maximum likelihood estimate and the shaded areas represent all the possible outcomes of the propagation using every samples from the posterior distribution.

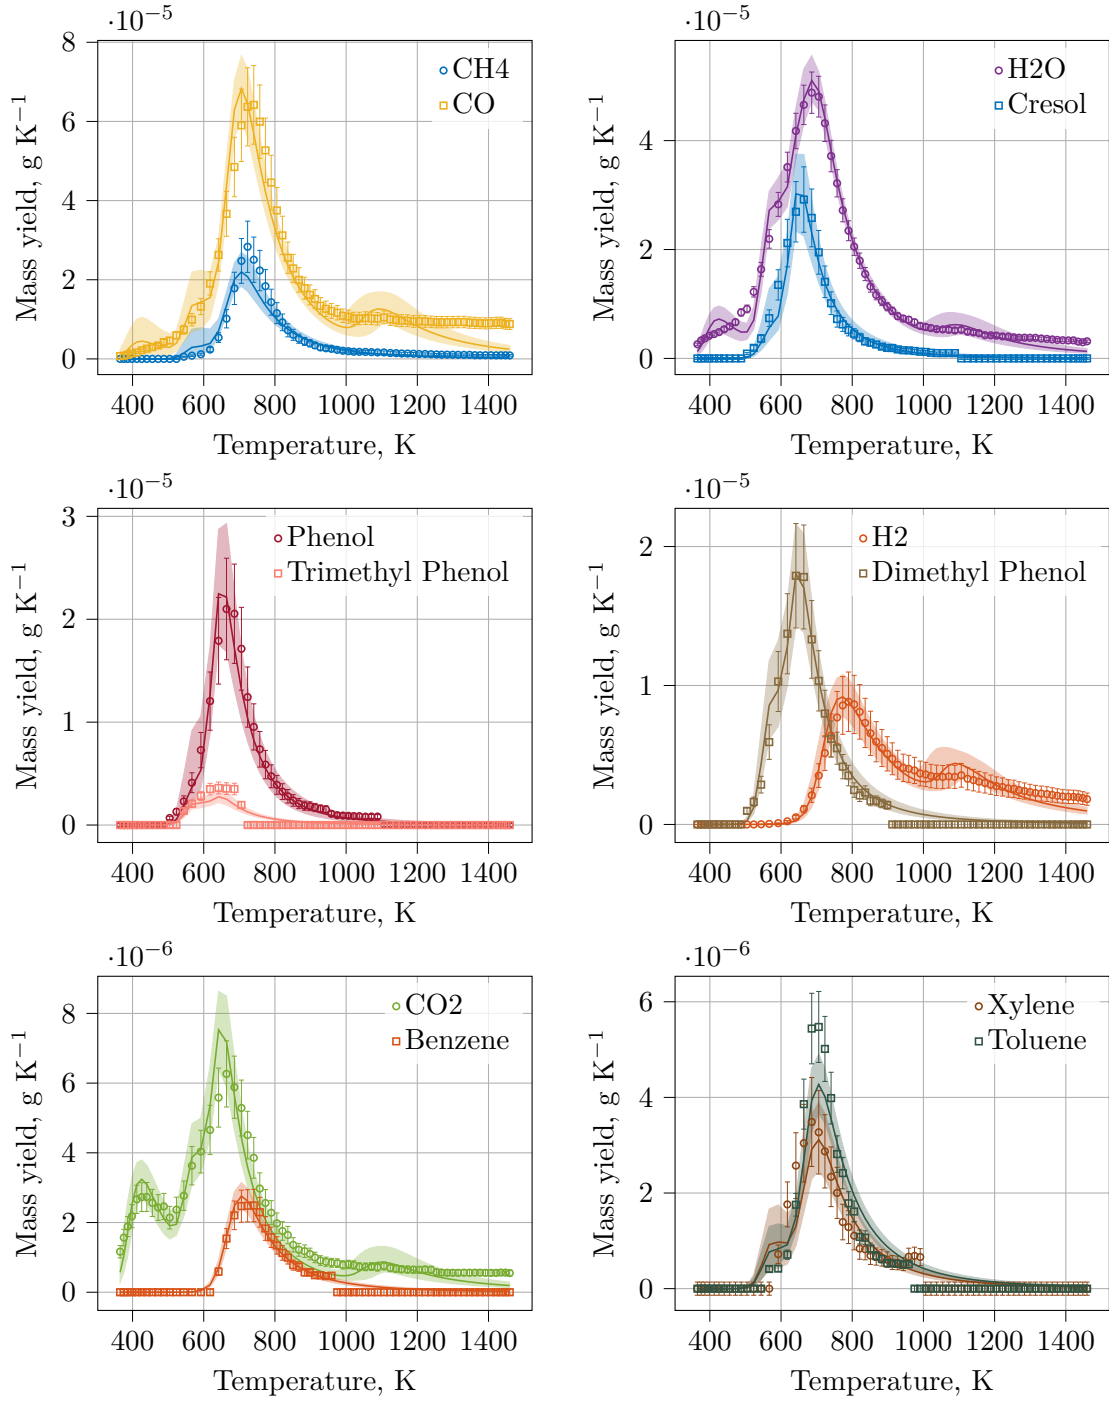


Figure 7.11: Mass yields as a function of temperature for each species obtained after the propagation of the MCMC samples with $\vartheta = 0.7n_t$. The symbols with the error bars represent the experimental results from Bessire and Minton [2017]. The continuous lines are the results of the propagation for the maximum likelihood estimate and the shaded areas represent all the possible outcomes of the propagation using every samples from the posterior distribution.

release of pyrolysis gases in a depolymerization process (Miller and Bellan [1997]). Thus, the initial reactant $\rho_1^{(s)}$ decomposes into two different branches: a slow process for low heating rates with reaction rate $k_{1,1}$ and a fast one for high heating rates with reaction rate $k_{1,2}$. The slow branch progressively depletes the reactant starting at lower temperatures ($\mathcal{E}_{1,1} < \mathcal{E}_{1,2}$) and requires time to be completed. Therefore, at low heating rates, the decomposition is dominated by this slow process that produces the intermediate solid $\rho_2^{(s)*}$. At high heating rates, the fast process becomes active and dominant. This process decomposes most of $\rho_1^{(s)}$, producing the intermediate solid $\rho_3^{(s)*}$, while impeding the production of $\rho_2^{(s)*}$. Each intermediate solid ($\rho_2^{(s)*}$ and $\rho_3^{(s)*}$) further reacts (at rates $k_{2,1}$ and $k_{3,1}$) and decomposes into a charred solid ($\rho_4^{(s)}$ and $\rho_6^{(s)}$), and releases pyrolysis gases ($\rho_5^{(g)}$ and $\rho_7^{(g)}$).




Figure 7.12: Proposed competitive pyrolysis model for the thermal degradation of PICA at 10 and 366 K min⁻¹.

As described in Section 3.4.2 and from Eq. (3.44), the reaction scheme depicted in Fig. 7.12 can be written as a system of ODEs as follows:

$$\frac{d}{dt} \begin{pmatrix} \rho_1^{(s)} \\ \rho_2^{(s)*} \\ \rho_3^{(s)*} \\ \rho_4^{(s)} \\ \rho_5^{(g)} \\ \rho_6^{(s)} \\ \rho_7^{(g)} \end{pmatrix} = \begin{bmatrix} -(k_{1,1} + k_{1,2}) & 0 & 0 & 0 & \dots & 0 \\ k_{1,1} & -k_{2,1} & 0 & 0 & \dots & 0 \\ k_{1,2} & 0 & -k_{3,1} & 0 & \dots & 0 \\ 0 & \gamma_{2,1,4} k_{2,1} & 0 & 0 & \dots & 0 \\ 0 & \gamma_{2,1,5} k_{2,1} & 0 & 0 & \dots & 0 \\ 0 & 0 & \gamma_{3,1,6} k_{3,1} & 0 & \dots & 0 \\ 0 & 0 & \gamma_{3,1,7} k_{3,1} & 0 & \dots & 0 \end{bmatrix} \begin{pmatrix} \rho_1^{(s)} \\ \rho_2^{(s)*} \\ \rho_3^{(s)*} \\ \rho_4^{(s)} \\ \rho_5^{(g)} \\ \rho_6^{(s)} \\ \rho_7^{(g)} \end{pmatrix}. \quad (7.23)$$

This system of ODEs can be simplified by removing the dependent variables $\rho_4^{(s)}$, $\rho_5^{(g)}$, $\rho_6^{(s)}$ and $\rho_7^{(g)}$ such that L is a (3×3) matrix of full rank. In addition, the mass conservation constraint $\gamma_{i,j,l+1} = (1 - \gamma_{i,j,l})$ allows us to remove the two unknown mass coefficients for the pyrolysis gases. In total, considering the four reaction rates $k_{i,j}$, each of them having two unknown parameters, the four mass coefficients and the two consistency relations, the total number of parameters to be calibrated is equal to 10.

More complex models that still provide a good fit to the experimental data were investigated. However, increasing the model complexity by adding more reactions did not improve significantly the cost function minimization. Therefore, it was decided to favor simplicity and computational cost over a slight improvement on the cost function (MacKay [2003]). Further iterations on this first scheme are foreseen when more experimental data become available.

The pyrolysis model parameters were first calibrated in Torres-Herrador et al. [2019a] using a shuffled complex evolutionary algorithm, which is a global optimization algorithm similar to a genetic optimization with a local search (Duan et al. [1993]). The graphical results of the calibrated model are shown in Fig. 7.13 and the nominal value of the parameters resulting from the optimization process are provided in Table 7.2. It can be seen from Fig. 7.13 that the calibrated model can reproduce the experimental data at both heating rates, and the density as

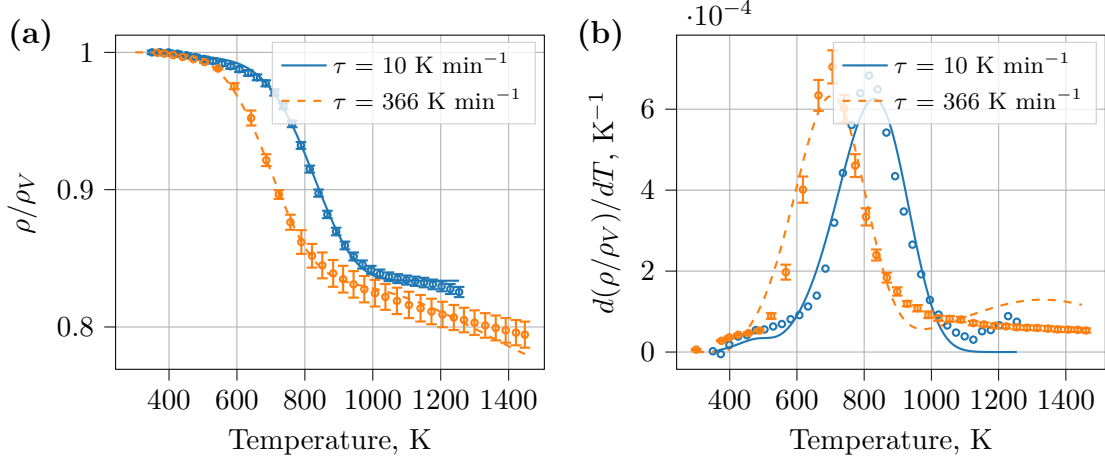


Figure 7.13: Comparison between the experiments and the calibrated model for the nominal parameter values for the two different heating rate scenarios. On the left (a) are the curves for the normalized sample density and on the right (b) the normalized production rate curves as a function of temperature. The experimental data from Wong et al. [2016] (10 K min^{-1}) and from Bessire and Minton [2017] (366 K min^{-1}) are represented by circular markers with their errorbars when available.

Parameter	Value	Parameter	Value
$\log_{10}(\mathcal{A}_{1,1})$	2.019	$\mathcal{E}_{1,1}$	32618.482
$\log_{10}(\mathcal{A}_{1,2})$	14.292	$\mathcal{E}_{1,2}$	143273.910
$\log_{10}(\mathcal{A}_{2,1})$	0.442	$\mathcal{E}_{2,1}$	51783.980
$\log_{10}(\mathcal{A}_{3,1})$	0.993	$\mathcal{E}_{3,1}$	31087.851
$\gamma_{2,1,5}$	0.163	$\gamma_{3,1,7}$	0.244

Table 7.2: Values for the calibrated parameters obtained from the deterministic optimization.

a function of time and its derivative are well reproduced by the proposed model. The shift in the production peak towards lower temperatures with increasing heating rate is also captured. This is in contrast with kinetic mechanisms based solely on multi-component parallel reactions that would not achieve to reproduce the experimental curves at both heating rates.

Consistently with the proposed model, the activation energy $\mathcal{E}_{1,1}$ is lower than $\mathcal{E}_{1,2}$ (Table 7.2), indicating that the conversion towards $\rho_2^{(s)*}$ starts earlier than the conversion towards $\rho_3^{(s)*}$. The pre-exponential factor $\mathcal{A}_{1,1}$ is lower than $\mathcal{A}_{1,2}$ which indicates that the production of $\rho_2^{(s)*}$ is slower than that of $\rho_3^{(s)*}$. In summary, the calibrated kinetic parameters are all coherent with the underlying hypotheses of the model.

7.5.2 Bayesian parameter inference for the competitive pyrolysis model

The likelihood function is built from Eq. (7.9) and the two observations are the mass loss and the production rate ($n_{\text{obs}} = 2$). The value of the free-parameter ϑ was not further studied here and we set $\vartheta = n_t$. As in Section 7.4 the weights were further relaxed at the extremities of the curves to avoid the calibrated parameter to be sensitive to modeling errors at these locations. The Markov chain is initiated using the results from the deterministic optimization from Table 7.2. The RWMH is used, the pre-exponential factors \mathcal{A}_i are taken in logarithmic values and the covariance matrix for the Gaussian proposal function was obtained after a few trial runs. Owing to the relatively short burn-in enabled by using the previous optimization result, we run the chain for 10^5 iterations.




Figure 7.14: (a) Mass loss curve at a heating rate of 10 K min⁻¹ and (b) production rate curve at heating rate 366 K min⁻¹ and their uncertainty ranges. The continuous lines are the computed mean at each temperature. The shaded areas represent the 2.5–97.5 % probability regions. The markers represent the experimental data with their respective error bars.

Parameter	Mean	Std	CV	Parameter	Mean	Std	CV
$\log_{10}(\mathcal{A}_{1,1})$	2.4768	0.3027	0.0.1222	$\mathcal{E}_{1,1}$	26811.37	893.61	0.0333
$\log_{10}(\mathcal{A}_{1,2})$	23.4935	1.1618	0.0494	$\mathcal{E}_{1,2}$	183938.42	2369.64	0.0129
$\log_{10}(\mathcal{A}_{2,1})$	0.2219	0.1238	0.5579	$\mathcal{E}_{2,1}$	48796.41	1723.16	0.03531
$\log_{10}(\mathcal{A}_{3,1})$	1.1969	0.0821	0.0686	$\mathcal{E}_{3,1}$	33566.43	976.07	0.0291
$\gamma_{2,1,5}$	0.1648	0.0038	0.0234	$\gamma_{3,1,7}$	0.3190	0.0703	0.2202

Table 7.3: Values for the mean, standard deviation (std) and coefficient of variation (CV) for the calibrated parameters obtained from the Bayesian inference.

Fig. 7.14 presents the simulated data obtained from the model evaluations at the sample values and compare them to the experimental data in order to assess the validity of the model. We represented the mean value (solid line) and the 2.5–97.5 % probability region (shaded areas). The model features an overall good agreement with the experimental data but some discrepancies as in Fig. 7.13 are again observed at high temperatures for 10 K min⁻¹ and at low temperatures for 366 K min⁻¹. These discrepancies are not captured by the uncertainties on the parameters and this suggests that further complexities could be captured by improving models that will help reducing the overall model error.

Second, in Table 7.3, the values obtained for the mean, standard deviation and coefficient of variation (CV) are summarized for the marginal distributions directly estimated from the Markov chains. The CV provides a measure of the relative dispersion of the parameter values with respect to the mean. The mean values may be different from the previous deterministic values used to initiate the Markov chains because of the differences in the weights used in the genetic optimization and the Bayesian inference. The mean value for $\gamma_{2,1,5}$ is very close to its deterministic value because this parameter is found to be the less correlated with all the other ones. The higher values of the CVs for the reaction branch (2,1) for both $\mathcal{A}_{2,1}$ and $\mathcal{E}_{2,1}$ suggest that these two parameters are more difficult to identify and are more uncertain, although this reaction produces the component $\rho_4^{(s)}$ at the two heating rates (see later in Fig. 7.16) and thus more information is provided from the experimental data. This higher uncertainty is attributed to the slower process of the reaction (2,1) characterized by the low value of $\mathcal{A}_{2,1}$.

A graphical representation of the full correlation matrix is provided in Fig. 7.15. The correlation matrix is obtained from the estimated covariance matrix normalized by the standard

deviations of the random parameters, such that diagonal terms of the correlation matrix are equal to one and off-diagonal terms range between -1 and 1. We note that the Arrhenius parameters \mathcal{A} and \mathcal{E} for the reactions (2,1) and (3,1) are found to be highly correlated with a value for the correlation coefficient that is close to 1. This results in a wide range of values for the pre-exponential factor \mathcal{A} and the activation energy \mathcal{E} that can give satisfactory fits to the data, the so-called kinetic compensation effect, that was also observed in the previous section. However, the pair of parameters \mathcal{A} and \mathcal{E} for reactions (1,1) and (1,2) that are related to the production of the intermediate components are found to be less correlated and less sensitive to the kinetic compensation effect. Besides the kinetic compensation effect for \mathcal{A} and \mathcal{E} , additional compensation mechanisms characterized by a high value for the correlation coefficient are observed: this is the case, for example, for the pair of parameters $\mathcal{A}_{1,1}$ with $\gamma_{3,1,7}$ or $\mathcal{A}_{2,1}$ with $\gamma_{3,1,7}$. There is only limited correlation with parameter $\gamma_{2,1,5}$ and results in the marginal distribution for which sampling has well converged (the closer to a Gaussian distribution). However, we note the high positive correlation value for the pair of parameters $\mathcal{A}_{1,1}$ with $\mathcal{A}_{1,2}$.

Finally, Fig. 7.16 represents the evolution of the densities as a function of the temperature with uncertainty intervals. For the nominal pathways, it is observed that at low heating rate (Fig. 7.16(a)), the initial solid is mostly converted into $\rho_2^{(s)*}$, while the opposite occurs at high heating rates (Fig. 7.16(b)). These intermediate steps further react at higher temperatures generating the char yield. It is observed in Fig. 7.16(b) that when both mechanisms are in competition, it is more uncertain to know what goes into which branch, leading to larger intervals and which also manifests itself in the high positive correlation between $\mathcal{A}_{1,1}$ with $\mathcal{A}_{1,2}$. This strong interaction between $\mathcal{A}_{1,1}$ with $\mathcal{A}_{1,2}$ is intrinsic to the competitive mechanisms: the uncertainties in these two parameters combine themselves and result in large uncertainty intervals in the densities when there is competition between the two reactions. On the other hand, for the pyrolysis at 10 K min^{-1} where the competition between the two reactions is less balanced, uncertainty ranges in the densities are tighter because the reaction is characterized mainly by the first branch of the mechanisms and is less influenced by the second one, as it can be corroborated on Fig. 7.16(a).

The results of this section show that a pyrolysis reaction scheme including a competitive mechanism can be used to represent the decomposition of the phenolic resin of PICA at two different heating rates. We were able to obtain samples from the posterior distribution using Markov chains and compare the model evaluations with the experimental data, thus taking into account the correlations between the input parameters. Not taking into account the correlation between the input parameters may lead to wrong results for the propagation and samples drawn directly from the Markov chain should be used for the uncertainty analysis. It was found that due to the nature of the competitive scheme and because a component can be both produced and consumed, there are additional kinetic compensation effects beyond the pair of parameters \mathcal{A} and \mathcal{E} , thus making the inference process even more complex.

7.5.3 Extrapolation to other heating rates

The proposed competitive model is compared to additional experimental results in order to assess the extrapolation capabilities of the model. Besides the 366 K min^{-1} used for the calibration, three other heating rates for the decomposition of PICA are available from the work of Bessire and Minton [2017], namely, at 186, 762 and 1500 K min^{-1} . We take advantage of those additional experimental results that were not used to calibrate the model in order to compare and assess its validity when extrapolated to those heating rates. From the 10^5 samples of the Markov chain, we select evenly 10^3 samples that are propagated through the pyrolysis model at the given heating rate. The propagation of the uncertain parameters allow us to establish whether parameter uncertainties may explain the observed discrepancies or if there is modeling error.

The comparison is presented in Fig. 7.17. For 186 K min^{-1} (Fig. 7.17(a)), small discrepancies are observed below 800 K and above 1200 K that may be due to modeling error. However, there




Figure 7.15: Graphical representation of the correlation matrix. The size of the squares is proportional to the magnitude of the elements of the matrix. On the colored graph, blue squares represent a positive correlation while red squares represent a negative correlation and the magnitude of the coefficient is represented on the side bar. The diagonal is characterized by the value of the correlation equal to 1 and off-diagonal terms vary between -1 (negative correlation) and 1 (positive correlation).

is a clear discrepancy between 800 K and 1300 K due to the model and in particular due to the second branch that is being activated and is delaying the decomposition at high temperatures. At 762 K min^{-1} (Fig. 7.17(c)) there is a clear model error at low temperatures between 600 K and 800 K, but the model seems to be reproducing the mass loss within good ranges at higher temperatures. At the highest heating rate (1500 K min^{-1} , Fig. 7.17(d)), it shows substantial discrepancies, particularly below 800 K, and applying the current calibrated model to this heating rate is obviously wrong, but the total mass loss seems to be well reproduced.

To sum up, the model performs well within the calibrated range of heating rates, that is $[10, 366] \text{ K min}^{-1}$, but may fail to reproduce the mass loss of the sample outside of this range, in particular at high temperatures for low heating rates and low temperatures for high heating rates. It performs well within the calibration domain as the proposed model possesses all the features related to the mechanisms observed in this range, but it is not the case for the higher heating rates and it seems that other unknown mechanisms are presents. More complex model with potentially other competitive branches could be developed in order to capture these other features of the decomposition mechanism.

These results emphasize that care must be taken when extrapolating the pyrolysis model outside of the calibrated range of heating rates, as unknown mechanisms may be present. Therefore, the model should be refined by adding more articulated branching evolution of the present scheme in order to model these discrepancies and calibrate it with these curves. The new calibrated results would then look probably closer to the exact experimental results. Discrepancies

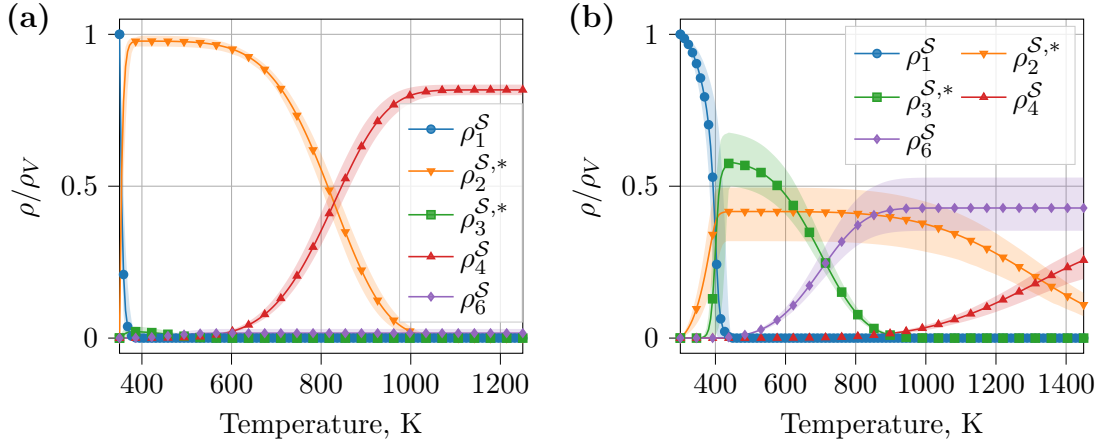


Figure 7.16: Evolution of the different solid densities for the calibrated model with their uncertainty intervals at (a) a heating rate of 10 K min^{-1} and (b) at a heating rate 366 K min^{-1} . It can be observed how the dominant pathway changes as the heating rate increases.

could also come from other mechanisms than competitive reactions and we should account for interactions between the gaseous products and the carbonaceous char, for instance, which are unaccounted for in the present model. We refer the reader to the experimental paper for a detailed discussion of the complexities that characterize the charring process of PICA (Bessire and Minton [2017]). Developing a dedicated finite rate kinetic model for the pyrolysis of PICA that fits the experimental results at all heating rates was beyond the scope of this study, but is the next necessary step towards the development of comprehensive predictive material response models.

7.6 Conclusions

In this chapter, we presented the inference of the parameters of a multi-component model that accurately reproduces the resin decomposition the production of gaseous species of PICA at 366 K min^{-1} . To explore efficiently the whole posterior, gradient-informed sampling algorithms are used, and a method for the adaptation of covariance matrix and a reparameterization of the equations are applied. The results on the sampling of the posterior show that a reparameterization based on the structure of the equations is an efficient method to improve the convergence of the RWMH algorithm. However, this reparameterization is particular to the present pyrolysis equations and requires the tuning of scaling factors, which can be challenging in high dimensions where the number of scaling factors increases. Therefore, gradient-informed sampling algorithms are used and it is showed that the adaptation of the covariance matrix based on the previous sample history improves by far their convergence rate, and can even fail if not performed. The adaptation of the covariance matrix only for a few iterations already improves significantly the convergence rate of both the HMC and ISDE, with similar performances for the two methods. Note that in the case of the HMC algorithm, the reparameterization can be chosen freely, while the ISDE is constrained to have an unbounded parameter set. In the case where the HMC algorithm is used with the reparameterization based on the pyrolysis equations, it shows very good convergence properties. A six-equation pyrolysis model is used to reproduce the production of gaseous species observed experimentally and the resulting 48-parameters posterior is sampled efficiently using the adaptive version of the ISDE algorithm. The resulting uncertainty-quantified model allows us to reproduce the data set together with the estimated observed uncertainty

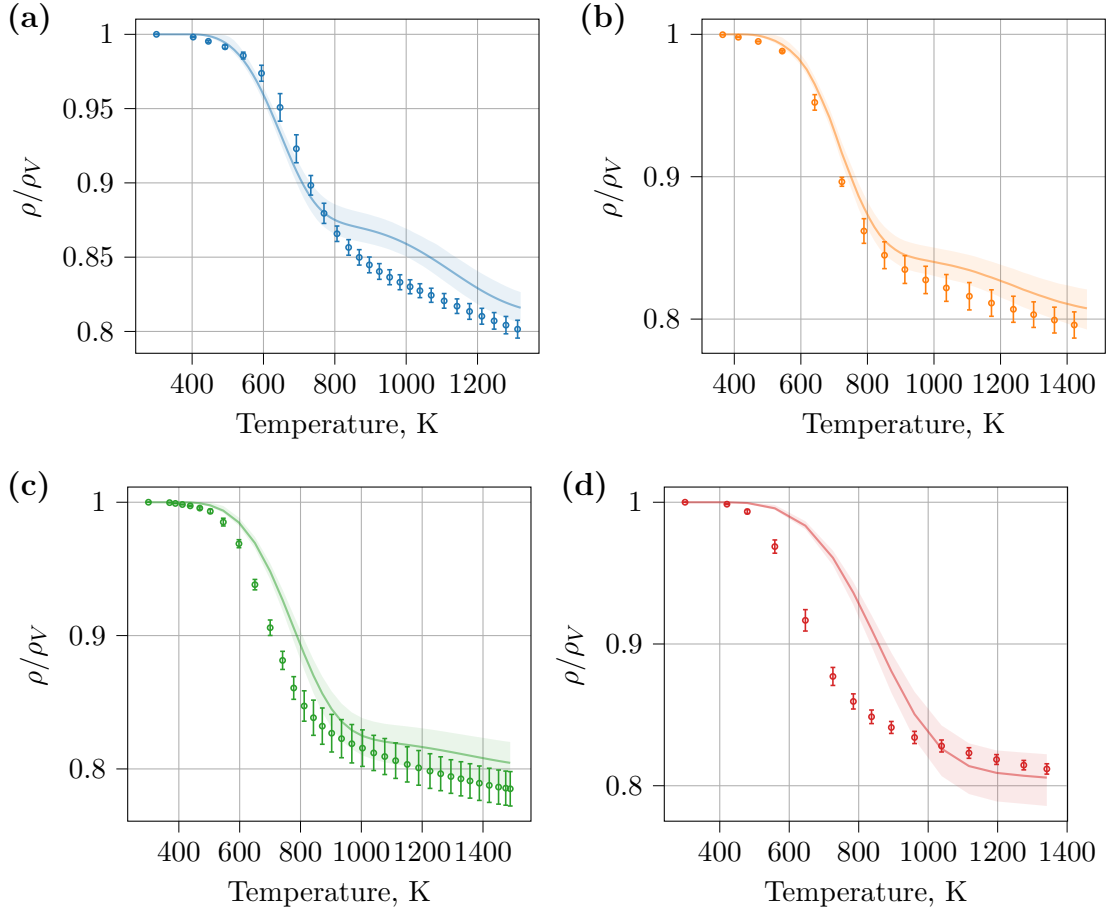


Figure 7.17: Comparison of the calibrated competitive pyrolysis model with the experimental results from [Bessire and Minton \[2017\]](#) at different heating rates: (a) 186 K min⁻¹; (b) 366 K min⁻¹ (at which the model was calibrated); (c) 762 K min⁻¹; (d) 1500 K min⁻¹.

ranges, although with some exceptions at low and high temperatures where model inadequacy is too significant due to limitations in the ability of the model to represent the physics. Results are shown for two values of the free parameter ϑ , that is introduced in the approximate likelihood to control the concentration of the posterior distribution. The choice of a value of $\vartheta = 0.7n_t$ is motivated by the quality of the overlapping of the uncertainty intervals and the observations.

Secondly, the Bayesian approach for parameter inference is applied to a more complex pyrolysis model that includes a competitive mechanism and is used to simulate the experimental data at two distinct heating rates. Some discrepancies are observed between the physical model and the experiments and it is due to the choice of the physical model that might not be perfect. It was observed that the Markov chains for several parameters were slow to converge, even for the relatively simple model with few parameters considered here. Larger uncertainty intervals are observed when the two competitive branches are active during the reaction process and additional interaction effects between parameters. Finally, we have shown the limitation of devolatilization models to predict the degradation of PICA at other heating rates.

Hence, as a future work, more complex competitive models could be developed that would allow to capture the high temperature mechanisms at low heating rate and the low temperature mechanisms at high heating rate, that were not yet accurately reproduced by the proposed

model. To improve the sampling efficiency, the use of gradient-informed algorithms with possible covariance matrix adaptation could also be investigated. Preliminary results, not presented here, showed that gradient-informed algorithms were longer to run and the time step difficult to calibrate. Indeed, a high value of the time step may provide large numerical errors due to model and gradients computation, and a low value leads to slow mixing of the MCMC chain that was prohibitive in terms of computational time.

Independently of the type of modeling for pyrolysis used, there are several interesting directions for future work that we identified. Firstly, the initial mass of the sample, used for the inference of the materials properties from the experiments of [Bessire and Minton \[2017\]](#), could also have an influence on the uncertainty in the inferred parameters. This initial mass was not communicated in [Bessire and Minton \[2017\]](#) and one had to deduce its value from the nominal experimental production curves.

Secondly, the influence of the free parameter ϑ was observed, but no clear indication for the choice of its value other than the visual inspection of the calibration curves was provided and should be further investigated.

Thirdly, in order for the iterations in the adaptive phase of the Markov chains to be used in the posterior distribution, the ergodicity of the chains generated by the gradient-informed algorithms with the adaptive procedure should be demonstrated. Numerical evidences in this work showed encouraging results in this direction.

Finally, the evaluations of the model and its gradients (there is no analytical solution for the competitive mechanisms) are not always available analytically and may be computationally expensive, in particular in high dimensions. Surrogate models could be constructed to speed-up the parameter inference process, such as polynomial chaos expansions ([Le Maître and Knio \[2010\]](#); [Xiu and Karniadakis \[2002\]](#)) for low dimensional stochastic inputs. Efficient gradient computations could be made, for instance, by using automatic differentiation ([Griewank and Walther \[2008\]](#)) or adjoint-based methods ([Cao et al. \[2003\]](#); [Bosco \[2019\]](#)). Moreover, the numerical resolution of the system of ODEs introduced numerical errors. Methods that account for the discretization errors, such as the multilevel Monte Carlo method, could be considered for the estimation of the model ([Dodwell et al. \[2015\]](#); [Beskos et al. \[2017\]](#); [Heng et al. \[2021\]](#)) and its gradients ([Jasra et al. \[2021\]](#)).

Bayesian inverse analysis of the Mars Science Laboratory atmospheric entry data for aeroheating and material response estimation

In this chapter, we infer the environment conditions at the surface of the MSL heat shield using the flight data acquired during its entry into Mars' atmosphere. The flight data considered contain in-depth temperature measurements and the reconstruction of the surface conditions is formulated as an inverse problem. We adopt a Bayesian probabilistic approach to account for measurement uncertainties and the PATO software is used along with a polynomial chaos expansion to approximate the computational model. In order to limit the error due to the in-depth material response, PICA material properties are first calibrated using the data from the shallowest thermocouple as boundary condition. Then, together with the calibrated material properties, the environment conditions at the surface of the heat shield are reconstructed. This chapter summarizes the work carried out at NASA Ames Research Center in 2019.

8.1 Introduction

Thermal protection system (TPS) thickness is typically determined by computing a nominal thickness based on the predictions from numerical simulations of the aerothermal flow around the spacecraft during the atmospheric entry and the response of the TPM exposed to the flow, then mission margins are added to this thickness ([Wright et al. \[2014\]](#)). Accurately predicting the aerothermal environment is of utmost importance as it strongly influences the final design of the TPS, but it is a challenging task that is still an active research domain.

The predictions are subjected to large uncertainties due to inaccuracies in the models used to represent the physics and their parameters, whose uncertainties impact the mission margins as they reflect some engineer's confidence in the simulations. For instance, for the Mars Science Laboratory (MSL) TPS recession, a design margin of 150 % was applied ([Bose et al. \[2014\]](#)). Ground test experiments are essential to assess the reliability of the TPS design and to improve our confidence in the models used, but the actual flight environment may not be exactly reproduced in ground facilities. Therefore, another way to improve our confidence in the models used in the numerical simulations is to compare the predictions with post-flight data, when available. This will ultimately allow to reduce the uncertainty in the predictions and TPS design for future




Figure 8.1: (a) Completed MSL heat shield and (b) MEDLI sensor locations (Cozmuta et al. [2011]).

missions (Gazarik et al. [2008]).

On 5 August 2012, the Mars Science Laboratory (MSL) entered the atmosphere of Mars and safely dropped off the 900 kg Curiosity rover on the surface of the planet. One of the key component of the success of the mission is the TPS that was implemented on MSL (Edquist et al. [2014]), shown in Fig. 8.1(a). The MSL heat shield was designed to withstand the high heat load of the reentry and in particular the fully turbulent boundary layer transition (BLT) that was expected to occur, which was unprecedented for past Mars missions. This led to the use of the PICA material, developed by NASA, that already proved its efficiency for high speed entries such as for the Stardust mission (Stackpoole et al. [2008]; Kontinos and Stackpoole [2008]).

The scientific objectives of MSL are multiple, among which is the measurement of the aerothermal environment and subsurface heating. To this end, the MSL heat shield was equipped with numerous instrumentation devices. The Mars Science Laboratory Entry, Descent and Landing Instrumentation (MEDLI), which was integrated inside the heat shield of the spacecraft, consists in pressure transducers (Mars Entry Atmospheric Data System (MEADS)), and thermocouples for temperature measurement and isotherm sensors for recession measurements (MEDLI Integrated Sensor Plug (MISP)). For MISP, the temperature data were successfully delivered, but the recession measurements did not behave as expected (Mahzari et al. [2013]).

The MISP sensors consist in circular plugs formed from the heat shield TPS that are integrated at seven locations of interests, as shown in Fig. 8.1(b). They were arranged in such a way that they would allow to detect the boundary-layer transition and to measure the in-depth temperature, thereby providing the measurements for enabling the reconstruction of the surface heating (Edquist et al. [2014]). Each MISP itself contains four in-depth thermocouples (TCs) embedded at increasing depths below the heat shield surface, which are numbered from TC1 at the top to TC4 at the bottom, with respective depths of 0.254 cm, 0.508 cm, 1.143 cm and 1.778 cm from the surface (Cozmuta et al. [2011]). The temperature measurements of each TC were recorded during the entry phase. As an illustration, we provide in Fig. 8.2 the temperature measurements of TC1 and the in-depth temperature measurements at the MISP3 position.

From Fig. 8.2(b), it can be seen that the boundary-layer transition occurs at five of the seven measurement locations. At MISP1 and MISP4 locations, which are the closest to the stagnation line, the transition was not observed. The time at which the transition occurs is around 65 s, except for MISP5 which is around 71.5 s. For the latter, the bump that characterizes the transition is barely noticeable and it is not certain that turbulent transition occurred.

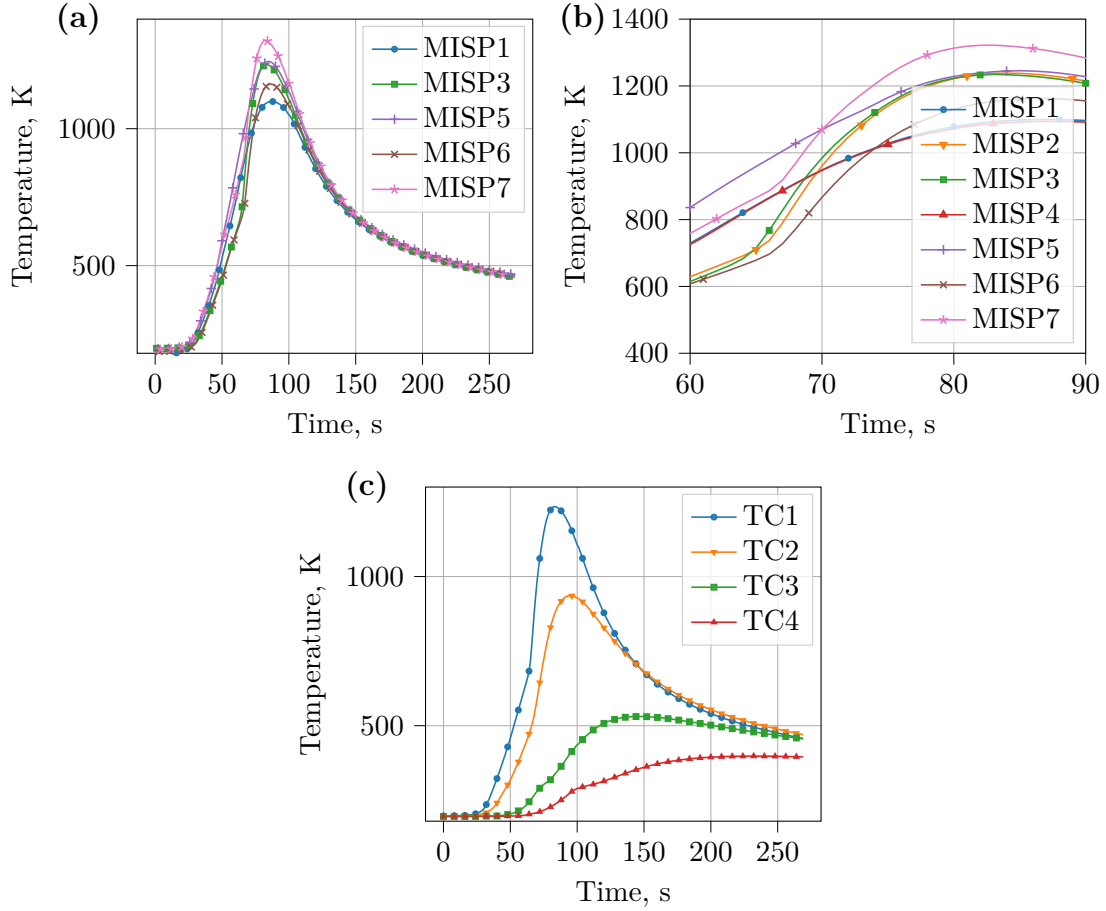


Figure 8.2: (a) Temperature data for the thermocouples closest to the surface (TC1) at every plug (MISP) locations. MISP2 and MISP4 were removed from the graph for clarity, which data are very closed to MISP3 and MISP1, respectively. (b) Zoom on the 60–90 s time frame where turbulent transition is observed. (c) The four thermocouple temperature data (TC1 to TC4) at the MISP3 location.

The overall objectives of the on-board instrumentation system is to better understand the surface environment and the material response during the atmospheric entry. Analyzing the post-flight data of the entry phase of the MSL will help to refine the models for both the entry environment and the material response. For MSL, the thermocouples that provide the temperature measurements are embedded within the heat shield, hence only allowing an indirect observation of the conditions at the surface. Thus, the problem of reconstructing the surface conditions must be formulated as an inverse problem.

There are a few studies that performed an inverse analysis for the reconstruction of the surface environment conditions and material properties from the post-flight analysis of the MEDLI data (Bose et al. [2014]; Mahzari et al. [2015]). Mahzari et al. [2015] perform an inverse parameter estimation by using a numerical optimizer that features a least-square minimization algorithm solved by the Levenberg-Marquardt method (Mahzari et al. [2011]). Optimization methods are powerful techniques to find the best parameter values that fits the post-flight data. However, they do not allow to account for the uncertainty from the measurement errors or from model inadequacy. Mahzari et al. [2015] also performed sensitivity of material properties and showed

that the heat capacity and thermal conductivity are driving parameters with a strong correlation, henceforth making their simultaneous calibration difficult.

In this chapter, we propose to build-up on the previous study of [Mahzari et al. \[2015\]](#) and address the inverse analysis of the MISP flight temperature measurements for the reconstruction of material properties and environment conditions in a probabilistic framework. In order to consider the uncertainties associated with the observations, a Bayesian approach is adopted. The model parameters are considered to be random variables and the parameter distribution is obtained by means of Bayes' formula. The collected flight data are used to update the prior parameter probability distribution and the measurement uncertainty is accounted through the likelihood function. The material response is computed using the PATO software. The posterior PDF can then be used for uncertainty propagation (Fig. 8.3), or for the comparison with CFD simulations to assess the reliability of the models. In order to reduce the influence of the in-depth material response on the reconstruction of the surface environment conditions, we first calibrate the material properties that will provide best-fit parameter values but also the probability distribution for quantifying their uncertainty. In the present work, the heat capacity is fixed and the variability in the measurement errors is captured only through the thermal conductivity parametric uncertainty. The material thermal conductivity is inferred by using the top thermocouple data as boundary condition, referred to as the TC driver approach ([Mahzari et al. \[2015\]](#)), such that this calibration does not depend on the surface environment. Then, we reconstruct the environment using the data from the two first thermocouple measurements using the calibrated material properties.

This chapter is organized as follows. We first start with a description of the methodology, that is the material-response code PATO, the Bayesian approach and the polynomial chaos approximation (Section 8.2). Then, we present the results for the calibration of the material thermal conductivity (Section 8.3), and finally the results for the reconstruction of the environment conditions (Section 8.4).

8.2 Methodology for model parameter calibration and rebuilding of surface conditions from flight data

For the design of TPSs and the analysis of material response, studies are usually performed in a one-dimensional framework, where only the heat transfer towards the subsurface is considered. We perform a similar one-dimensional analysis at each MISP location. The general methodology used in this chapter for the inference of the material properties and the reconstruction of the surface conditions is depicted in Fig. 8.3.

8.2.1 Calibration and propagation

First, a surrogate model, or emulator, for the material-response solver PATO is built. Although a single computation of the one-dimensional model is relatively cheap, running thousands of them for the Bayesian calibration could become rapidly unfeasible. Therefore, we apply the PCE method, which was described in Section 5.5.2, to build a surrogate (cheaper) model that will be used for the Bayesian calibration. The support for the surrogate is the support of the joint prior probability distribution function, which is a uniform distribution. The joint prior PDF is obtained from the tensor product of marginal uniform distributions assumed to be the prior functions for the unknown variables. The PCE expansion described in Eq. (5.70) is for a single scalar output, but can be easily extended to multidimensional outputs by collecting the model response as a function of time $\mathbf{Y} = (\eta(t_1, \boldsymbol{\xi}), \eta(t_2, \boldsymbol{\xi}), \dots, \eta(t_f, \boldsymbol{\xi}))$, where \mathbf{Y} is a random vector with values in \mathbb{R}^f , for which a PCE is built at every time steps and the coefficients are thus varying as a function of time.

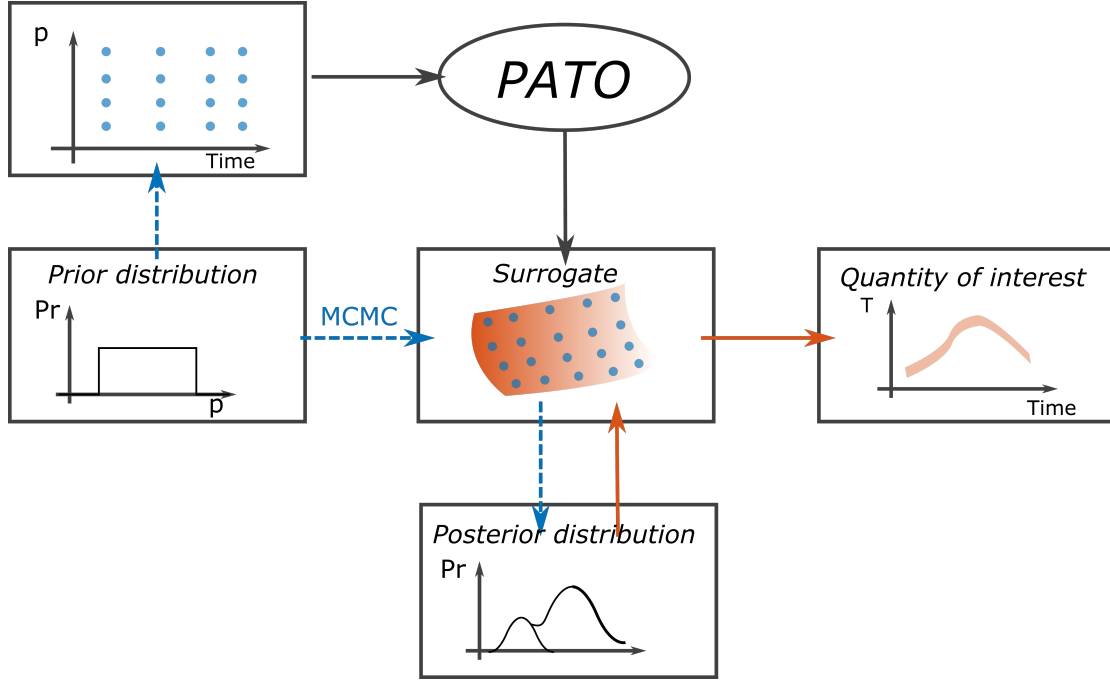


Figure 8.3: Schematic illustration of a one-dimensional parameter inference and uncertainty propagation. A surrogate model is first built as a function of time from selected design points from the prior distribution using Dakota. Posterior distribution is then inferred using Bayesian inference (Quantification of Uncertainty for Estimation, Simulation, and Optimization (Queso) package) and finally is propagated through the same surrogate model to obtain the quantity of interest.

Next, the posterior distribution is computed by means of Bayes' formula and the MCMC method is used to draw samples from it. The samples drawn from the posterior can be propagated through the surrogate model and used to characterize the quantities of interest (posterior predictive checks), such as the thermocouple data, or used to compute the surface conditions. The characterization and the propagation of uncertainty is performed using the numerical toolbox Design Analysis Kit for Optimization and Terascale Applications (Dakota), developed at Sandia National Laboratories (Adams et al. [2018]).

8.2.2 Material-response solver

The model used to simulate the material response is the PATO solver, developed at NASA Ames Research Center (Lachaud and Mansour [2014]; Lachaud et al. [2017]). PATO implements the volume-averaged Navier-Stokes equations (in the following, we will drop the symbol for the averaging operator $\langle \cdot \rangle$ for clarity) and therefore solves a momentum equation inside the material (Darcy's law), that was not allowed with previous material solver such as the FIAT solver (Chen and Milos [1999]).

Energy equation

The volume-averaged energy conservation equation implemented within PATO reads, in terms of temperature, as (Lachaud et al. [2017]):

$$\sum_{i=1}^{N_\phi} [\epsilon_i \rho_i c_{p,i} \partial_t T] - \partial_{\mathbf{x}} \cdot (K_{\text{eff}} \cdot \partial_{\mathbf{x}} T) = - \sum_{i=1}^{N_\phi} [h_i \partial_t (\epsilon_i \rho_i)] - \partial_t (\epsilon_g \rho_g h_g - \epsilon_g p) + \partial_{\mathbf{x}} \cdot (\epsilon_g \rho_g h_g \mathbf{v}_g), \quad (8.1)$$

where N_ϕ denotes the number of phase in the material, that are the fibers and the pyrolyzing phases. The second-order tensor K_{eff} is the effective thermal conductivity. Quantities $c_{p,i}$ are the specific heat capacities at constant pressure ($\text{J kg}^{-1} \text{K}^{-1}$), h_g is the specific enthalpy of the gas mixture (J kg^{-1}), ϵ_i are the volume fractions, in particular ϵ_g is the porosity, p is the pressure (Pa) and \mathbf{v}_g is the gas velocity (m s^{-1}). The quantities ρ_i denote the intrinsic densities, hence the products $\epsilon_i \rho_i$ represent the averaged densities. The energy dissipated due to the viscous interaction between the solid phases with the fluid, and the species diffusion are neglected in Eq. (8.1) (Lachaud et al. [2017]; Meurisse et al. [2018]).

Model for the effective thermal conductivity

The effective thermal conductivity K_{eff} depends on the temperature and the state of decomposition of the material. The effective thermal conductivity is assumed to be a linear function of the advancement of the reaction and the values of the effective thermal conductivities for the virgin and char materials are interpolated as

$$K_{\text{eff}}(T, \alpha) = K_{\text{eff},v}(T) (1 - \alpha) + K_{\text{eff},c}(T) \alpha, \quad (8.2)$$

where α is the global advancement coefficient of the pyrolysis reactions occurring within the material and varies from 0 in its virgin state to 1 in its charred state. The advancement coefficient is a function of the densities of the pyrolyzing sub-phases that are described by the system of ODEs introduced in Section 3.3.2.

The effective thermal conductivity should account for conduction in the solid, conduction in the gas, and the effective radiative heat transfer, which includes for instance the radiative heat exchange between the fibers through the porous media. The values for the virgin and charred effective thermal conductivities $K_{\text{eff},v}$ and $K_{\text{eff},c}$ are usually tabulated as a function of temperature (Covington et al. [2008]). However, in order to perform model parameter calibration, we need to provide a parameterized function for both $K_{\text{eff},v}$ and $K_{\text{eff},c}$. For the present one-dimensional analysis, assuming homogeneous and isotropic material properties, the second order tensor K_{eff} is a scalar, denoted by k (we drop the eff subscript for clarity), and the following third order polynomial is assumed for both the virgin and char states:

$$k_s = k_{\text{cond},s} + k_{\text{rad},s} \left(\frac{T}{T_{\text{ref}}} \right)^3, \quad (8.3)$$

where the index s stands either for the virgin state ($s = v$) or the char state ($s = c$) and T_{ref} is a reference temperature for normalization equal to 273.15 K.

Eq. (8.3), $k_{\text{cond},s}$ encompasses a contribution from the conduction in the solid and within the gas phase, although the latter is expected to be small as the pressure is relatively low. The term $k_{\text{rad},s}$ encompasses the contribution coming from the radiation between solid phases within the material separated by the gas phase. At low temperatures, radiation will be small and the main contribution to the effective thermal conductivity is due to conduction. At higher temperatures, the radiation is the dominant mode.

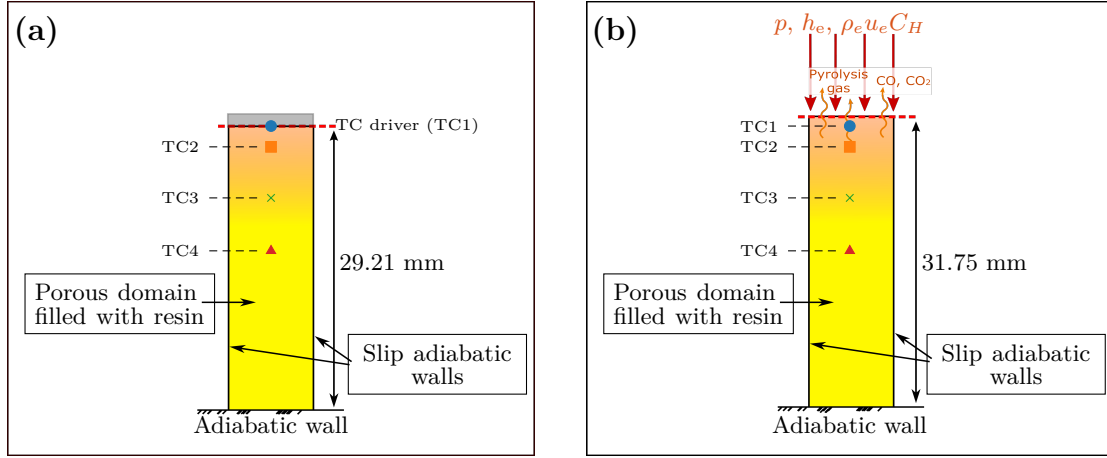


Figure 8.4: (a) One-dimensional computational domain with TC1 as boundary condition, used for the calibration of material thermal properties. The part of the material above TC1 is not simulated in the TC driver approach. (b) One-dimensional computational domain with the heating on top of the material for the calibration of the environmental conditions.

8.2.3 Computational domain

Depending on either we consider the calibration of the material properties or the reconstruction of the environment conditions, the computational domain and the boundary condition will be slightly be different. The one-dimensional computational domains for the two cases are shown in Fig. 8.4. For the calibration of material properties (Fig. 8.4(a)), that uses the TC driver approach, the temperature at the TC1 location (the shallowest thermocouple) is enforced from the measured temperature as a function of time, and thereby constitute the top boundary condition. The domain above TC1 is not simulated in this case and the total length of the sample is thus equal to 29.21 mm. The mesh is fixed as there is no surface recession.

The computational domain for the calibration of the environmental conditions is shown in Fig. 8.4(b). The top boundary conditions are imposed from the environmental conditions that are the wall ressure p , the boundary layer edge enthalpy h_e and the heat transfer coefficient $\rho_e u_e C_H$. In this case, the surface is receding and the moving mesh is updated based on the recession rate to account for the ablating surface. The total length of the sample is the total PICA thickness, which is equal to 31.75 mm.

In both cases, the mesh is composed of 36 cells in depth with the minimum cell size at the top and a cell width expansion ratio equal to 0.1. The bottom wall is assumed to be adiabatic and the sides of the domain, that are used to replicate a one-dimensional domain, are slip walls with adiabatic conditions. The pyrolysis considered is the model of Goldstein [1969] described in Section 3.5.2. Properties of the pyrolysis gas were pre-computed assuming a gas at chemical equilibrium and were tabulated as a function of pressure and temperature and provided as input to PATO.

8.3 Results for the calibration and propagation of material thermal property uncertainty

The estimation of the effective properties of a pyrolyzing material is not trivial. The properties vary with the temperature but also with the advancement of the pyrolysis reactions that defines the state of decomposition of the material (from virgin to char), which itself depends on the heating conditions. It is usual to infer material properties from ground test experiments, yet

they might be representative only of the heating conditions applied during the experiments and for which the properties are calibrated. Their extension to flight conditions, which are not exactly the same as the experimental conditions, might not be straightforward, such that the properties estimated from ground tests are left with large uncertainties. Therefore, prior to the reconstruction of the environment conditions, it is necessary to calibrate first the material properties using the actual flight data.

8.3.1 Calibration of material thermal properties

Let \mathbf{m} be a realization of the random vector of model parameters \mathbf{M} with values in \mathbb{R}^{n_u} defined on a probability space $(\Omega, \mathfrak{B}, \mathbb{P})$. For the calibration of the parameters of the effective thermal conductivity from Eq. (8.3), we have $n_u = 4$ with $\mathbf{m} = (k_{\text{cond},v}, k_{\text{rad},v}, k_{\text{cond},c}, k_{\text{rad},c})$. We select a uniform prior probability density π^0 that reflects the lack of knowledge in the parameter values and for which the support needs to be selected carefully. On the one hand, a support for the prior that is too large will result in a polynomial chaos approximation that lacks accuracy. On the other hand, a support that is too small may influence the posterior distribution. Based on previous simulation results with varying prior supports, the following joint uniform distribution is finally chosen:

$$\pi^0(\mathbf{m}) = \begin{cases} 1.33 \times 10^{20} & \text{if } \mathbf{m} \in [0; 0.5] \times [0; 10^{-10}] \times [0; 0.5] \times [0; 3e - 10], \\ 0 & \text{else.} \end{cases} \quad (8.4)$$

The likelihood function currently used in Dakota is a Gaussian likelihood, which is similar to Eq. (7.8), and reads

$$\pi(\mathcal{Y}^{\text{obs}}|\mathbf{m}) = \frac{1}{\sqrt{(2\pi)^{n_u} \prod_{i=1}^{n_{\text{obs}}} \prod_{j=1}^{n_t} \sigma_{ij}^2}} \exp \left(-\frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \sum_{j=1}^{n_t} \sigma_{ij}^{-2} (\bar{y}_{ij}^{\text{obs}} - g_{ij}(\mathbf{m}))^2 \right), \quad (8.5)$$

where n_{obs} represents all the TCs measurements for MISP1, MISP2, MISP3, MISP4 and MISP6, except from TC1 which is the boundary condition. Data for TC3 and TC4 at MISP5 and MISP7 were not recorded because of data channel limitations (Mahzari et al. [2013]) and were therefore not included in the likelihood function.

Model error is not accounted and the weights represent the standard deviations on the data, denoted by σ_{ij} . The values for the σ_{ij} need to be estimated from the raw measurements in order to compute the likelihood function. However, the raw measurements were not provided and the standard deviations on the measurements were not communicated, and therefore we assumed a constant value of 10 K on all data. This value is based on the MSL project science (PS) requirement that specifies the objectives and the desired accuracy of the measurements to be made using the MISP system (Cozmuta et al. [2011]). In particular, the PS-368 requirement states that the measured subsurface material temperature response should lie within $\pm 12\%$, as a function of time in absolute value. Assuming that it represents the region of the distribution encompassed by 2σ around the mean, which represents 95.5 % of the probability for a Gaussian distribution and is therefore assumed to be equal to 20 K, our assumption slightly exceeds the 12 % at low temperatures, but is much more restrictive at high temperatures and is thus a good compromise. We note that the choice for the standard deviations is important as their value directly affects the resulting posterior distribution on the parameters, and other assumptions could be considered or if available, directly estimated from the raw measurements. The measurements data and the standard deviations are plotted in Fig. 8.7.

The calibration is performed using Dakota (Adams et al. [2018]) with the Queso package for the Bayesian calibration. First, an emulator is built by means of a polynomial chaos expansion for the probability density π^0 for each η_{ij} , $1 \leq i \leq n_{\text{obs}}$, $1 \leq j \leq n_t$. The PCE coefficients are computed using a nested sparse grid quadrature with a second-order sparse grid level quadrature

rule computed using the isotropic Smolyak formula. For a uniform distribution, it corresponds to the Gauss-Patterson quadrature rule. Then, the Bayesian calibration is run using a delayed-rejection random-walk Metropolis–Hastings algorithm. The initial point for the Markov chain is obtained by approximating the maximum a posteriori (MAP) point from the emulator using a full Newton optimization algorithm, allowing us to start the chain from a region of high probability and resulting in a short burn-in phase in the MCMC chain. The proposal covariance is initialized by computing the inverse of the local Hessian, which is corrected in case where it is indefinite. The chain is run for 10^4 iterations and the firsts 200 samples are discarded. The 9800 remaining samples are used to compute the statistics of interests, such as the credible and prediction intervals. The results for the MCMC chain are shown in Fig. 8.5. The diagnostic of the mixing of the chains is performed visually by assessing the marginal paths of the samples, which seems to have appropriately explored the posterior distribution. As shown in Fig. 8.6, the marginal posterior distributions estimated from the marginal paths are similar to a Gaussian function.




Figure 8.5: MCMC chains for the parameters of the effective thermal conductivity model. The samples are shown after the burn-in phase.

In Fig. 8.6, the bivariate posterior distributions are also Gaussian-shaped with almost no correlation. However, it is observed from the pairwise joint samples plots that $k_{\text{rad},v}$ and $k_{\text{cond},v}$ are clearly correlated, as well as $k_{\text{rad},c}$ and $k_{\text{cond},c}$. Hence, the virgin and the char effective thermal conductivities do not depend on each other, while the conductive and radiative effects covary. It is therefore difficult to infer directly on their value based on the mere knowledge of the

temperature from thermocouple measurements. The coupling between radiation and conduction will be the most visible in the intermediate range temperature, where both conduction and radiation effects are of a similar order of magnitude.

In order to isolate the different contributions from conduction and radiation, one should use an additional output signal to the temperature measurements. In Eq. (8.3), we assumed that the conductivity term was a constant function of temperature. Based on this assumption, we could for instance infer on the virgin and char values based on pure conduction experiments on the char and virgin materials.

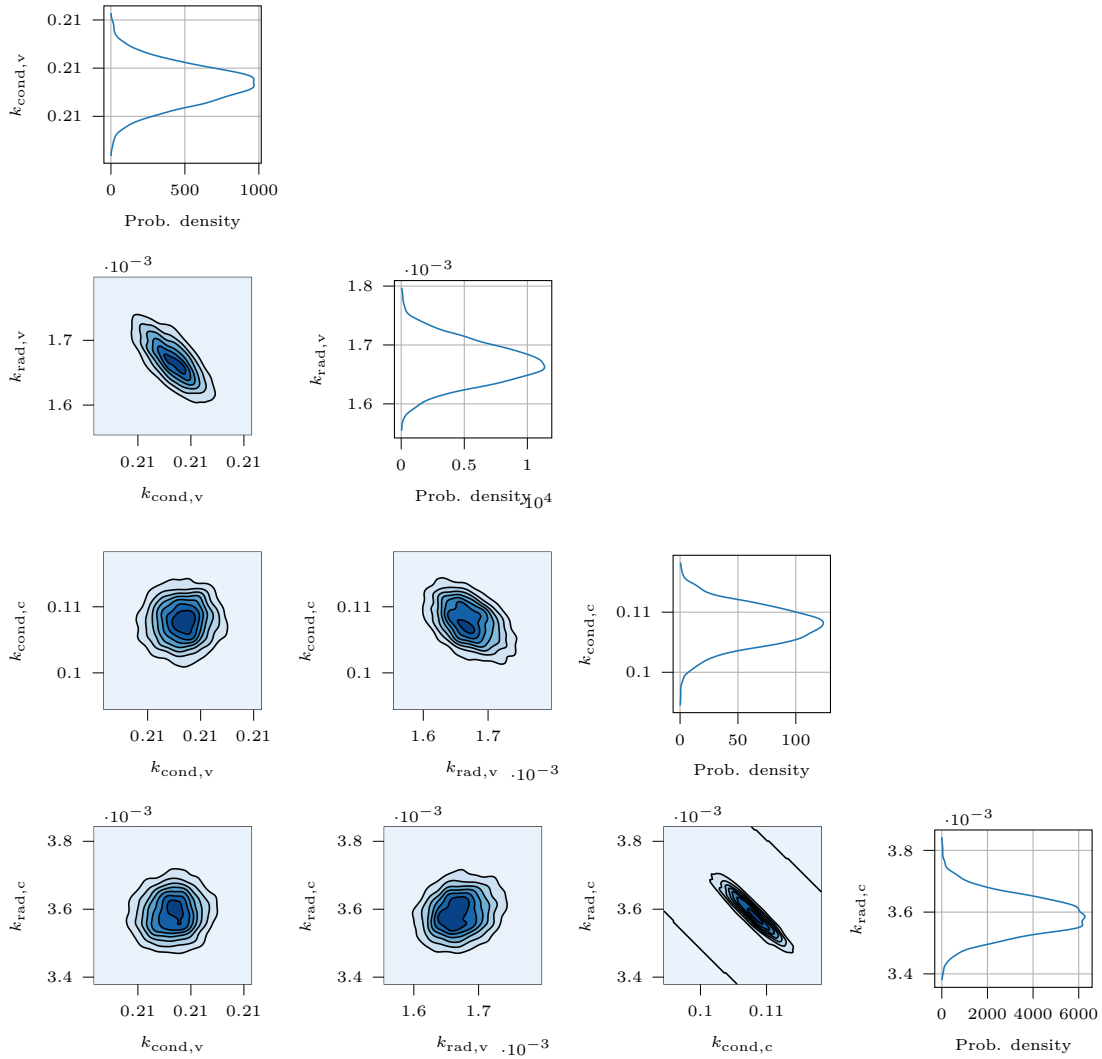


Figure 8.6: Marginal and bivariate posterior distributions for the effective thermal conductivities model parameters. Units are in $\text{W m}^{-1} \text{K}^{-1}$ for the effective thermal conductivity parameters and the probability densities have the inverse units of their variable(s).

8.3.2 Posterior predictive checks

The samples drawn from the posterior PDF obtained after the Bayesian calibration are propagated using Monte Carlo through the surrogate model build from the PC expansion. The results for the thermocouple measurements at each MISP location are shown in Fig. 8.7. The intervals represent the 95 % probability region using the samples from the MCMC chain and considering the experimental uncertainties assigned (prediction intervals). It can be observed that all the simulated temperatures are in good agreement with the experimental measurements and the simulated uncertainty ranges fall reasonably well within the post-flight data uncertainty ranges. On the contrary, the uncertainty ranges resulting only from the parameter uncertainty (credible intervals) are an order of magnitude smaller and not visible on the graphs. This small parameter uncertainty results from the fact that the data present low variability and the uncertainty is mostly accounted for by the measurement noise. However, care must be taken as the actual post-processing of the data was not communicated and any prior filtering of the data could have reduced the inferred parameter uncertainty.

Concerning the reliability of the model to represent the data, we can observe that most of the mean values fall within the experimental uncertainty ranges, except for the maximum temperature that is underestimated by the model for MISP4 at TC2 (blue curve) and overestimated for MISP6 at TC3 (orange curve), and a general trend to overestimate the temperature for $t > 200$ s at TC for all plugs. Also, one can notice the small bump around 100 s at TC4 that is visible at almost all MISPs locations, and to a lesser extent at TC3 around 75 s, in which cases the simulated temperature are outside the flight data range that, showing that the model is not able to capture this feature. This bump can be attributed to the release of water vapor that is not accounted for in the model. This model error that is not accounted in the present analysis may have an effect on the resulting posterior by shrinking the distribution.

Finally, the graph in Fig. 8.8 analyzes the effective virgin and char thermal conductivities with their uncertainty ranges as a function of temperature and compares them to the values proposed in Covington et al. [2008] for PICA, which were obtained by matching FIAT model predictions to experimental arc-jet data, to the values from the TACOT material (Lachaud et al. [2011c]), which has similar properties to PICA, and to the values for the Zeram material. The results of the present calibration lie between the thermal conductivities for PICA and TACOT. It is interesting to observe that in the present analysis, the virgin and char conductivities does not have the same value, whereas it was the case for the properties of the other material, except for $T > 1100$ K for Zeram. Here, the thermal conductivity of the virgin material is lower than the char, then the two curves cross each other at 1035 K and finally the trend reverses afterwards. In that sense, it is closer to the Zeram values, where the virgin conductivity is first higher than the char conductivity. However for the Zeram, the two values collapse after 1110 K, because at higher temperature there is almost no resin left. Assessing and measuring the conductivity of the material during pyrolysis is not trivial and is an on-going research area. Notably, the assumption of a virgin thermal conductivity at high temperatures is questionable: at those temperatures, the resin has already reached an advanced state of decomposition. Thus, the relation used from Eq. (8.3) is purely empirical and the calibrated parameters only correspond to the present flight conditions.

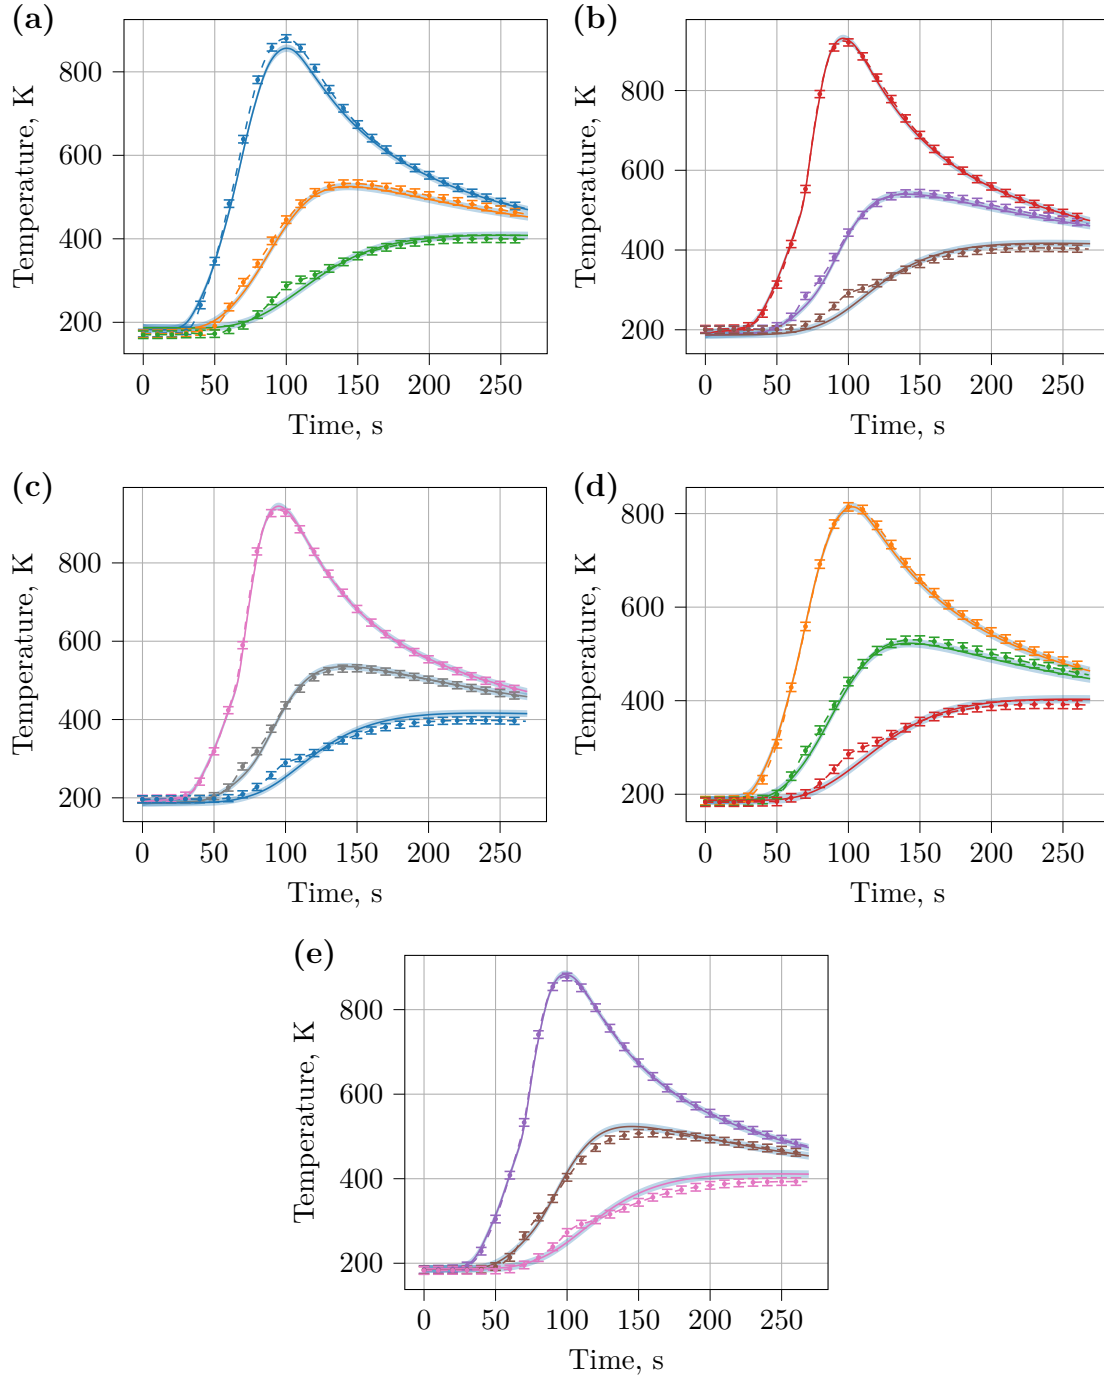


Figure 8.7: Propagation of material property uncertainty characterized by the MCMC samples shown for the simulated thermocouples data as a function of temperature for (a) MISP1, (b) MISP2, (c) MISP3, (d) MISP4, and (e) MISP6. The continuous lines represent the means and the thin shaded areas around the means are the 95 % prediction intervals. On each plot, the curves represent the measurements for TC2, TC3 and TC4 from top to bottom, respectively. The actual post-flight data are depicted by dots with error bars that represent two times the standard deviation.




Figure 8.8: Comparison of the calibrated thermal conductivities as a function of temperature with the values from [Covington et al. \[2008\]](#) (triangle up markers), values from TACOT (triangle down), Zuram database (circles). The continuous lines (blue) are the virgin conductivities and the discontinuous lines are (orange) are the char conductivities. The shaded areas represent the total probability interval.

8.4 Rebuilding of the environment conditions

In this section, we address the reconstruction of the environment from the thermocouple measurements and their uncertainty. The methodology is similar to as in the previous section, except that now the computational domain (Fig. 8.4(b)) represents the whole thermal protection material thickness and the top boundary condition is the aerothermal environment. We also note that, while in the previous section a common value for the thermal conductivity at all MISPs was used thanks to the assumption of the homogeneity of the material properties, this can no longer be the case for the surface environment conditions which evolve along the heat shield surface. Thus the reconstruction of the environment conditions has to be performed separately at each MISP position. There are two additional challenges that need to be addressed, related to the selection of the prior domain and to the turbulent transition at the surface of the material.

Challenge in selecting the prior domain for the PCE construction

Because the boundary conditions are unknown, there might be an issue in building the surrogate model due to the divergence in the material response solutions at a given depth within the material, in particular for the thermocouple closest to the surface (TC1). A small variation in the boundary conditions may ultimately lead to a recession of the surface that reaches the thermocouple position and results in its destruction (burn out). For instance, a combination of a high boundary layer enthalpy and a high heat transfer rate may lead to a situation where the surface of the material actually reaches TC1. When such a condition occurs, PATO automatically sets the thermocouple value to zero, as illustrated in Fig. 8.9, thereby degrading the quality of the fit of the polynomial chaos expansion when incorporating such a solution in the model approximation.




Figure 8.9: Illustration of the case of a higher surface heating that leads to a recession of the surface below the location of TC1 and before TC2. The numerical solver PATO sets the temperature numerically to zero, which value should not be used in the PCE construction.

Anticipating a priori which conditions will lead to the burn out of the thermocouple is not obvious. Fortunately, the burn out of thermocouples was not observed during the atmospheric entry of MSL and the prior domain, that defines the bounds of the initial parameter support, must be selected carefully in such a way that no burn out is observed. In practice, the prior domain is obtained by trial-and-error: when the parameter values at the extremities of the prior distribution lead to a surface recession that ultimately reaches the thermocouple, the prior domain is slightly reduced until no burn out of the top thermocouple is observed.

Challenge with turbulent transition

Pre- and post-flight analyses demonstrated that turbulent transition was present at most MISP locations before peak heating. BLT and transition times can be observed (with some delay) on top thermocouple (TC1) temperature data for which a sudden change can be observed in the temperature curve, as seen in Fig. 8.2(b). BLT is still visible on deeper TCs but its impact is milder. Because of the transition to turbulence, the functional assumed for the transfer coefficient can vary from one location to the other.

We assume the heat transfer coefficient to be a sum of N_M Gaussian-shaped functions parameterized as:

$$\rho_e u_e C_H = \sum_{i=1}^{N_M} C_{H,i}^a \exp\left(-\frac{(t - t_{\mu,i})^2}{t_{\sigma,i}^2}\right), \quad (8.6)$$

where C_H^a is a constant that has the units of $\rho_e u_e$ ($\text{kg m}^{-2} \text{s}^{-1}$), t_{μ} (s) denotes the mean of the Gaussian, which controls its position, and t_{σ} (s) denotes the standard deviation of the Gaussian and controls its width. While the assumption of a single Gaussian-shaped function seems reasonable at the MISP locations where no transition occurs, it may not be able to represent the heat transfer coefficient where transition occurs. Indeed, the transition induces a change in the curve for the heat transfer coefficient that can not be captured by a single Gaussian function. By having a sum of Gaussian functions, one can represent rapid changes in the heat transfer coefficient, for instance, by tuning the value of the standard deviation t_{σ} of the individual functions: a low value will result in a sharp increase of the heat transfer coefficient, hence allowing to represent the transition to turbulence, and a high value will represent a slower increase. In practice, we will not use more than two components in the sum from Eq. (8.6), where the second function is narrower and will help to capture the a rapid change in the curve.

The BL edge enthalpy does not depend on whether the BLT occurs or not, and is parameterized by the following third order polynomial:

$$h_e = h_{e,0} + h_{e,1} \left(\frac{t}{t_0}\right) + h_{e,2} \left(\frac{t}{t_0}\right)^2 + h_{e,3} \left(\frac{t}{t_0}\right)^3, \quad (8.7)$$

where t_0 is a reference time for normalization taken to be equal to 1 s. We assume the pressure to be known (which in practice has little uncertainty), whereas the parameters for modeling the heat transfer coefficient and the enthalpy are considered to be uncertain.

8.4.1 Results of the calibration and posterior predictive checks

In the following, the environment conditions are calibrated at MISP3 and MISP4 positions. The material conductivity is fixed at the MAP value obtained in the previous section.

At MISP3, the transition to turbulence occurs and a two-component Gaussian mixture is assumed for the heat transfer coefficient, which results in 6 parameters to be calibrated. For the enthalpy, only the leading coefficient $h_{e,0}$ is assumed to be unknown for simplicity. Thus, we have $n_u = 7$ with $\mathbf{m} = (C_{H,1}^a, t_{\mu,1}, t_{\sigma,1}, C_{H,2}^a, t_{\mu,2}, t_{\sigma,2}, h_{e,0})$. The support \mathcal{S} of the uniform prior PDF is given in Table 8.1 which provides the lower (L_b) and upper (U_b) bounds of the support in each dimension. The prior expresses as

$$\pi^0(\mathbf{m}) = \begin{cases} \frac{1}{U_{b,i} - L_{b,i}} & \text{if } \mathbf{m} \in \mathcal{S}, \\ 0 & \text{else.} \end{cases} \quad (8.8)$$

The MCMC chains are shown for $C_{H,1}^a$, $t_{\mu,1}$, $t_{\sigma,1}$ and $h_{e,0}$ in Fig. 8.10. A total of 10^4 MCMC iterations are run and the 200 first iterations are discarded. All the parameters seem to have converged with a good mixing (also for the $C_{H,2}^a$, $t_{\mu,2}$, $t_{\sigma,2}$ that are not showed). The estimated

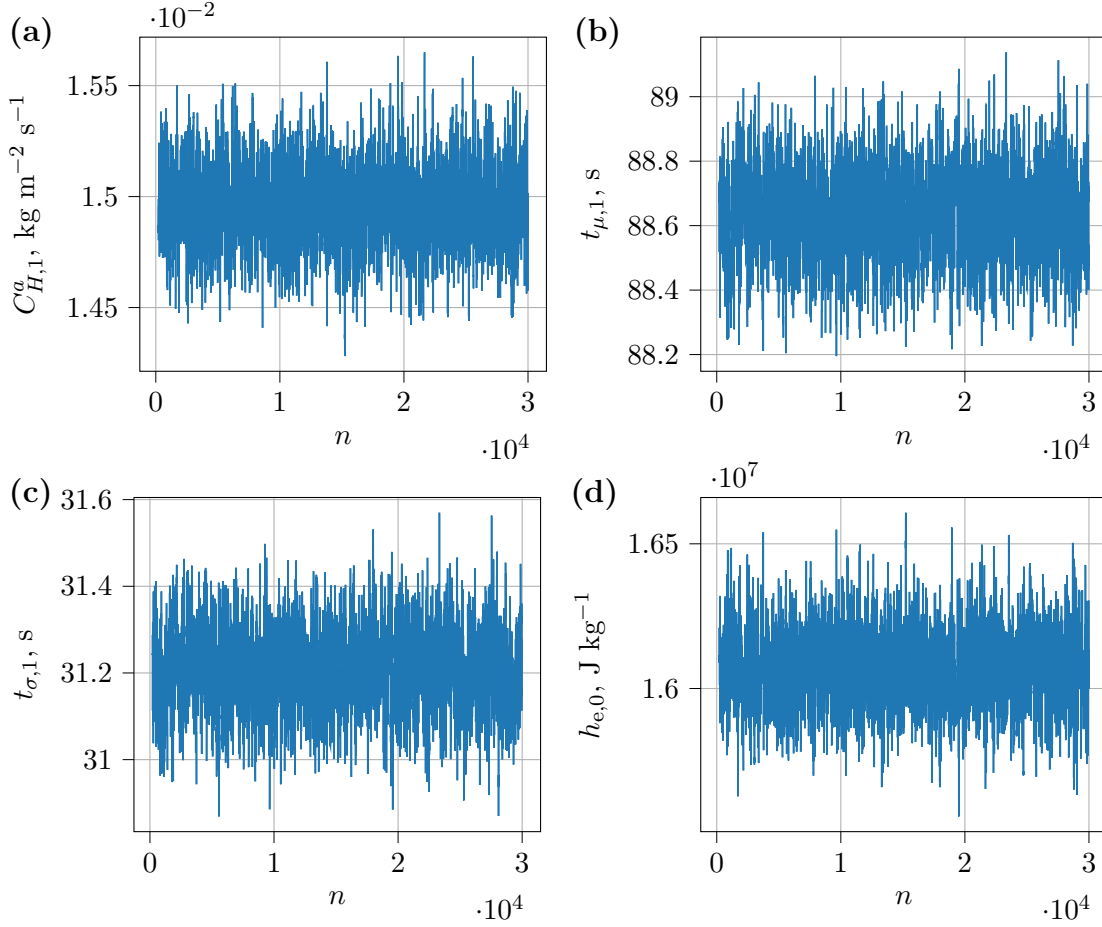


Figure 8.10: MCMC chains for selected parameters for the heat transfer coefficient and the BL edge enthalpy at MISP3.

marginal and bivariate posterior probability densities for the same parameters are shown in Fig. 8.11. We can observe a negative correlation between $h_{e,0}$ and the three parameters from the Gaussian function. This behavior is expected, as an increase in the boundary layer edge enthalpy value must result in a decrease in the heat transfer coefficient value in order to keep the overall heat transfer constant.

The simulated temperatures for TC1 and TC2 at MISP3 obtained at the sample values are shown in Fig. 8.15(a) and are compared to the actual flight data (dotted lines). The propagated temperatures are in good agreement with the thermocouple temperature data. Some discrepancies are observed at TC1, which is the closest to the surface, but it is almost in a perfect agreement at TC2. Indeed, the effects of the surface conditions are less visible on deeper thermocouples. The transition to turbulence is also well captured by the two-component Gaussian function for the heat transfer coefficient, with once again a better agreement for TC2, and the computed uncertainty ranges reproduce the data well.

The propagation of the MCMC samples through Eq. (8.6) and Eq. (8.7) are shown in Fig. 8.12 and compared to the values provided from CFD simulations. In this case, the comparison of the calibrated environment with the CFD does not provide a satisfactory agreement, except for $t < 48$ s and $t > 110$ s for the heat transfer coefficient. The curves provided from the CFD simulations feature discontinuities whereas the calibrated curves are smooth.

Therefore, whereas the reconstructed environment seems to be in good agreement with the

Table 8.1: Values of the bounds of the prior support for MISP3.

	$C_{H,1}^a$ ($\text{kg m}^{-2} \text{s}^{-1}$)	$t_{\mu,1}$ (s)	$t_{\sigma,1}$ (s)	$C_{H,2}^a$ ($\text{kg m}^{-2} \text{s}^{-1}$)	$t_{\mu,2}$ (s)	$t_{\sigma,2}$ (s)	$h_{e,0}$ (J kg^{-1})
L_b	0.005	50	5	0.035	50	5	$0.7\text{e}7$
U_b	0.03	100	35	0.052	100	20	$2\text{e}7$

calibration data from TC1 and TC2 and allows to simulate with a good accuracy the temperature profiles, it seems that it is not in agreement with the CFD results. Moreover, the measurement errors only result in small uncertainty on the parameter and can not explain the observed differences.

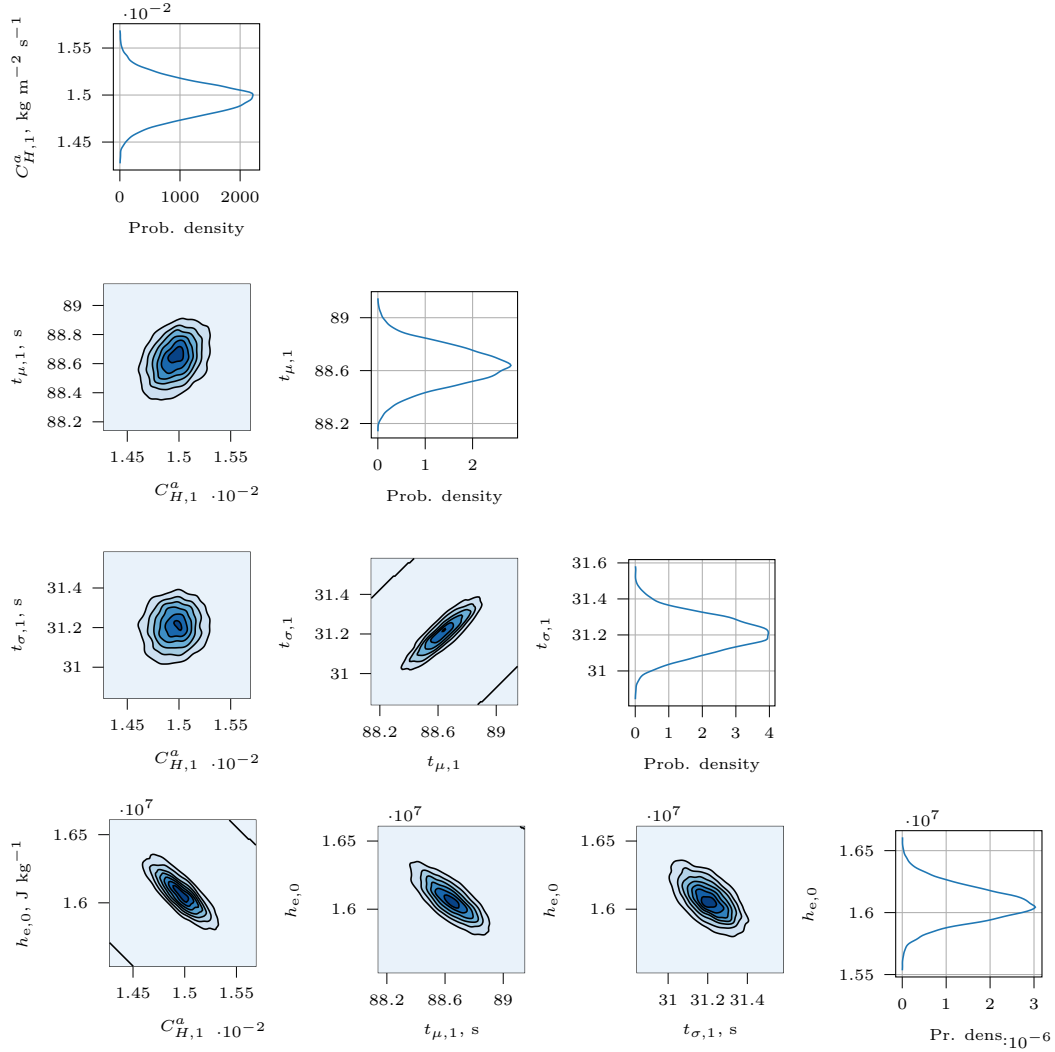


Figure 8.11: Selected marginal and bivariate posterior probability densities for the parameters of the heat transfer coefficient and BL edge enthalpy at the surface of MISP3. The probability densities have the inverse units of their variable(s).




Figure 8.12: Comparison of CFD (red line with markers) and reconstructed environment using Bayesian calibration at MISP3 for (a) heat transfer coefficient (b) BL edge enthalpy. The continuous lines (in blue) are the median values and the shaded areas are the total (100 %) probability interval that represent the uncertainty due to the overall model parameter variability.

At MISP4, where no transition to turbulence is observed, a single Gaussian function for $\rho_e u_e C_H$ (in the following, we drop the index $i = 1$) is used for the heat transfer coefficient. Only two parameters are assumed to be uncertain, namely C_H^a and t_μ , whereas t_σ is fixed to 27.88 s. The burn-out of TC1 was observed when considering the latter to be unknown. Again, for the enthalpy only the leading coefficient $h_{e,0}$ is assumed to be unknown. Hence, $n_u = 3$, $\mathbf{m} = (C_H^a, t_\mu, h_{e,0})$ and $\pi^0(\mathbf{m})$ is the uniform distribution from Eq. (8.8) with support $[0.01; 0.04] \times [50; 100] \times [0.7e7; 2e7]$. The delayed-rejection Metropolis–Hastings algorithm is run for 10^4 iterations and the MCMC chains for the three parameters are shown in Fig. 8.13. The chains are well mixed for the three parameters. Bivariate and marginal posteriors are shown in Fig. 8.14. As for MISP3, we can observe a negative correlation between $h_{e,0}$ and the two other parameter, which is even more pronounced in the case of the $h_{e,0}$ and C_H^a parameter pair. This high correlation suggests that one cannot simultaneously identify the value of these two parameters using only the current thermocouple observations, and results in an increase in the parameter uncertainty.

From the result in Fig. 8.15(b), the single Gaussian function seems to be in agreement with the general trend of the temperature curve, except for a shift in TC1 that is observed around the maximum temperature and the temperature that is slightly overestimated for TC2 around 50 s.

Both for MISP3 and MISP4, one can note that the computed probability intervals are larger than the ones computed for the calibration of the material properties from Fig. 8.7. Compared to Fig. 8.7 less temperature data are used in the likelihood function. Indeed, only the two top TC data are used at a single MISP, while in Fig. 8.7 all the MISPs and the three bottom TC data are included. In agreement with the Bayesian tenets, this results in less confidence on the calibrated parameters and propagated probability intervals that are smaller.

8.5 Summary and perspectives for future work

In the first part of this chapter, we addressed the inference of material thermal conductivity in a Bayesian framework. We observed correlations between the virgin radiative and conductive coefficients, and between the char radiative and conductive coefficients, showing that the radiative and conductive contributions could not be separated. However, no correlation between the

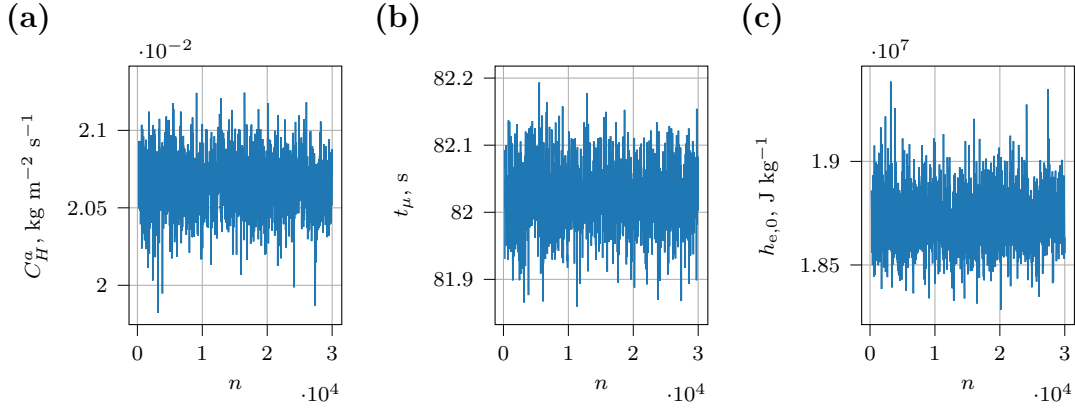


Figure 8.13: MCMC chains for the parameters for the heat transfer coefficient and the BL edge enthalpy at MISP4.

virgin and the char coefficients was observed, showing that the virgin and char contributions can be isolated (at the temperatures considered).

We were able to simulate the thermocouple data and their uncertainty with good accuracy. The PDFs on the calibrated parameters showed little uncertainty on the parameter values, which can be understood from the relatively high number of data used for the inference. Hence, we could highlight model error, such as the release of water vapor that was not accounted by the material response model, which could in turn have an effect on the posterior distribution. This model error could be included in the uncertainty analysis by incorporating an additional term in the statistical model (Kennedy and O’Hagan [2001]), or by providing more flexibility in the likelihood function by introducing an additional free parameter as discussed in Chapter 7.

In the second part of this chapter, we presented the results for the reconstruction of the environment (convection at the wall and BL enthalpy) using data from the two shallowest thermocouples. We emphasized the issue in the assessment of the prior distribution for the construction of the PCE due to the high sensitivity of the material recession on environment conditions and the possible burn out of TCs, and the issue of turbulent transition on the assumed boundary condition functional. We showed that the environmental conditions could be reconstructed from the TC measurements using a one or two-component Gaussian function for the heat transfer, depending on whether turbulent transition occurs or not, and a polynomial function for the boundary layer edge enthalpy. Again, the Bayesian inverse procedure allowed us to simulate with good agreement the flight thermocouple temperature data and their uncertainty. However, the comparison with the surface conditions provided from CFD results did not show to be in agreement with the reconstructed curves. Discrepancies between the predicted and reconstructed environments could be due to physical phenomena that are not accounted in the CFD model, but it would be first needed to investigate other methods for the reconstruction of the environment. Future studies could investigate different functions for the heat transfer coefficient and the BL edge enthalpy, or, for the latter, include more uncertain parameters. We further note that the rebuilding was performed using MAP value for the material parameter conductivity, and its uncertainty should also be included in the analysis, for instance, by inferred all parameters at the same time. We also observed a correlation between the parameter from the BL enthalpy and the parameters from the heat transfer coefficient, which could indicate that one cannot simultaneously identify their values based only thermocouple data, and that other measurements should be used, such as the surface recession, which was not usable in the case of MSL.

The comparison with CFD results was performed only at MISP3, but it should also be done at other locations. For the comparison, it would also be valuable to compute the uncertainty in

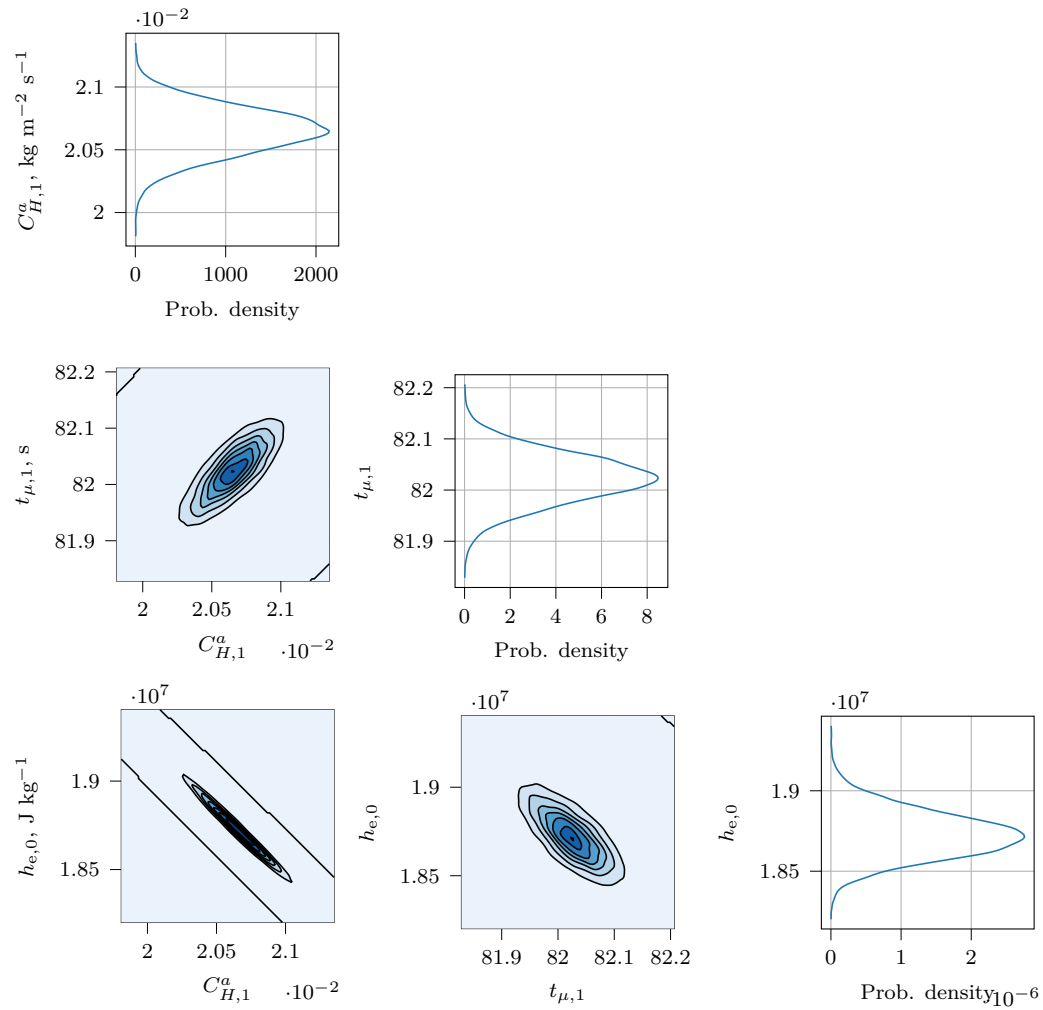


Figure 8.14: Marginal and bivariate posterior distributions for model parameters of the environment conditions at the surface of MISP4. The probability densities have the inverse units of their variable(s).




Figure 8.15: Comparison of the calibration data and the posterior predictions. TC1 (blue) and TC2 (orange) measurements with data (dotted line with standard deviation), PATO results (continuous line) for (a) MISP3 and (b) MISP4.

the surface condition by propagating the uncertainty in the post-flight parameters measurements (e.g. angle of attack or vehicle velocity) directly using the CFD solvers (e.g. LAURA (McNeil and Peter [1996]) or DPLR (Wright et al. [2009])), see for instance Hosder and Bettis [2012], for which the methodology could be used to obtain input heat flux uncertainty.

Finally, for both the calibration of the material properties and environment conditions, a constant value of 10 K was assumed for the standard deviations on the measured temperatures. It could be interesting to investigate the effect on the resulting posterior distribution when this value varies as a function of temperature (e.g. by considering a $\pm 12\%$ value as a function of time). If the raw flight data can be provided, an estimation of the standard deviation could be computed, or it could be included as a hyperparameter and inferred together with the material parameters. Obviously, larger standard deviations on the temperatures would result in a wider uncertainty on the parameters and the propagated quantities of interest.

Propagation of uncertainty in pyrolysis model parameters for a Martian atmospheric reentry

This last chapter is concerned with the propagation of uncertainty in the pyrolysis kinetic parameters and mass coefficients to study the uncertainty in the heat shield recession, temperature and gas production. We demonstrate the use of the identified six-reaction pyrolysis model in a 3D numerical simulation of heat shield decomposition during a Martian entry. The pyrolysis model parameters are considered as uncertain input parameters and the posterior probability distribution resulting from the Bayesian approach to parameter inference is fully characterized by the samples from the Markov chain. The samples are propagated through the high-fidelity PATO numerical solver using the Monte Carlo approach. I would like to acknowledge J. Meurisse for the discussions and for generating the meshes for the MSL simulations.

“Prediction is very difficult, especially if it’s about the future.”

Niels Bohr

9.1 Introduction

For the assessment of TPS margins, there is a need to predict the contribution of the uncertain parameters on the surface recession and heat shield temperature, which drive the design of the TPS and for which the predictions are still largely uncertain. Among all the input parameters, pyrolysis model parameters are expected to play a significant role ([Rivier et al. \[2019\]](#)). Indeed, the effective material properties are directly influenced by the resin decomposition rate through the advancement coefficient, which in turn influence recession and sub-surface temperature. The species production, which is proportional to the mass fractions, is also expected to have a large impact on the boundary layer gas composition which may encompass species that can further radiate and increase the overall heat shield temperature. Predicting the boundary layer gas composition is crucial to understand the contribution of the radiative heat flux to the heat shield, which can constitute an important part of the total heat flux. For instance, recent measurements from the Schiaparelli lander highlighted the important contribution of the radiative heat flux on the back cover of the heat shield ([Gülhan et al. \[2019\]](#)). A maximum radiative contribution of 61 % of the total heat flux towards the back cover was measured at 115 s after the beginning of

the atmospheric entry. Therefore, accurately predicting the boundary layer gas composition will allow to better characterize species radiation. Hence, in this chapter, we seek to propagate the uncertainty in pyrolysis model parameters, characterized by PDFs on the input model parameters, through numerical solvers in order to assess the output PDFs on recession, temperature, as well as on the species production at the surface of the heat shield.

Several studies already considered the propagation of input PDFs for quantifying the uncertainty in the modeling of TPS for atmospheric entry. [Copeland et al. \[2012\]](#) used the direct Monte Carlo method to perform uncertainty quantification in one-dimensional material response simulations using the FIAT code. The input PDFs considered were the material density, specific heat and thermal conductivity (both char and virgin), which were characterized by means of normal distributions with hyperparameters estimated from available experimental data. They compared the Pearson correlation coefficients and Sobol indices for global sensitivity analysis and showed that a nonlinear analysis of the uncertainty is necessary. In order to determine the distribution of the quantity of interest with a good accuracy, a large number of Monte Carlo simulations may be required which can be computationally expensive. In order to reduce the computational time of the Monte Carlo method, [Anzalone et al. \[2016\]](#) applied the Quadrature Methods of Moments, in which moments of the input distributions are propagated to obtain moments on the quantity of interest, and they studied the influence of material parameters on the surface recession rate using a one-dimensional ablation code. Frequently, the numerical model considered is also computationally expensive, and methods that seek to approximate the computationally expensive model by a (cheaper) surrogate model are preferred. Polynomial chaos methods have received a lot of attention in the past years, as for instance in [Hosder and Bettis \[2012\]](#) who performed uncertainty propagation and sensitivity analysis for non-ablating TPMs using a non-intrusive polynomial chaos method. They studied the uncertainty in the heat flux at the surface of the material due to uncertainties in the freestream, the recombination efficiency due to wall catalysis and binary-collisions integral using CFD simulations (using the DPLR code). [Turchi et al. \[2017b,a\]](#) used polynomial chaos for uncertainty propagation and apply an anchored analysis of variance (ANOVA) ([Tang et al. \[2015\]](#)) method for sensitivity analysis to study the influence on surface mass blowing and surface temperature using the VKI Stagnation Line code. [Rivier et al. \[2019\]](#) also used the polynomial chaos and anchored-ANOVA methods using PATO for the PICA ablative TPM. Independent uniform input probability distributions were assumed and it was shown that parameters from pyrolysis reactions, such the activation energies, the pre-exponential factors or the pyrolysis gas composition, contribute to the uncertainty on the material recession, char and virgin fronts and surface temperature.

Past studies usually consider standard distributions as input PDFs and neglect potential correlations between the parameters. In Section 7.4, we inferred the uncertainty in pyrolysis model parameters in a Bayesian framework and probability distributions were directly obtained from experimental data. We did not introduce additional stochastic terms in the pyrolysis model in order to allow the identified pyrolysis model to be used in material response codes for numerical simulations of heat shield decomposition in atmospheric entries without requiring stochasticity to be implemented explicitly inside these codes. Notably, our approach allows the uncertainties associated with the identified pyrolysis model to be propagated through such a numerical simulation in an atmospheric entry by using a Monte Carlo approach implemented as a wrapper around the code. The probability distribution on the pyrolysis model parameters can be non-standard with correlated variables, of high dimension and the samples drawn through MCMC methods can be correlated. The Monte Carlo approach for uncertainty propagation has the additional advantage of being insensitive to the curse of dimensionality and of being applicable to any type of distribution provided that we can draw samples from it. The deterministic pyrolysis model is thereby provided to the material-response code PATO, which is then run for multiple values of the model parameters drawn from the posterior. In this chapter, we demonstrate such a use of the identified pyrolysis model in a numerical simulation of heat shield decomposition during a Martian entry.

This chapter is organized as follows. In Section 9.2, we describe the numerical set-up considered for the simulations of the Martian atmospheric entry and we compare the result for the nominal parameter values using the six-reaction pyrolysis model and the two-reaction pyrolysis model from Goldstein [1969]. In Section 9.3, we present the results of the uncertainty propagation on the heat shield recession, the surface and in-depth temperatures, and the surface gas composition. Finally, the conclusions are presented in Section 9.4.

9.2 Numerical model

We consider a reentry in the atmosphere of Mars and a thermal protection system similar to the MSL heat shield. Deterministic simulations of an atmospheric entry on Mars using PATO were already performed in Meurisse et al. [2018] and we consider a similar numerical set-up. The MSL heat shield is a rigid 4.5 m diameter spherically-blunted half-angle cone made up of 113 PICA tiles that are assembled together with a silicone elastomer bonding (Edquist et al. [2009]; White et al. [2013]), see Fig. 8.1(a). A uniform PICA thickness of 31.75 mm was established to be sufficient to withstand the reentry at an acceptable risk (Wright et al. [2014]).

The PATO software (Lachaud and Mansour [2014]), developed at NASA Ames Research Center, is used to perform 3D simulations of the full heat shield during an atmospheric entry. PATO solves the conservation equations of solid mass, gas mass, gas momentum and total energy for a porous reactive material using the volume-averaged formulation, which was described in Section 2.3. Assuming local thermal equilibrium at the pore scales, material pyrolysis decomposition, production of gases from the pyrolysis reactions and ablation, and homogeneous chemical reactions in the gas phase are modeled. The heterogeneous chemical reactions (gas-surface interactions) are represented using equilibrium chemistry models that are preferred for design due to the lack of reliability of available finite-rate chemistry models and data (Meurisse et al. [2018]).

9.2.1 Pyrolysis model in PATO

The parameters of a six-reaction pyrolysis model were inferred in Section 7.4. However, because of the difference between the pyrolysis model implemented in PATO (see Lachaud et al. [2017]) and the model described in Eqs. (3.11)–(3.13), the values for the mass fractions F , for which the MAP values are summarized in Table C.1, need to be slightly adapted.

The pyrolysis model implemented in PATO is a general model that consider $1 \leq i \leq N_\phi$ decomposing phases, each of which has different sub-phases to model its $1 \leq j \leq N_p$ degradation mechanisms (that are the fictitious pyrolyzing components of the resin, as described in Section 3.3.2). The total production rate ($\text{kg m}^{-3} \text{s}^{-1}$) of species (or element) $1 \leq k \leq N_g$ in the gas mixture is expressed in terms of density variables and writes as

$$\pi_k(t) = \sum_{i=1}^{N_\phi} \sum_{j=1}^{N_p} \gamma_{ijk} \epsilon_{i,0} \rho_{i,0} w_{ij} \frac{d\alpha_{ij}}{dt}(t), \quad (9.1)$$

where α_{ij} denotes the advancement coefficient of the reaction, $\epsilon_{i,0} \rho_{i,0}$ the initial averaged density and γ_{ijk} the mass coefficient of species k for the phase i and sub-phase j .

The model from Eqs. (3.11)–(3.13) is expressed as a function of the total initial mass of the sample m_0 , or dividing by the volume, the total solid density $\epsilon_0 \rho_0$, and not in terms of the single initial resin content ($i = 1$) $\epsilon_{1,0} \rho_{1,0}$ for carbon/phenolic TPM. The calibrated parameters need to be corrected in order to be compatible with the model implemented in PATO. The total solid density is the sum of the fiber and resin densities, $\epsilon_0 \rho_0 = \epsilon_{f,0} \rho_{f,0} + \epsilon_{1,0} \rho_{1,0}$ ¹, hence, Eq. (9.1) for

¹For the PICA TPM, the averaged densities are $\epsilon_f \rho_f = 160 \text{ kg m}^{-3}$ and $\epsilon_0 \rho_0 = 280 \text{ kg m}^{-3}$.

Table 9.1: Nominal value for the elemental mass fractions of the pyrolysis gas produced by the six-reaction pyrolysis model.

	C	H	O
R ₁	0.243	0.034	0.723
R ₂	0.524	0.063	0.413
R ₂	0.664	0.075	0.261
R ₄	0	1	0
R ₅	0.714	0.073	0.213
R ₆	0.323	0.039	0.638

the single phenolic pyrolyzing phase ($N_\phi = 1$ and we drop the index i), can be written as

$$\pi_k(t) = \sum_{j=1}^{N_p} \gamma_{jk} \epsilon_0 \rho_0 \left(1 - \frac{\epsilon_{f,0} \rho_{f,0}}{\epsilon_0 \rho_0} \right) w_j \frac{d\alpha_j}{dt}(t), \quad 1 \leq k \leq N_g. \quad (9.2)$$

By equating Eq. (3.13) (divided by the volume) and Eq. (9.2), the mass fractions F_{ij} calibrated in the previous chapter can be expressed as

$$F_{ij} = \gamma_{ij} w_i \left(1 - \frac{\epsilon_{f,0} \rho_{f,0}}{\epsilon_0 \rho_0} \right), \quad 1 \leq i \leq N_p, \quad 1 \leq j \leq N_g. \quad (9.3)$$

The inputs for the pyrolysis model in PATO are the γ_{ij} and the w_i . One can compute their value by summing over the species

$$\sum_{j=1}^{N_g} F_{ij} = \left(1 - \frac{\epsilon_{f,0} \rho_{f,0}}{\epsilon_0 \rho_0} \right) w_i \underbrace{\sum_{j=1}^{N_g} \gamma_{ij}}_{=1}, \quad (9.4)$$

and the relative weights w_i can be computed from

$$w_i = \frac{\sum_{j=1}^{N_g} F_{ij}}{\left(1 - \frac{\epsilon_{f,0} \rho_{f,0}}{\epsilon_0 \rho_0} \right)}. \quad (9.5)$$

Therefore, each F_{ij} is divided by a factor equal to approximately 0.43 which represents the ratio of the mass (density) of the initial resin with the total mass (density) of the sample, and the γ_{ij} can readily be retrieved through Eq. (9.3).

9.2.2 Pyrolysis gas chemistry

The chemistry of the gas phase is assumed to be at equilibrium, and only the mass fractions of the elements produced by the pyrolysis have to be provided, which are carbon, hydrogen and oxygen. The equilibrium chemistry assumption allows to reduce the number of independent variables that are solved by the numerical solver compared to a finite-rate chemistry model and reduces the computational time of the simulations. The open source library Mutation++ (Scoggins et al. [2020]), developed at the von Karman Institute for Fluid Dynamics, is used to compute thermodynamic and transport properties of the equilibrium gas from the local pressure, temperature and elemental composition.

In the mass conservation Eq. (2.35), the element production resulting from the pyrolysis decomposition appears as a source term. Often, material codes assign to each pyrolysis reaction

a constant pyrolysis gas composition with elemental mass fractions equal to 0.51, 0.14 and 0.36 for carbon, hydrogen and oxygen, respectively, resulting in a 50 % char residue after completion of the pyrolysis reactions (Sykes [1967]). In Section 7.4 we inferred on the species production and the elemental composition can be retrieved for each of the pyrolysis reaction, knowing the mass fractions and the composition of the pyrolysis gas from the reaction scheme in Eqs. (7.17)–(7.22). Indexing the elements by $k = 1, 2, 3$ for carbon, hydrogen and oxygen, respectively, the mass fraction for the production of the element $1 \leq k \leq N_g$ from the i -th pyrolysis reaction, $1 \leq i \leq N_p$, writes as

$$z_{ik} = \frac{\sum_{j=1}^{N_g} \gamma_{ij} Z_k(S_j) M_k}{\sum_{j=1}^{N_g} \gamma_{ij} Z_C(S_j) M_C + Z_H(S_j) M_H + Z_O(S_j) M_O}, \quad (9.6)$$

where $Z_k(S_j)$ is the number of atoms of the element k in the species S_j . The values that correspond to the MAP, further referred to as the nominal values, for the elemental pyrolysis gas composition produced by each of the pyrolysis reaction computed from Eq. (9.6) is summarized in Table 9.1. The input PDF for the elemental pyrolysis gas composition for each reaction is obtained by computing Eq. (9.6) for each of the sample drawn from the input PDF.

In the case of the six-reaction pyrolysis model, inferred from the experiments of Bessire and Minton [2017], the elemental compositions of the gas produced by each pyrolysis reaction (Table 9.1) are different. Comparatively, the global elemental composition of the overall pyrolysis reaction is equal here to 0.61, 0.07 and 0.34 for carbon, hydrogen and oxygen, respectively, hence, more hydrogen and less oxygen are produced by the six-reaction model. The char residue left from the pyrolyzing resin is equal to 40 %, which is less than what is assumed in current state-of-the-art models.

9.2.3 Mesh and issues regarding the in-depth resolution

The mesh, shown in Fig. 9.1, is a monolithic simplification of the tiled configuration of the manufactured heat shield. A comparison of a 3D tiled with a 3D monolithic heatshield model is performed in Meurisse et al. [2018], and showed differences for the in-depth temperature but the two model yielded very closed results for the surface temperature. The physical time interval takes place between 50 s and 100 s along the reentry path, during which the pressure peak and the peak heating on the heat shield occur. The aerothermal boundary conditions are provided from hypersonic CFD simulations, which are considered to be deterministic (Meurisse et al. [2018]).

Mesheres with 26k cells (15 cells in depth), 64k (37) and 103k (60) were investigated. Issues were observed for the six-reaction pyrolysis model for the 26k and 64k cells, as illustrated in Fig. 9.2, which shows the surface pyrolysis flux at the nose and the lee side. As it will be showed in the next section, the pyrolysis flux that is produced with the six-reaction pyrolysis model increases faster than with previous pyrolysis models, such as the two-equation model from Goldstein [1969]. During the first 5 s of the simulation, the pyrolysis production rises quickly and intensively, and the mesh with 26k and 64k cells does not provide a satisfactory solution, and refining the mesh allows to better capture these high fluxes. The mesh with 103k cells was found to provide sufficient accuracy, for which neither oscillations (as in Fig. 9.2(a)) nor rapid increases (as in Fig. 9.2(b)) in the fluxes were observed for the nominal (MAP) solution.

9.2.4 Comparison with a two-equation pyrolysis model

The comparison between the nominal temperature obtained with the six-reaction model and the two-reaction model from Goldstein [1969] (see Section 3.5.2) is shown in Fig. 9.3. The surface temperature at the nose and in-depth temperatures at the nose and the wind sides, whose positions are summarized in Table 9.2, are represented on the graph.




Figure 9.1: Illustration of the mesh for the MSL heat shield. View from the front (a) and the side (b).

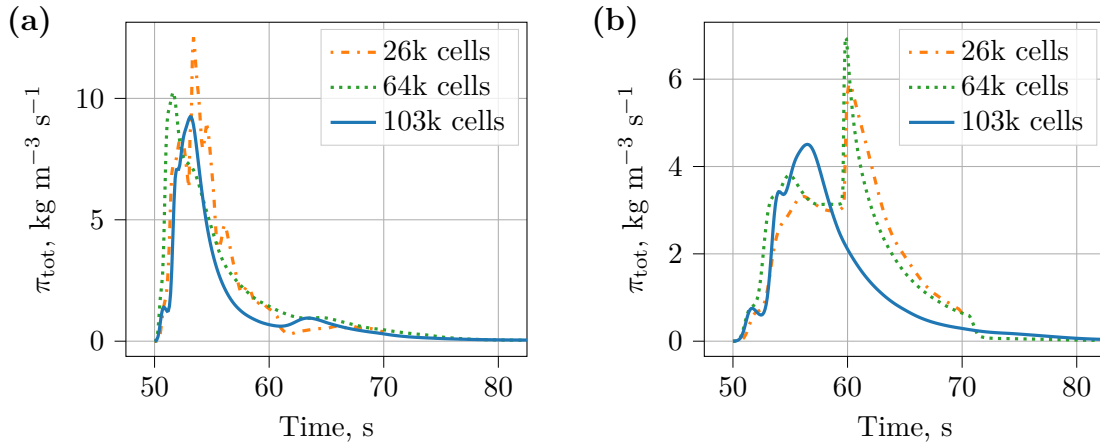


Figure 9.2: Total pyrolysis production rate with the six-reaction model for increasing in-depth mesh cells. (a) Nose, and (b) lee side.

Discrepancies between the results obtained with the two-reaction model and the six-reaction model can be observed. At the surface and for the two shallowest probes (probe 1 and probe 5), between 52 s and 64 s, the predicted temperature is higher with the two-reaction model, and between 64 s and 74 s it is lower than the temperature predicted with the six-reaction model. Deeper below the surface, the two-reaction model is predicting a higher temperature than the six-reaction model over the whole time domain.

In Fig. 9.4, we compare the nominal solutions for the pyrolysis production rate. At all probes, it is observed that the six-reaction model is producing pyrolysis gas much faster and more acutely than the two-reaction model. Pyrolysis reactions are endothermic, hence they absorb the heat and the temperature of the material increases more slowly than if there was no reaction. Therefore, because the pyrolysis reactions are slower with the two-reaction model than with the six-reaction model, the temperature reaches first higher values with the two-reaction model, as it was observed in Fig. 9.3. Then, close to the surface, the pyrolysis reactions reach completion, and the temperature increases faster with the six-reaction model, whereas they are just triggered with the two-reaction model, thus attenuating the increase in the temperature

Table 9.2: Heat shield embedded probe positions. Probes 1–4 are below the heat shield “nose” position and probes 5–8 are below the “wind side” (see later in Table 9.3).

Position	x (m)	y (m)	z (m)
Probe 1	0.00254	0.0	0.0
Probe 2	0.00508	0.0	0.0
Probe 3	0.01143	0.0	0.0
Probe 4	0.01778	0.0	0.0
Probe 5	0.726	-0.0025	-2.184
Probe 6	0.728	-0.0025	-2.183
Probe 7	0.734	-0.0025	-2.181
Probe 8	0.74	-0.0025	-2.179

curves.

Deeper below the surface, the heating rate is lower and the pyrolysis reactions are slower, characterized by production peaks that are wider and less intense. The temperature predicted by the two-reaction model is globally higher than with the six-reaction model, as seen in Fig. 9.3 at probes 2 and 6, and probes 3 and 7.

At probes 4 and 8, the two deepest, the temperature remains relatively low and the pyrolysis reactions has almost not started yet. The lower temperature observed with the six-reaction model here is due to the lower temperature that is conducted in-depth as a consequence from the above pyrolysis reactions. Hence, the six-reaction model predicts a lower subsurface temperature, with a relative temperature difference between the two models of $(568 - 536)/536 \times 100 \approx 6\%$ and $(580 - 552)/552 \times 100 \approx 5\%$ at probe 4 and probe 8, respectively, at $t = 100$ s.

To sum up, it seems that for the probes close to the surface, the temperature are approaching the same value at $t = 100$ s, no matter the pyrolysis model assumed. However, deeper below the surface, the pyrolysis reaction scheme can influence the sub-surface temperature.

9.3 Results of the propagation of uncertainty

From the Markov chains computed in Section 7.4.2 for each case of the free parameter ϑ , we select 500 evenly spaced samples to be propagated. Selecting all the samples would not be realistic in terms of computational time and would result in a poor convergence of the statistics of interests due to the correlation between the samples, inherent to the MCMC procedure. It is possible to compute an effective number of sample draws, see Gelman et al. [2014] p. 286 as discussed in Section 5.4.7, and it could be done with the results from the ISDE simulation.

To represent the uncertainty on the quantity of interest, 95 % probability regions (2.5 and 97.5 percentiles) are computed using the results of the 500 Monte Carlo samples. It is important to mention that for $\vartheta = 0.7n_t$, 14 samples out of the 500 led to a solution that features a similar issue as discussed in Section 9.2.3, and were not used for the estimation of the uncertainty ranges. These samples were identified based on the value of the advective pyrolysis heat flux that reach unacceptable values in such a case.

9.3.1 Temperature and gas elemental composition

For the results in the following section, the mesh with 103k cells is used. Each single PATO simulation runs on 34 cores from the NIC5 cluster¹ from the University of Liège, which features AMD EPYC Rome 7524 32-Core CPUs (2.9 GHz), and each simulation lasts around 1h20.

¹Provided by the Consortium des Équipements de Calcul Intensif (CÉCI).




Figure 9.3: Heat shield nose surface and in-depth temperature probes. Nominal solution for the six-equation pyrolysis model (continuous lines) compared to the model of Goldstein [1969] implemented in PATO (discontinuous lines). The surface probes at the nose are the top two lines with no markers.

Allocating more than 34 cores does not further improve significantly the computational time. In total, 1000 simulations are run in parallel on the cluster² (500 simulations for each value of the free parameter ϑ).

From the temperature plot in Fig. 9.5, the maximum of uncertainty is reached at the surface and near the peak heating (Fig. 9.5(a)). The uncertainty decreases just below the surface but then increases again with depth. For these in-depth temperatures, maximum uncertainty is also reached near the peak heating conditions for the shallowest probes (Fig. 9.5(b)), whereas the maximum uncertainty is reached at $t = 100$ s for the deepest probes (Fig. 9.5(c)). This was already observed in Chapter 8 and in Copeland et al. [2012].

Hence, the pyrolysis parameters have a relatively low impact on temperature, and the maximum is around 1 % at the surface. We noted little influence on the in-depth probes, with the maximum uncertainty at the peak temperature when it is reached, or at the end of the simulation for the deepest probes. However, one has to note the important effect of going from the two-reaction pyrolysis model to the six-reaction pyrolysis model as it was observed from Fig. 9.3, showing that the structure of the pyrolysis model can have a larger impact on temperature than the model parameter uncertainty.

Fig. 9.6 shows the elemental gas composition at the surface of the heat shield. The results of the simulations are probed at the four locations on the heat shield given in Table 9.3. Uncertainty ranges are shown for the two values of ϑ : the inner (darker) shaded areas are for $\vartheta = n_t$ and the outer (brighter) shaded areas are for $\vartheta = 0.7n_t$. A rapid change of the composition at the beginning can be observed, where the pyrolysis production is at his highest value and the most severe (see Fig. 9.4) at the beginning of the heating. After 85 s, where the heating and the

²As a comparison, these 1000 simulations would represent more than 5 years of computational time on a single CPU.




Figure 9.4: Heat shield nose surface and in-depth probes for the total pyrolysis production rate. The nominal solution for the six-equation pyrolysis model (continuous lines) is compared to the model of Goldstein [1969] implemented in PATO (discontinuous lines). The surface probes at the nose are the top two lines with no markers.

Table 9.3: Heat shield surface probe position.

Position	x (m)	y (m)	z (m)
Nose	0.0	0.0	0.0
Wind side	0.72346	-0.0025	-2.184
Lee side	0.72346	0.0025	2.184
Middle side	0.72346	2.184	0.0

pressure start to decrease, the oxygen concentration increases, and the hydrogen and carbon concentration decreases. The uncertainty ranges also increase, except at the nose which remains almost constant.

The numeric values for the median and the uncertainty ranges for the elemental composition at the nose are summarized in Table 9.4. Regarding the uncertainty ranges, we can see that there is a maximum of uncertainty around 50.4 s, very rapidly after the heating starts where intense gas production occurs. At this maximum, the ratio of the thickness of the 95 % probability region and the median value is around 63.25 %, 96.86 % and 33.37 % for carbon, hydrogen and oxygen, respectively. Around the peak heat, at $t = 75$ s, these values are around 4.88 %, 9.03 % and 7.86 %, respectively. The numeric values for $\vartheta = 0.7n_t$ are also shown, and it can be seen that it affects significantly the relative thickness of the probability ranges, as it ranges up to 78.19 %, 125.57 % and 41.06 % at 50.4 %. Tables for the uncertainty ranges at the wind side, lee side and middle side, where the pressure and heating are lower, are provided in Appendix D. At these other locations, the observations around the maximum uncertainty region and the peak heating are similar than at the nose. One can note further from the graph from Fig. 9.6 the uncertainty regions are however wider at $t = 100$ s, and the increase/decrease of the elemental composition

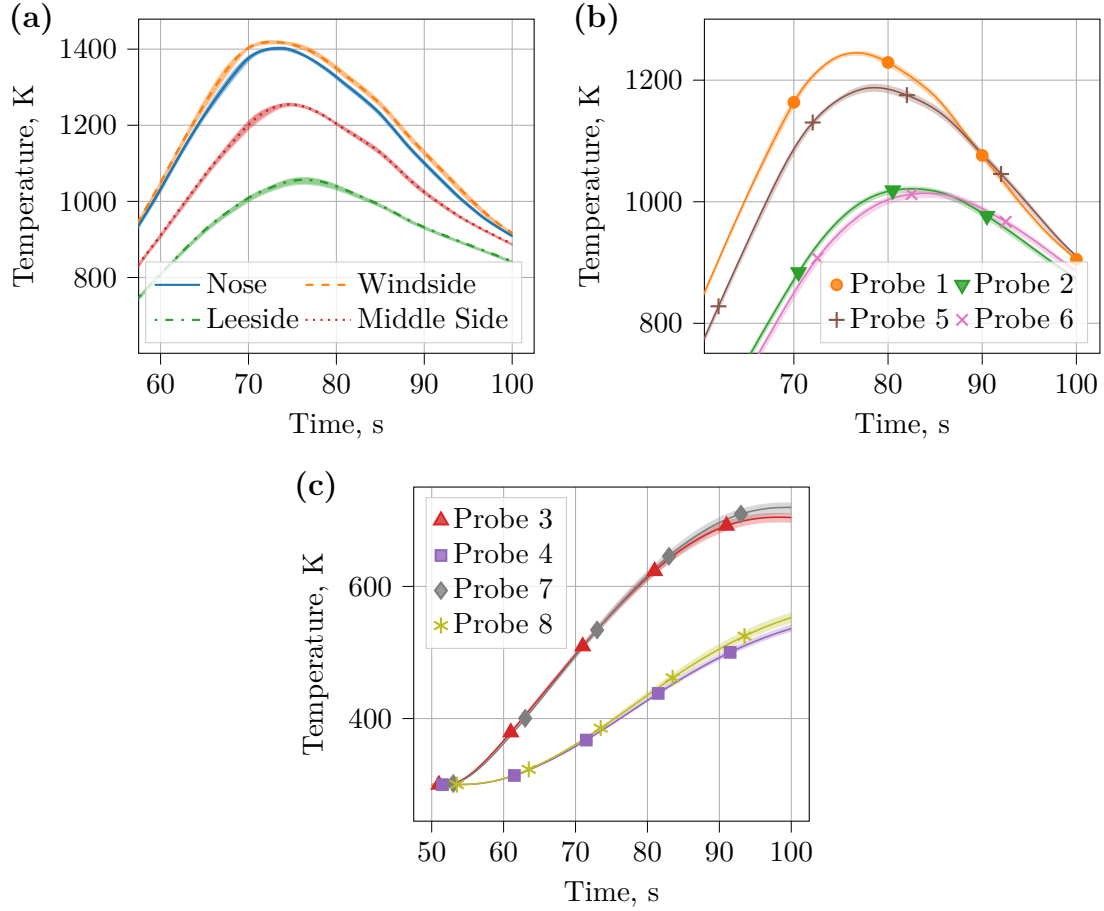


Figure 9.5: Heat shield embedded temperature probes with uncertainty ranges. (a) two shallowest probes at the nose and wind side, and (b) the two deepest probes at the nose and wind side.

after 85 s is less marked: median value remains almost constant, but some simulation results also predict the same behavior as at the nose, thus the larger uncertainty ranges.

9.3.2 Recession and pyrolysis production

The recession of the surface of the heat shield as a function of time at the probe locations is shown in Fig. 9.7. For $\vartheta = 0.7n_t$ (Fig. 9.7(a)), it shows an increase of the uncertainty as time increases. At 100 s after the beginning of the reentry, the median recession is around 0.922 mm at the nose region, 0.961 mm at the wind side, 0.531 mm at the middle side, and 0.195 mm at the lee side, and the results for the median are similar for $\vartheta = n_t$ (not showed on the graph for clarity). In Fig. 9.7(b), we provide the estimation of the probability density functions at $t = 100$ s using a Gaussian kernel-density estimation method for the two values of ϑ . The PDFs feature one mode and are almost centered around the median. For $\vartheta = 0.7n_t$, the supports of the PDFs are slightly wider, showing that the uncertainty is larger. Indeed, the thickness of the 95 % probability regions are around 0.038 mm (the ratio of this thickness and the median is equal to 4.14 %) and 0.034 mm (3.58 %) at the nose and the wind side, respectively, while for $\vartheta = n_t$ these values are equal to 0.03 mm (3.27 %) and 0.026 mm (2.7 %), respectively. At

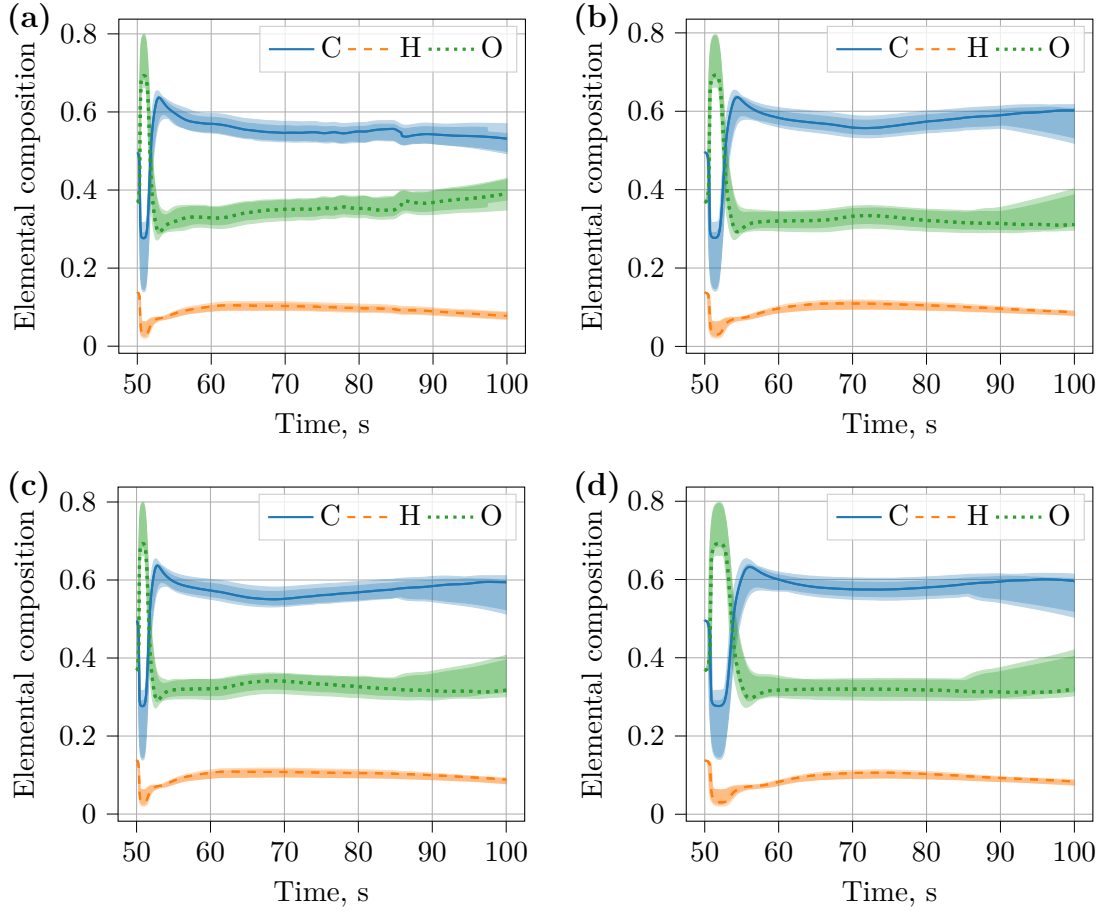


Figure 9.6: Results of the propagation of uncertainty for the gas elemental composition at the surface at the (a) nose, (b) middle side, (c) wind side (d) lee side. The continuous, discontinuous and dotted lines are the results of the simulation at the MAP for C, H and O, respectively, and the shaded areas represent the total uncertainty (95 % probability region) estimated from the Monte Carlo samples for the two values of ϑ .

the lee side, the uncertainty is relatively higher for both values of ϑ , which is around 0.019 mm (9.70 %) and 0.023 mm (11.86 %) for $\vartheta = n_t$ and $\vartheta = 0.7n_t$, respectively.

These results show that the value of the free parameter ϑ does not affect significantly the median and the mean of the recession of the surface of the heat shield. However, as expected, decreasing the value of ϑ increases the resulting uncertainty on the recession. In general, we note that the level of uncertainty on the recession due to pyrolysis model parameters is relatively low. Recession is related to the pyrolysis model parameters only through the effective material properties, themselves related to the reaction rates through the porosity of the material. Hence, other parameters like the intrinsic material thermal conductivity, the heat capacity, or the initial porosity, are foreseen to increase further the general uncertainty level on the recession.

The impact of pyrolysis model parameters can also be observed on other quantities. The results for the advective pyrolysis heat flux and the total pyrolysis flux at the surface of the heat shield are shown in Fig. 9.8. Advective pyrolysis heat flux is maximum at the nose and wind side, and the corresponding uncertainty is also maximum at those two locations before the peak heating (Fig. 9.8(a)). It is the opposite in Fig. 9.8(b) for the total pyrolysis production,

$\vartheta = n_t$					
50.4 s			75 s		
C	0.273	[0.228;0.401] (63.25 %)	0.546	[0.530;0.557] (4.88 %)	
H	0.065	[0.048;0.111] (96.86 %)	0.099	[0.094;0.103] (9.03 %)	
O	0.665	[0.489;0.711] (33.37 %)	0.355	[0.343;0.371] (7.86 %)	
$\vartheta = 0.7n_t$					
50.4 s			75 s		
C	0.275	[0.215;0.430] (78.19 %)	0.550	[0.530;0.566] (6.44 %)	
H	0.063	[0.040;0.119] (125.57 %)	0.099	[0.092;0.105] (13.19 %)	
O	0.667	[0.451;0.725] (41.06 %)	0.352	[0.335;0.373] (10.63 %)	

Table 9.4: Probabilistic predictions for the gas elemental composition at the surface of the heat shield (nose) for the two values of ϑ at $t = 50.4$ s (maximum uncertainty) and $t = 75$ s (around peak heating). The values are the median predictions with their 2.5–97.5 % probability intervals between squared brackets, and the ratio of the thickness of the probability interval and the median between parentheses.

for which the uncertainty ranges are the largest at the lee side and middle side, while they are thinner at the nose and wind side.

9.4 Conclusion and future work

In this chapter, we assessed the impact of the uncertainty in the pyrolysis model parameters, characterized by input PDFs inferred from ground-experiments, on the temperature, gas elemental composition and surface recession of an ablating TPS for a Mars atmospheric entry.

The surface gas elemental composition has the largest uncertainty. We noted the relatively low impact of pyrolysis model parameter uncertainty on temperature, with a maximum of 1 % variation at the surface, but we emphasized a more significant impact when going from the two-reaction pyrolysis model to the six-reaction pyrolysis model, with a relative temperature difference up to 5 %. Hence, the structure of the pyrolysis model can have a larger impact on temperature than the model parameter uncertainty.

The propagated uncertainty (2σ) on the recession is equal to 0.922 ± 0.019 mm (with 4.14 % the ratio of the thickness of the 95% probability interval and the median) at the nose and up to 11.85 %, 0.195 ± 0.023 mm, at the lee side, with the resulting distributions that are almost Gaussian. The comparison of the values obtained using either $\vartheta = n_t$ or $\vartheta = 0.7n_t$ shows that the resulting uncertainty range is about 1 % higher for the latter, thus providing more conservative margins on the recession.

It is a first step towards uncertainty quantified simulations for atmospheric entry. As only the parameters from the pyrolysis model were considered, the uncertainty provided represent a lower estimate of the uncertainty bounds, and should therefore be applied with care. As future work, more uncertain parameters, such as the material thermal conductivity, the heat capacity, or the porosity of the material, should be included in the analysis. It would also be desirable to perform a sensitivity analysis of the pyrolysis model parameters on the quantity of interest, in order to identify the most influential parameters and eventually reduce the number of uncertain input parameters for future predictions.

It is important to mention that the current pyrolysis model was calibrated at a single heating rate (366 K min^{-1}) and that more complex models could be developed in order to account for the various heating rates encountered over the heat shield (Torres-Herrador et al. [2019a]). The model error on the pyrolysis model at 366 K min^{-1} was partly included in the input parameters by relaxing the weights in the likelihood function, but the error due to the extrapolation of this

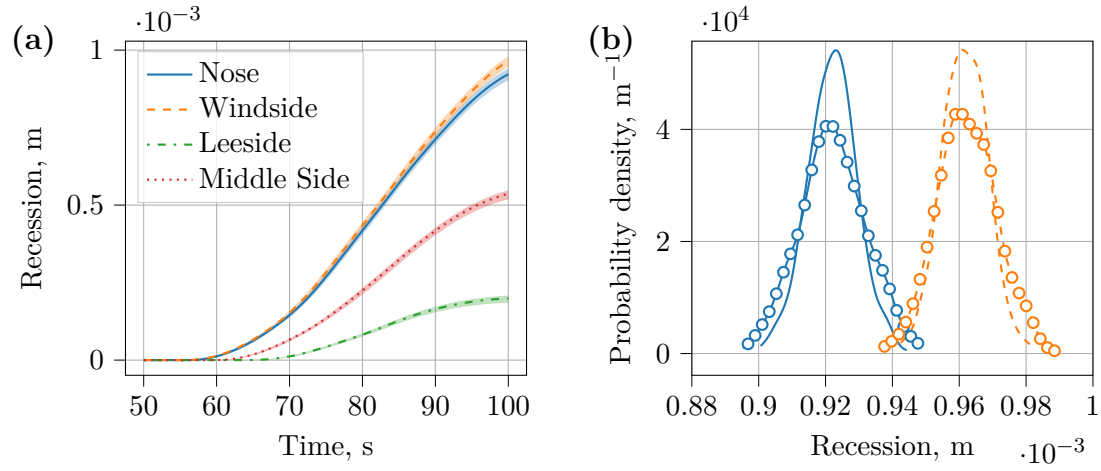


Figure 9.7: Results of the propagation of uncertainty in pyrolysis model parameters on the recession of the material at four different locations on the heat shield. (a) Recession as a function of time shown for $\vartheta = 0.7n_t$, and (b) estimation of the probability density functions at $t = 100$ s for the nose region (continuous line) and wind side (dashed line) for $\vartheta = n_t$ (no markers) and $\vartheta = 0.7n_t$ (circle markers).

model to other heating rates was not accounted for, which could also have an influence on the uncertainty on recession.

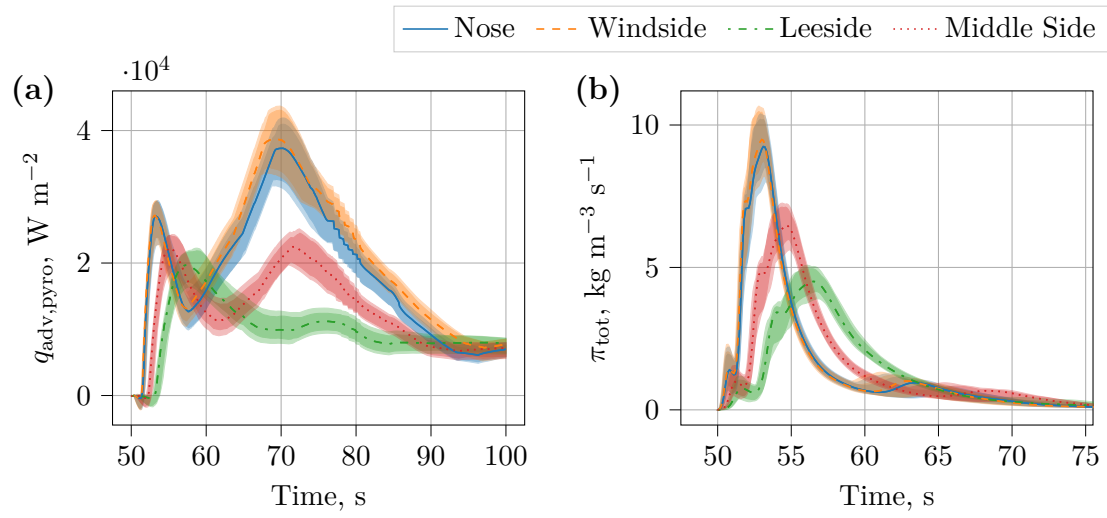


Figure 9.8: (a) Advective pyrolysis heat flux and (b) total pyrolysis production rate with uncertainty ranges for the two values of ϑ .

Part IV

Conclusions and directions for future work

Chapter 10

Conclusions and perspectives

10.1 Summary and conclusions

This thesis contributed to the development of an uncertainty-quantified model for pyrolysis-ablation of porous fiber/resin ablative materials. We presented the characterization of input parameter uncertainty using recent experimental data on the pyrolysis of thermal protection materials. We also presented the propagation of the uncertainty for an atmospheric entry problem by showing the response of a thermal protection system during a Mars entry and clarified the role played by the uncertainties.

The work was divided into three main parts, in which we contributed to the development of a high-fidelity deterministic numerical model for pyrolysis-ablation, reviewed methods for uncertainty quantification and developed a python code for performing Bayesian inference, and finally applied methods from uncertainty quantification to practical applications relevant to TPMs decomposition and atmospheric entry.

In Part I (Chapters 2–4), the modeling of atmospheric entry flows and pyrolysis-ablation was reviewed. In Chapter 2, we reviewed the VANS equations that can be used to simulate aerothermal flows through and around porous ablative TPMs.

In Chapter 3, we reviewed two important models for representing the resin pyrolysis of TPMs, namely the multi-component and competitive models. For multi-component models, we contributed to the derivation of an analytical solution which allowed us to save computational time for the calibration of pyrolysis experiments and to build a new detailed pyrolysis model.

In Chapter 4, the framework of the VANS equations within a unified flow-material approach was applied to simulate an ablation-pyrolysis experiment of a material subjected to a high-enthalpy flow. This method provides a high-order solution of the flow around and within the material, as well as a detailed description of the flow-material interaction at the interface. Yet being still computationally expensive, it is a very promising approach to study with a high fidelity the decomposition of ablative porous TPMs.

Part II (Chapters 5–6) was concerned with the theory of uncertainty quantification and its implementation in numerical codes. In Chapter 5, we reviewed the theory and methods for uncertainty quantification, with a particular focus on methods for Bayesian inference. We reviewed state-of-the-art MCMC algorithms, such as gradient-informed sampling algorithms, and their implementation.

In Chapter 6, we introduced a new software for performing Bayesian inference, and described the structure of the code, and showed its use with an example of parameter inference. This code

was developed for research purposes and for having a versatile toolbox in which new methods can be easily implemented thanks to the ease of use of the Python coding language.

Finally, Part III (Chapters 7–9) was concerned with the application of the uncertainty quantification methods to problems relevant to TPMs decomposition and atmospheric entry, and answered the key objectives of the thesis stated in the introduction (Section 1.5).

First, in Chapter 7, a Bayesian approach involving an approximate likelihood function was applied to infer the model parameters of chemical reactions relevant to pyrolysis using multi-component and competitive models. For multi-component models, we proposed to use gradient-informed sampling algorithms, along with a method for the adaptation of the covariance matrix and a reparameterization of the pyrolysis equations, in order to sample efficiently from the posterior distribution. The results showed that the adaptation of the covariance matrix based on the previous sample history improves the convergence rate of the sampling algorithms, even with a few samples, as compared to when there is no adaptation. A six-equation pyrolysis model was found by iteratively increasing the number of equations in the model until the experimental data were sufficiently accurately reproduced. The resulting uncertainty-quantified model allowed us to reproduce the data set together with the estimated observed uncertainty ranges and we showed that we were able to control the concentration of the posterior distribution by using an additional parameter introduced in the approximate likelihood. For the competitive models, we showed that this kind of model was able to simulate pyrolysis at two distinct heating rates and the model parameter uncertainty was also obtained. Yet, even with the competitive models, we showed that model errors were still present when extrapolating to other heating rates.

Then, in Chapter 8, we investigated the use of Bayesian methods for the analysis of the post-flight data of the MSL mission. Thermocouple data within the heat shield were simulated with a good accuracy and we were able to reconstruct the environment conditions at the surface. The results for the thermocouple highlighted a model discrepancy at moderately low temperature, that could be due to water vapor release. The comparison of the present results of the post-flight analysis and the results from CFD showed discrepancies between the predicted and the reconstructed environment, which could be due to non-predicted physical behaviors. However, the comparison with more CFD results at other MISP locations should be performed.

Finally, in Chapter 9 we assessed the impact of the uncertainty in the pyrolysis model parameters on the recession of the surface of an ablating TPS for a Mars atmospheric entry with the configuration of the Mars Science Laboratory mission heat shield. We showed that the surface gas elemental composition has the largest uncertainty. We compared the results for two values of the additional parameter in the approximate likelihood and we showed that more conservative margins could be assessed on the recession by providing a smaller value of the parameter. We also emphasized that the pyrolysis model structure could have a large impact on the temperature, whereas pyrolysis model parameters had a relatively low impact on the temperature.

10.2 Suggestions and perspectives for future work

This work is a first step towards uncertainty-quantified simulations for atmospheric entry. Throughout this manuscript, we discussed several improvements and directions for future research. Some of them are summarized and further discussed below, and we provide a more holistic viewpoint.

10.2.1 Pyrolysis model inference

In Chapter 7, we used a formulation of the likelihood function that involves a weighted distance, but which led to the introduction of a regularization of the weights at low and high temperatures. Alternative formulations of the likelihood function could be sought, such as formulations that use data transformation or that take into account the time-dependent nature of the data

by evaluating how well the model reproduces the observations at each individual step, see for instance in [Bennett et al. \[2013\]](#) for a review of different metrics.

We adopted an approach whereby we kept the structure of the pyrolysis model unchanged (we do not introduce additional stochastic terms) and we sought to incorporate a representation of the uncertainties into the model parameters. However, this approach of encapsulating all the uncertainties in the model parameters has its limitations. Indeed, the reference deterministic model needs to be close enough to the observations, otherwise the model discrepancy cannot be represented only by accommodating uncertainty in the model parameters and additional terms should be included in the statistical model ([Kennedy and O’Hagan \[2001\]](#)). Alternatively, pyrolysis models can still be improved. The use of competitive mechanisms in the pyrolysis scheme was already a first step towards a more accurate modeling, however, other unmodeled physical phenomena, such as outgassing at low temperatures or the plateau in the production curves, are not yet represented.

We presented the results for parameter inference for two different pyrolysis models. Building a suitable pyrolysis model (both for multi-component and competitive pyrolysis models) and obtaining an initial guess was found to be particularly difficult, and the automatization of this procedure would be valuable. For instance, for multi-component pyrolysis models, we emphasized the difficulty of finding an appropriate reaction pattern when accounting for the species production, due to the high numbers of different models that were providing satisfactory fitting to the data, hence justifying the need for an iterative procedure. In order to go towards a more automated procedure, different models could be evaluated from a set of plausible models while running, for instance, a genetic optimization (as in [Torres-Herrador et al. \[2019b\]](#)) for finding the best initial guess. The selection of the model could then be performed using methods for model comparison and model selection, for instance, using Akaike’s criterion for statistical model inference ([Burnham and Anderson \[2002\]](#)) or in a Bayesian framework, e.g., see in [Burnham and Anderson \[2002\]](#); [MacKay \[2003\]](#); [Link and Barker \[2006\]](#), or in [Strong and Oakley \[2014\]](#).

Generalization to other heating rates

The heating rates encountered during an actual atmospheric entry largely vary within the heat shield and the impact of using multi-component models at these multiple heating rates is not known. Implementing competitive pyrolysis models in high-fidelity numerical solvers would be a first step towards a more accurate modeling of resin decomposition at the heating rates encountered during an atmospheric entry. The samples inferred in this work using Bayesian inference for the competitive pyrolysis model could be readily propagated and the results compared with the current state-of-the-art models as well as with the results of Chapter 9.

Yet, we emphasized that extrapolating the models at heating rates outside of the calibration range could still display significant model errors. Thus, more elaborated competitive models valid in a wider heating rate range should also be investigated.

10.2.2 Inferring the input parameter distributions

During pre-flight analyses, there is a need to characterize the model parameter distributions. In this work, we focused mostly on the characterization of uncertainty from pyrolysis model parameters. However, more uncertain parameters, such as the material thermal conductivity, the heat capacity, or the porosity of the material, should be included in the analysis, as they can also have a non-negligible impact on the quantities of interest ([Rivier et al. \[2019\]](#)). They can be inferred from ground-based experiments either by direct measurements, or indirect measurements. For the latter, a similar Bayesian probabilistic inverse analysis as the one applied in this work could be performed.

Material properties are not the only source of uncertainties, and uncertainties related to the atmospheric entry environment or the flight parameters should also be characterized. These parameters affect the whole external aerothermal flow conditions and a coupled flow-material

approach or a unified approach would be required in order to study the material response. When simulating the external flow, the computational cost can become expensive and methods to reduce the cost of uncertainty quantification analyses will be needed, or by reducing the number of uncertain parameters. This is discussed below.

10.2.3 Directions in uncertainty propagation

In Chapter 9, the direct Monte Carlo method was applied, thanks to the model in PATO that was relatively fast to compute ($\approx 1\text{h}20\text{min}$). For finer meshes, or when the evaluations of the computational model are time consuming, such as for instance in Argo, the convergence rate of the statistical descriptors may be too slow and become prohibitive, and more efficient methods for uncertainty propagation should be investigated. In order to reduce the computational cost and improve the accuracy of the estimation of the statistical descriptors, improvements of the standard Monte Carlo method or surrogate-based methods could be sought.

Advanced sampling methods can be used to reduce the computational cost of the standard Monte Carlo method by carefully selecting the samples to be propagated such that it improves the accuracy of estimated statistical descriptors. These methods are for instance the quasi-Monte Carlo method (Dick et al. [2013]) which is based on deterministic sequences of numbers rather than (pseudo-)random sequences of numbers, or the Latin hypercube sampling (Loh [1996]).

Multilevel Monte Carlo methods (see a review in Giles [2015]) are a more recent extension of the Monte Carlo methods that are used when the solution of the numerical model can be computed on different level of accuracy. It can offer improved computational efficiency by smartly computing the solution at sample values on different resolutions, such as for instance when a hierarchy of computational grids are available. Most simulations are performed at a low accuracy with a lower computational cost and only a few simulations are performed at a high accuracy and a high cost. See for instance in Mishra et al. [2012] or Pisaroni et al. [2017] for recent advances and applications to CFD problems.

Surrogate-based propagation methods, such as the PCE methods, can be used to replace the expensive computational model, but they suffer from the curse of dimensionality and finding quadrature rules for general probability distribution is not trivial. Firstly, to mitigate the curse of dimensionality, the dimension of the input probability distribution could be reduced by identifying non-influential parameters of the distribution by performing a sensitivity analysis of the pyrolysis model parameters on the quantity of interest. Secondly, concerning the PCE methods, they are well-suited for standard distributions with independent variables, such that the polynomial bases that are optimal for multivariate distributions can be built as the tensor product of univariate orthogonal polynomials that are optimal for the corresponding univariate standard distributions (Xiu and Karniadakis [2002]). When considering non-standard distributions with potentially dependent stochastic variables, one can still use polynomial bases that are optimal for the standard distributions, but the accuracy of the statistical descriptors can be degraded (Jakeman et al. [2019]). In order to build polynomials that are optimal for non-standard distributions, one possibility is to use the arbitrary polynomial chaos (aPC), also called data-driven polynomial chaos, which is an extension of PCE to correlated random variables with non-standard distributions or unknown distributions, as for instance when the samples are provided from a Bayesian inference analysis (Witteveen and Bijl [2006]). In the aPC approach, the expansion is not built using conventional polynomials obtained from the Askey scheme but are constructed numerically. The aPC method is used in Jakeman et al. [2019] in a non-intrusive way, and in Paulson et al. [2017] using a stochastic Galerkin projection. It was investigated in Lacroix [2020] for a probability distribution for pyrolysis model parameters obtained from Bayesian inference with moderately low dimensions, and encouraging results in this direction were showed.

10.2.4 Validation of the models

The comparison of the numerical predictions with the experimental observations will clarify the extent to which high-fidelity numerical simulations of TPS ablation are able to reproduce experimentally the observed behaviors, whereas our uncertainty bounds will allow to specifically assess the extent to which discrepancies can be explained as a consequence of uncertainty in the physico-chemical models. Any unexplained discrepancies will provide useful guidance to future efforts to improve the physico-chemical models and the computational methods as well as to design experimental setups.

Within Argo, a comparison of a Plasmatron experiment from [Helber et al. \[2016b\]](#) was already compared with the numerical simulations in [Schrooyen et al. \[2017\]](#) for a carbon pre-form material. A comparison with a problem relevant to pyrolysis-ablation is still missing, and the comparison with the experiments would require to have material properties for pyrolyzing materials, which are still lacking from the literature. The Zoram database, currently under development ([Rothermel et al. \[2014\]](#)), is a first step towards the development of such detailed material properties.

Besides the replication of ground-test experiments, analyzing the post-flight data as in Chapter 8 will help to refine the models used in the simulations for predicting the atmospheric entry for future missions. The comparison of the reconstructed environment from post-flight data with more CFD results will clarify the validity of the models currently used in state-of-the-art computer models, will further enrich these models and allow to have a quantification of the error, which will be useful for future uncertainty-quantified pre-flight analyses. This will ultimately improve our confidence in the models used, hence reducing the uncertainty margins and benefiting the design of TPSs for future space missions.

Part V

Appendices

Appendix A

Analytical solutions and link between competitive and multi-component parallel pyrolysis models

A.1 Computation of derivatives for multi-component parallel models

The computation of derivatives is also required for gradient-informed algorithms to compute the solution of inverse problems. Derivatives are expressed here as a function of temperature and for the linearly increasing temperature case as it will be used in an illustrative example in the next section. The chain rule provides the following results:

$$\frac{\partial \nu_j}{\partial \mathcal{A}_i} = m_0 F_{ij} \left((1 - \alpha_i)^{n_i - 1} \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} \left(n_i \mathcal{A}_i \frac{-\partial \alpha_i}{\partial \mathcal{A}_i} + (1 - \alpha_i) \right) \right), \quad (\text{A.1})$$

$$\begin{aligned} \frac{\partial \alpha_i}{\partial \mathcal{A}_i} = & - \left((1 - n_i) \left(-\frac{\mathcal{A}_i T}{\beta} \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} - \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} + C_i \right) \right)^{\frac{n_i}{1-n_i}} \dots \\ & \times \left(-\frac{T}{\beta} \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} - \frac{\mathcal{E}_i}{\mathcal{R}T} \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} + \frac{\partial C_i}{\partial \mathcal{A}_i} \right), \end{aligned} \quad (\text{A.2})$$

$$\frac{\partial C_i}{\partial \mathcal{A}_i} = \frac{T_0}{\beta} \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} + \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\} \frac{\mathcal{E}_i}{\beta \mathcal{R}}, \quad (\text{A.3})$$

$$\frac{\partial \nu_j}{\partial \mathcal{E}_i} = m_0 F_{ij} \left((1 - \alpha_i)^{n_i - 1} \mathcal{A}_i \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} \left(n_i \frac{-\partial \alpha_i}{\partial \mathcal{E}_i} - \frac{(1 - \alpha_i)}{\mathcal{R}T} \right) \right), \quad (\text{A.4})$$

$$\begin{aligned} \frac{\partial \alpha_i}{\partial \mathcal{E}_i} = & - \left((1 - n_i) \left(-\frac{\mathcal{A}_i}{\beta} T \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} - \frac{\mathcal{A}_i}{\beta} \frac{\mathcal{E}_i}{\mathcal{R}} \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} + C_i \right) \right)^{\frac{n_i}{1-n_i}} \dots \\ & \times \left(-\frac{\mathcal{A}_i}{\beta \mathcal{R}} \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} + \frac{\partial C_i}{\partial \mathcal{E}_i} \right), \end{aligned} \quad (\text{A.5})$$

$$\frac{\partial C_i}{\partial \mathcal{E}_i} = \frac{\mathcal{A}_i}{\beta \mathcal{R}} \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T_0} \right\}, \quad (\text{A.6})$$

$$\begin{aligned} \frac{\partial \nu_j}{\partial n_i} = & m_0 F_{ij} A_i \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} (1 - \alpha_i)^{n_i-1} \dots \\ & \times \left(n_i \frac{\partial \alpha_i}{\partial n_i} + (1 - \alpha_i) \log(1 - \alpha_i) \right), \end{aligned} \quad (\text{A.7})$$

$$\begin{aligned} \frac{\partial \alpha_i}{\partial n_i} = & - \left(((1 - n_i)(B_i + C_i))^{\frac{1}{1-n_i}} \dots \right. \\ & \times \left. \left(\frac{-B_i + (1 - n_i) \frac{\partial C_i}{\partial n_i} - C_i}{(1 - n_i)^2 (B_i + C_i)} + \frac{\log \{(1 - n_i)(B_i + C_i)\}}{(1 - n_i)^2} \right) \right), \end{aligned} \quad (\text{A.8})$$

$$\frac{\partial C_i}{\partial n_i} = - \frac{(1 - \alpha_0)^{1-n_i}}{(1 - n_i)^2} - \frac{(1 - \alpha_0)^{1-n_i} \log(1 - \alpha_0)}{1 - n_i}, \quad (\text{A.9})$$

$$B_i = - \frac{\mathcal{A}_i}{\beta} \left(T \exp \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} + \frac{\mathcal{E}_i}{\mathcal{R}} \text{Ei} \left\{ \frac{-\mathcal{E}_i}{\mathcal{R}T} \right\} \right). \quad (\text{A.10})$$

Note that in Eq. (A.7), the last term on the right-hand side may pose some numerical issues because of the indetermination. This is the case, for instance, after sufficiently large time after the reaction is completed, or when the different reactants have distant activation energies leading to the reaction i being completed much earlier than other ones. In both cases, $\alpha_i \rightarrow 1$ and we have to resolve $(1 - \alpha_i) \log(1 - \alpha_i)$ for which the limit $0 \times \infty$ is of indeterminate form. This limit can be determined using l'Hôpital's rule for computing the limit of the ratio of two functions and is equal to 0. Practically speaking, in our numerical implementation, every time one of the advancement of reaction factor α_i reaches a value of $\varepsilon_\alpha = 1 - 10^5$, we set to zero the quantity $(1 - \alpha_i) \log(1 - \alpha_i)$.

Derivatives with respect to F_{ij} are trivial. Issues regarding the computation of the solution when $\alpha_i \rightarrow 1$ in $(1 - \alpha_i)^{n_i-1}$ with $0 < n_i < 1$ that may arise are avoided because we assumed that $n_i > 1$. Finally, because the pyrolysis reactions are considered to be independent (parallel), we have $\partial \pi / \partial \mathcal{A}_i = \partial \pi_i / \partial \mathcal{A}_i = \partial(\nu_j(1 - \gamma_{C,i}) / \gamma_{ij}) / \partial \mathcal{A}_i$ (same for \mathcal{E}_i, n_i, w_i).

A.2 Discussion on the existence of a solution for generalized competitive mechanisms

The system Eq. (3.44) is linear and the coefficients will be generally non-constant due to the temperature dependence of the reaction rates. The solution is of the form

$$\mathbf{m}(t) = e^{\int_{t_0}^t L(\tau) d\tau} \mathbf{m}_0 \quad (\text{A.11})$$

which can be checked by differentiating it. When the coefficient matrix L is constant (isothermal pyrolysis) the matrix-exponential method can be used and the solution of the system writes

$$\mathbf{m}(t) = e^{L(t-t_0)} \mathbf{m}_0 \quad (\text{A.12})$$

where the exponential of a given square matrix X is defined as

$$\exp(X) = \sum_{k=0}^{\infty} \frac{1}{k!} X^k. \quad (\text{A.13})$$

This solution can be used to simulate isothermal pyrolysis or a when the temperature is a piecewise constant function as in [Wong et al. \[2015b\]](#), with initial conditions $\mathbf{m}_0 = (m_{\text{init}}, 0, \dots, 0)^T$.

For a non-constant coefficient matrix L , the system of ODEs from Eq. (3.44) has to be integrated numerically. Indeed, in order to have a solution of the type A.11, we need to prove that L commutes with its integral $I_L = \int_{t_0}^t L(t)dt$ for the exponential matrix to be valid ([Nemytskii and Stepanov \[1960\]](#)). Two matrices X and Y commute if $XY = YX$. There is not a condition or criterion for a matrix to commute with its integral, although we can cite the work of [Epstein \[1963\]](#). We demonstrate here using a counterexample that the non constant coefficient matrix does not commute with its integral and so that a solution of the type Eq. A.11 cannot be found. We consider the simple case of a two-step reaction mechanism $\rho_1 \rightarrow \rho_2 \rightarrow \rho_3$. The non-singular 2×2 coefficient matrix is $L = \begin{pmatrix} -k_1 & 0 \\ k_1 & -k_2 \end{pmatrix}$ and $\mathbf{m} = (m_1, m_2)^T$. It is worth noting that only the activated variables m_1 and m_2 need to be considered for the resolution of the system, as m_3 is a linear combination of the previous variables (the physical meaning is that m_3 can be obtained using mass conservation of the two reactants m_1 and m_2); considering ρ_3 in the system of ODEs would make the coefficient matrix singular. The two matrix products are

$$L I_L = \begin{pmatrix} k_1 I_1 & 0 \\ -k_1 I_1 - k_2 I_1 & k_2 I_2 \end{pmatrix}, \quad I_L L = \begin{pmatrix} k_1 I_1 & 0 \\ -k_1 I_1 - k_1 I_2 & k_2 I_2 \end{pmatrix}. \quad (\text{A.14})$$

The condition for L to commute with I_L is $k_1 I_2 = k_2 I_1$. The result for L constant is immediately verified as we have $k_1 k_2 = k_2 k_1$. For the non constant case, we can show that the condition is $\exp(-\mathcal{E}_1/(\mathcal{R}T)) \text{Ei}(-\mathcal{E}_2/(\mathcal{R}T)) = \exp(\mathcal{E}_2/(\mathcal{R}T)) \text{Ei}(-\mathcal{E}_1/(\mathcal{R}T))$ (where Ei is the exponential integral function) which is not verified. As a conclusion, when considering multi-step non isotherm reactions, an analytical solution of the form A.11 will generally not be found.

A.3 From generalized competitive models to multi-component models

A.3.1 Linear multi-component, multi-step competitive reactions

We consider the most general case of multi-component multi-steps competitive reactions. For this general case, we can again use the matrix formalism of Eq. 3.38 where now

$$L = \begin{pmatrix} L_1 & 0 & \cdots & 0 \\ 0 & L_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & L_n \end{pmatrix} \quad (\text{A.15})$$

is a block diagonal matrix. Each block L_i is a lower triangular square matrix of different dimension n_i representing a particular component of the pyrolysis mechanism. The null entries represent matrices which dimensions are consistent with the block diagonal. Each block is decoupled and can be solved separately. The generalized initial conditions are now

$$\mathbf{m}_0 = (\underbrace{m_{1,\text{init}}, 0, \dots, 0}_{n_1}, \underbrace{m_{2,\text{init}}, 0, \dots, 0}_{n_2}, \dots, \underbrace{m_{n,\text{init}}, 0, \dots, 0}_{n_n})^T. \quad (\text{A.16})$$

The previous case of pyrolysis reactions for aerospace applications was in fact considering 1×1 matrices $L_i = k_i$ with $\mathbf{m}_0 = (m_{1,\text{init}}, m_{2,\text{init}}, \dots)^T$ and is therefore a particular case of this

general formulation. For this case, because L is diagonal it commutes with its integral and it is possible to find an analytical solution. It was also possible to find a closed form of the solution for the non-linear system of PDEs, both for L constant and for linear heating rates.

In the case of time dependent heating rates (linear, modulated, ...) with multi-step pyrolysis, or for non-linear reaction orders ($n \neq 1$), the numerical treatment of the system of ODEs is required.

A.3.2 Link with multi-component parallel models

Except from the non-linearity in the system of ordinary differential equations induced by the reaction order that is not accounted for in the competitive mechanisms presented here, the multi-component model presented in Section 3.3 can be seen as a particular case of this general model assuming that:

1. there are no competitive reactions ($N_{ci} = 1$, $1 \leq i \leq N$), which simplifies Eq. 3.43 into

$$\frac{dm_i}{dt}(t) = -k_i m_i(t) + \sum_{l=1}^N k_l \gamma_{l,i} m_l(t), \quad (\text{A.17})$$

2. each fictitious component produces only gas and no other solid component ($\gamma_{l,i} = 0 \forall l \in I_s$).
3. the gaseous products do not deposit ($\gamma_{l,i} = 0 \forall i \in I_s$) and there are no secondary reactions (frozen chemistry) among gaseous species ($k_i = 0 \forall i \in I_g$).

In matrix form, Eq. A.17, together with hypotheses 2 and 3, leads to the following system of ODEs for the solid species:

$$\frac{d\mathbf{m}_s}{dt}(t) = -K \mathbf{m}_s(t) \quad (\text{A.18})$$

where we denoted $\mathbf{m}_s = \{m_i, i \in I_s\}$ and $K = \text{diag}(k_1, k_2, \dots, k_{N_s})$. The matrix K is diagonal because the $m_{s,i}$ are independent from each other, following from hypothesis 2, and each ordinary differential equation for $i \in I_s$ can be solved independently. Once we have solved for the \mathbf{m}_s , the solution for the gaseous product then follows from the right-hand side of Eq. A.17 together with hypothesis 3 and we have

$$\frac{d\mathbf{m}_g}{dt}(t) = \Gamma^T \mathbf{m}_s, \quad \text{with } \Gamma = \begin{bmatrix} k_1 \gamma_{11} & k_1 \gamma_{12} & \cdots & k_1 \gamma_{1N_g} \\ k_2 \gamma_{21} & k_2 \gamma_{22} & \cdots & k_2 \gamma_{2N_g} \\ \vdots & \vdots & \ddots & \vdots \\ k_{N_s} \gamma_{N_s 1} & k_{N_s} \gamma_{N_s 2} & \cdots & k_{N_s} \gamma_{N_s N_g} \end{bmatrix}, \quad (\text{A.19})$$

where Γ is the $N_s \times N_g$ matrix of coefficients for the production of gaseous species. Matrix Γ completely defines the finite rate chemistry produced by the decomposition of the solid material and the elements of Γ are unknown parameters.

Appendix B

Additional implementation details

B.1 Effect of the reparametrization on the Metropolis–Hastings algorithm

The change of parameter is denoted $\tilde{\mathbf{p}} = s(\mathbf{p})$ and the inverse change is denoted by s^{-1} . Here, a given iteration in the Metropolis–Hastings algorithm is indexed by i and the proposed parameter value has the superscript $*$. The acceptance ratio is

$$r = \frac{\pi_0(\tilde{\mathbf{p}}^*)\pi(d|s^{-1}(\tilde{\mathbf{p}}^*))J(\tilde{\mathbf{p}}^*|\tilde{\mathbf{p}}_{i-1})}{\pi_0(\tilde{\mathbf{p}}_{i-1})\pi(d|s^{-1}(\tilde{\mathbf{p}}_{i-1}))J(\tilde{\mathbf{p}}_{i-1}|\tilde{\mathbf{p}}^*)}. \quad (\text{B.1})$$

The new proposal function J is defined on the \tilde{p} parameter space. The likelihood function is not affected by the change of parameter space since the code output is unaffected. The prior density needs however to be recomputed. Because the probability density is conserved, we can write for a general density of probability π

$$\pi_{\tilde{p}}(\tilde{p})d\tilde{p} = \pi_p(p)dp \quad (\text{B.2})$$

$$\pi_{\tilde{p}}(\tilde{p}) = \pi_p(p) \det \left(\frac{dp}{d\tilde{p}} \right). \quad (\text{B.3})$$

In terms of the transformed parameter \tilde{p} , we can write

$$\pi_{\tilde{p}}(\tilde{p}) = \pi_p(s^{-1}(\tilde{p})) \det \left(\frac{d(s^{-1}(\tilde{p}))}{d\tilde{p}} \right). \quad (\text{B.4})$$

For the random-walk Metropolis–Hastings algorithm (the random walk being performed now on the \tilde{p} space) and assuming a uniform π_p prior distribution, the acceptance ratio is therefore

$$r = \frac{\det \left(\frac{d(s^{-1}(\tilde{\mathbf{p}}))}{d\tilde{\mathbf{p}}} \Big|_{\tilde{\mathbf{p}}=\tilde{\mathbf{p}}^*} \right) \pi(d|s^{-1}(\tilde{\mathbf{p}}^*))}{\det \left(\frac{d(s^{-1}(\tilde{\mathbf{p}}))}{d\tilde{\mathbf{p}}} \Big|_{\tilde{\mathbf{p}}=\tilde{\mathbf{p}}_{i-1}} \right) \pi(d|s^{-1}(\tilde{\mathbf{p}}_{i-1}))}. \quad (\text{B.5})$$

For linear transformations, we can see directly that the acceptance ratio is unaffected by the reparametrization. However, in case of non-linear transformations, the determinant of the Jacobian must be computed explicitly. For instance, the non-linear transformation introduced in Section 7.2.3 for the pyrolysis equations, the determinant of the Jacobian writes as

$$\det = \prod_{i=1}^N \exp(\tilde{A}_i + \tilde{E}_i \bar{X}_i) = \prod_{i=1}^N A_i, \quad (\text{B.6})$$

where N is the number of reaction rates in the model.

Appendix C

Link with Bayesian statistics, additional graph of convergence for activation energy and table of input parameter values

C.1 Motivation for the expression of the approximate likelihood function

C.1.1 Link with Bayesian statistics

Equation 7.5 can be seen as an application of the Bayes' theorem to the calibration of model parameters, where π^0 and π are interpreted as probability density functions. The prior PDF π^0 is updated through the likelihood function to obtain the posterior PDF on the parameters $\pi(\cdot|\mathbf{y}^{\text{obs}})$. The function from Eq. (7.6) for the prior is improper in the sense that its value does not integrate to 1 and we will check numerically that the resulting posterior has a finite integral Gelman et al. [2014]. The likelihood function from Eq. (7.8) can also find an interpretation from a statistical point of view. Assuming that the discrepancy between the model predictions and the observations is due only to experimental errors and that these experimental errors (represented by the ε in Eq. (5.28)) are statistically independent and Gaussian with mean zero and variance $\left(s_{ij}^{(e)}\right)^2$, the likelihood is the multivariate (non-normalized) Gaussian distribution from

Eq. (7.8) with entries $w_{jji}^{-1} = \left(s_{ij}^{(e)}\right)^2 / n_r$ since

$$\prod_{i=1}^{n_{\text{obs}}} \prod_{j=1}^{n_t} \prod_{k=1}^{n_r} \exp \left(-\frac{1}{2} \left(\frac{y_{ijk}^{\text{obs}} - \eta_{ij}(\mathbf{m})}{s_{ij}^{(e)}} \right)^2 \right) \propto \exp \left(-\frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \sum_{j=1}^{n_t} \left(\frac{\frac{1}{n_r} \sum_{k=1}^{n_r} y_{ijk}^{\text{obs}} - \eta_{ij}(\mathbf{m})}{s_{ij}^{(e)} / \sqrt{n_r}} \right)^2 \right), \quad (\text{C.1})$$

and provides the probability density of observing the quantity \mathbf{y} given the parameters \mathbf{m} . However, for the case considered here, the discrepancy between the model predictions and the observations is not only due to experimental errors, but model inadequacy and parametric uncertainty are also present. Thus, to account for model inadequacy and parametric uncertainty, we allow w_{jji}^{-1} to be different from and greater than $\left(s_{ij}^{(e)}\right)^2 / n_r$.

The use of diagonal matrices $W_{::i}$ may not be fully justified within a Bayesian statistics perspective. Most likely, full matrices should be used and they could be obtained either from the

experiments or from the inference process together with the parameter distribution. The inverse of the weight matrices can be estimated, for instance, by computing the covariance matrices for the data. However, in general the covariance matrices computed from the experiments are singular when only few realizations are available, and inferring on large non-singular covariance matrices is complex and beyond the scope of this work.

One way to simplify the covariance matrix and make it non-singular is to reduce the correlation between the data such that the covariance matrix becomes diagonal. As off-diagonal entries depend on the acquisition frequency of the data, we could reduce the correlation by subsampling the initial measurement time series. However, selecting the subsampling ratio that results in a sufficient decorrelation without significant loss of information about the physical process is not trivial. Indeed, reducing too much the data set could lead to a loss of information that would result in an increase of the parameter uncertainty.

A case where reducing the data set is feasible without losing significant information is when important characteristics that summarize the physics of the model can be identified within the data set. Navarro et al. [2018] have been able to pursue such a strategy and identify three important characteristics in an application relevant to debris flow, namely, arrival time, location of the maximum height and decay behavior. However, finding important characteristics is challenging in the case of pyrolysis reactions because it is more difficult to identify common characteristics to all curves; from the pyrolysis experiments (see Fig. 3.3), it can be seen that some species have characteristics that are not common to all curves: CO_2 has two local maxima, or CO has a plateau above 1100 K. Reducing the data set may also lead to the presence of undesirable modes in the solution when parameters from nonlinear models are inferred. Using all the time measurements increases the chance to suppress the modes in the inferred distribution that are not present in the parameter distribution.

Therefore, instead of trying to reduce the data set, we will use all the data we have from the experiments. To take into account the dependence of the covariance matrix on the data, we use a likelihood function that will depend on the data acquisition frequency as discussed in the next section.

Finally, from a statistical point of view, Eq. (C.1) can be seen as the product of Gaussian functions centered at the different time observations averaged by the number of observations n_t . From that point of view, the Gaussian assumption might also be questionable as the data may not be distributed according to a Gaussian function at every time step. This assumption is difficult to verify based only on a small number of experimental repetitions and is thus not addressed in this work. However with more data points, the Gaussian assumption could be verified using, for instance, a kernel density estimation at each time observations.

C.1.2 Likelihood function for time series measurements

Because the weighting matrices $W_{:,i}$ are assumed to be diagonal, the likelihood function from Eq. (7.8) does not take into account the dependence structure. The concentration of the posterior depends on the data acquisition frequency and there may be a risk for the posterior to have a region of high density over a small domain if there are too many terms in the second sum and the weights are too large.

To choose the values of the w_{jji} , we introduced two normalizations. The first normalization takes into account the observed dispersion by choosing the weights w_{jji} inversely proportional to s_{ij}^2 . The second normalization takes into account the number of data acquired in the time series n_t . The motivation for this second normalization is discussed below.

Instead of starting from the discretized problem and the weighted norm from Eq. (7.7), we start from the continuous dynamical model and we define the norm for a time-dependent function that we will eventually discretize. Given a strictly positive and bounded weight function v , we

define the weighted norm of a square integrable time-dependent function $u : [0, t_f] \rightarrow \mathbb{R}$

$$\|u\|_v := \left(\int_0^{t_f} v(t) (u(t))^2 dt \right)^{1/2}, \quad (\text{C.2})$$

where t_f denotes the total time duration of the process. Defining a potential function $U_i(\mathbf{m})$ as the square of Eq. (C.2), with u equal to the difference between the continuous observations $\bar{y}_i^{\text{obs}}(\cdot)$ and the continuous model output $\eta_i(\cdot, \mathbf{m})$ for observation i , we have

$$U_i(\mathbf{m}) = \frac{1}{t_f} \int_0^{t_f} v_i(t) (\bar{y}_i^{\text{obs}}(t) - \eta_i(t, \mathbf{m}))^2 dt, \quad (\text{C.3})$$

where the ratio $1/t_f$ is used to normalize the potential over the integration domain. Next, we discretize the observations $\bar{y}_i^{\text{obs}}(\cdot)$ and we approximate the integral in Eq. (C.3) by a finite sum of n_t rectangular boxes of widths Δt_j that coincide with the times between the different experimental observations. Using this discretization in the potential in Eq. (7.7) results in

$$\phi_i(\bar{\mathbf{y}}_{i:}^{\text{obs}}, \mathbf{g}_{i:}(\mathbf{m})) = \sum_{j=1}^{n_t} \frac{\Delta t_j}{t_f} v_{ij} (\bar{y}_{ij}^{\text{obs}} - \eta_{ij}(\mathbf{m}))^2, \quad (\text{C.4})$$

and we end up with the following likelihood function for the n_r experimental runs

$$\begin{aligned} \exp \left(-\Phi(\mathbf{y}^{\text{obs}}, G(\mathbf{m})) \right) &= \exp \left(-\frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \phi_i(\bar{\mathbf{y}}_{i:}^{\text{obs}}, \boldsymbol{\eta}_{i:}(\mathbf{m})) \right) \\ &= \exp \left(-\frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \sum_{j=1}^{n_t} \frac{\Delta t_j}{t_f} v_{ij} (\bar{y}_{ij}^{\text{obs}} - \eta_{ij}(\mathbf{m}))^2 \right). \end{aligned} \quad (\text{C.5})$$

This last expression for the likelihood function is similar to Eq. (7.8) with $w_{ij} = v_{ij}(\Delta t_j/t_f)$. The factors Δt_j depend on the acquisition frequency of the experimental data that fixes the successive values of Δt_j , and, for a constant acquisition frequency, the ratio $\Delta t_j/t_f = 1/n_t$ is constant. In the limit where we have only two time instants in the time series ($t = 0$ and $t = t_f$), then $\Delta t_f/t_f = 1$ and the likelihood thus obtained with the single data at $j = 1$ is proportional to the evaluation of a Gaussian with mean $\bar{\mathbf{y}}_{:1}^{\text{obs}}$ and diagonal covariance matrix with diagonal entries (the data at $t = 0$ is the initial condition that provides no information).

C.1.3 Interpretation using the notion of push-forward measure

We wish to have a posterior that converges to a unique posterior when increasing the data acquisition frequency and that is representative of the variability of the experimental observations. For a given model, increasing the acquisition frequency should have no influence on the posterior. For instance, two contiguous observations separated by a very short time interval are probably providing the same information, and the resulting posterior using the two observations instead of a single one should not become thinner. The presence of the factor $\Delta t_j/t_f$ in Eq. (7.9) serves to ensure this behavior.

The presence of the factors $\Delta t_j/t_f$ and ϑ in Eq. (7.9) can also be linked with the concept of push-forward measure (Butler et al. [2018]), that is, a posterior for which the result of the forward propagation through the model matches the data-generating probability distribution. This concept differs from a purely statistical interpretation of the Bayesian formula, according to which collecting more data provides more confidence.

C.2 Influence of the initial covariance matrix

In Fig. C.1, we compare the effect of the initial estimations of the covariance matrix and its adaptation on the sampling of the bivariate posterior for \mathcal{A} and \mathcal{E} in the case of the one-reaction pyrolysis model and using the ISDE algorithm. Using the identity matrix and the inverse of the diagonal of the Hessian does not provide a satisfactory sampling of the posterior (top figures) as the sampling remains in the region of high probability. The adaptation of the covariance can be observed on the bottom left picture where we see a higher density of samples in the high probability region close to the starting value and is due to the time for the algorithm to learn from the scale and correlations. In this case, the exploration of the posterior is more efficient and has a comparable efficiency to when using the full Hessian matrix (bottom right). The first samples drawn in the high probability region could be discarded.




Figure C.1: Comparison of the posterior sampling with the different covariance matrix estimation for the ISDE algorithm. (a) identity covariance matrix, (b) ISDE with diagonal Hessian and no adaptation, (c) ISDE diagonal Hessian with adaptation, (d) ISDE with full Hessian. In total, 1000 samples are shown on each graph for the sake of clarity and emphasizing the adaptation of the covariance matrix. They were obtained from a single run of 10000 samples from which we took the first 2000 samples every 2 iterations. The sample values are compared to the direct evaluation of the posterior distribution rendered in the background. The “+” marker symbolizes the starting value.

It is not shown here, but it is worth noting that in the case where we apply first the change of variables based on the structure of the pyrolysis model, it induces a transformation and a rescaling of the posterior distribution. If the scaling factors appearing in the change of variable are chosen carefully, the tuning of the covariance matrix is made easier as it is close to the identity matrix. Therefore, because the change of variables proposed here is appropriate for the pyrolysis equations, the RWMH can be very efficient with a proposal covariance matrix that requires a minimum tuning.

The use of gradient-informed sampling algorithms was motivated by the high number of scaling factors that are unknown in high dimensions. Moreover, a change of variables well-suited to the equations considered is not always available. Nevertheless, even gradient-informed algorithms still require to be scaled according to the directions of the posterior distribution that are not known a priori. The results in this section show that the covariance matrix estimate is essential. Further improvements to the adaptation of the covariance matrix estimate for gradient-informed algorithms could be performed by considering locally adaptive covariance matrices based on a state-dependent scaling, as already suggested in [Rosenthal \[2010\]](#).

C.3 Graph of convergence for \mathcal{E}

Additional graphs for the convergence of the parameter \mathcal{E} are provided in Figs. C.2–C.3.

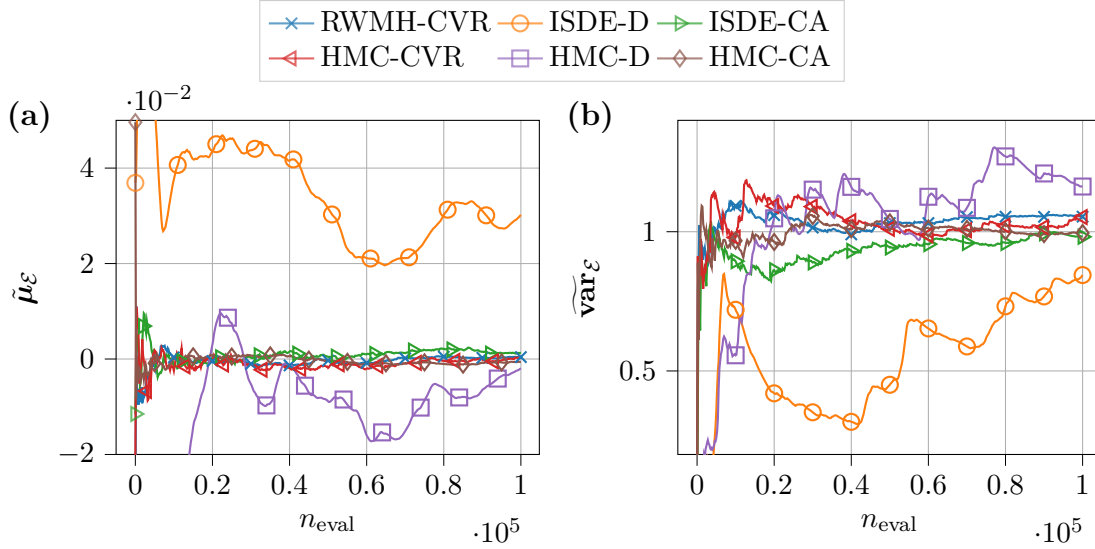


Figure C.2: Comparison of the convergence of the mean (a) and variance (b) of \mathcal{E} as a function of the number of function evaluations (after the burn-in phase) for the different sampling methods. Random-Walk Metropolis–Hastings (RWMH); RWMH with the change of variable based on the structure of the pyrolysis equations (RWMH-CVR); ISDE with the covariance matrix estimated from the diagonal of the inverse of the Hessian and change of variable based on the support of the posterior (ISDE-D); ISDE-D with the covariance matrix resulting from the adaptation process (ISDE-CA); HMC with the change of variable based on the structure of the pyrolysis equations (HMC-CVR); HMC with the covariance matrix estimated from the diagonal of the inverse of the Hessian and change of variables based on the support of the posterior (HMC-D); HMC-D with the covariance matrix resulting from the adaptation process (HMC-CA).

C.4 Parameter value table

Table C.1 provides the values for the six-equation pyrolysis model at the MAP (maximum a posteriori), which was approximated by the maximum value of the posterior reached during the MCMC iterations, for $\vartheta = n_t$ (left) and for $\vartheta = 0.7n_t$ (right). \mathcal{A} is in s^{-1} , \mathcal{E} is in J mol^{-1} , r and F are dimensionless.

	$\vartheta = n_t$	$\vartheta = 0.7n_t$
\mathcal{A}_0	$2.39130813546179 \cdot 10^8$	$3.33948906150348 \cdot 10^8$
\mathcal{E}_0	$7.64474934 \cdot 10^4$	$7.72205470 \cdot 10^4$
r_0	9.37226700	8.494309
F_{02}	$1.1380 \cdot 10^{-3}$	$2.9000 \cdot 10^{-4}$
F_{03}	$7.5300 \cdot 10^{-4}$	$7.1300 \cdot 10^{-4}$
F_{05}	$2.0350 \cdot 10^{-3}$	$1.5040 \cdot 10^{-3}$
\mathcal{A}_1	$3.928878043671898926 \cdot 10^{12}$	$5.033745123431458984 \cdot 10^{12}$
\mathcal{E}_1	$1.80048760 \cdot 10^5$	$1.81634729 \cdot 10^5$
r_1	5.499556	5.756187
F_{11}	$3.6940 \cdot 10^{-3}$	$4.7540 \cdot 10^{-3}$
F_{12}	$1.1827 \cdot 10^{-2}$	$1.1795 \cdot 10^{-2}$
F_{15}	$4.2520 \cdot 10^{-3}$	$4.2790 \cdot 10^{-3}$
F_{16}	$4.8600 \cdot 10^{-4}$	$5.0700 \cdot 10^{-4}$
F_{110}	$4.7200 \cdot 10^{-4}$	$4.5400 \cdot 10^{-4}$
F_{111}	$7.4400 \cdot 10^{-4}$	$7.3300 \cdot 10^{-4}$
\mathcal{A}_2	$4.703951292689136 \cdot 10^{15}$	$1.9854747464107464 \cdot 10^{16}$
\mathcal{E}_2	$1.82589107 \cdot 10^5$	$1.89029514 \cdot 10^5$
r_2	6.019494	7.154706
F_{21}	$5.2500 \cdot 10^{-4}$	$6.3700 \cdot 10^{-4}$
F_{22}	$1.9490 \cdot 10^{-3}$	$2.4920 \cdot 10^{-3}$
F_{23}	$3.4400 \cdot 10^{-4}$	$3.6900 \cdot 10^{-4}$
F_{24}	$1.0790 \cdot 10^{-3}$	$1.0210 \cdot 10^{-3}$
F_{25}	$3.4730 \cdot 10^{-3}$	$3.6410 \cdot 10^{-3}$
F_{26}	$2.0800 \cdot 10^{-4}$	$2.1200 \cdot 10^{-4}$
F_{27}	$1.5230 \cdot 10^{-3}$	$1.7790 \cdot 10^{-3}$
F_{28}	$1.3110 \cdot 10^{-3}$	$1.8200 \cdot 10^{-3}$
F_{29}	$4.1700 \cdot 10^{-4}$	$4.6200 \cdot 10^{-4}$
F_{211}	$1.3300 \cdot 10^{-4}$	$7.7000 \cdot 10^{-5}$
\mathcal{A}_3	$5.3413100200432137 \cdot 10^{10}$	$1.88989329806231567 \cdot 10^{11}$
\mathcal{E}_3	$1.73244332 \cdot 10^5$	$1.79307529 \cdot 10^5$
r_3	6.987248	7.905416
F_{30}	$2.4960 \cdot 10^{-3}$	$2.8290 \cdot 10^{-3}$
\mathcal{A}_4	$1.421255610884815104 \cdot 10^{18}$	$1.0922370953018785792 \cdot 10^{19}$
\mathcal{E}_4	$2.35776070 \cdot 10^5$	$2.45429327 \cdot 10^5$
r_4	6.970952	7.76469
F_{43}	$7.9000 \cdot 10^{-4}$	$8.7000 \cdot 10^{-4}$
F_{44}	$2.2850 \cdot 10^{-3}$	$2.7990 \cdot 10^{-3}$
F_{45}	$3.3010 \cdot 10^{-3}$	$3.6330 \cdot 10^{-3}$
F_{47}	$3.3010 \cdot 10^{-3}$	$3.0980 \cdot 10^{-3}$
F_{48}	$1.6810 \cdot 10^{-3}$	$1.5940 \cdot 10^{-3}$
F_{49}	$1.5800 \cdot 10^{-4}$	$1.2600 \cdot 10^{-4}$
\mathcal{A}_5	$8.81607232326933760 \cdot 10^{17}$	$1.911837989333075968 \cdot 10^{18}$
\mathcal{E}_5	$4.06141069 \cdot 10^5$	$4.07562253 \cdot 10^5$
r_5	9.071686	8.263785
F_{50}	$8.1300 \cdot 10^{-4}$	$7.5800 \cdot 10^{-4}$
F_{52}	$3.2850 \cdot 10^{-3}$	$2.8690 \cdot 10^{-3}$
F_{53}	$2.2000 \cdot 10^{-4}$	$2.5600 \cdot 10^{-4}$
F_{55}	$1.5020 \cdot 10^{-3}$	$1.1920 \cdot 10^{-3}$

Table C.1




Figure C.3: Comparison of the convergence of the mean (a) and variance (b) of \mathcal{E} for different values of n_f . The number of iteration in the adaptation phase are compared for the HMC and ISDE algorithms.

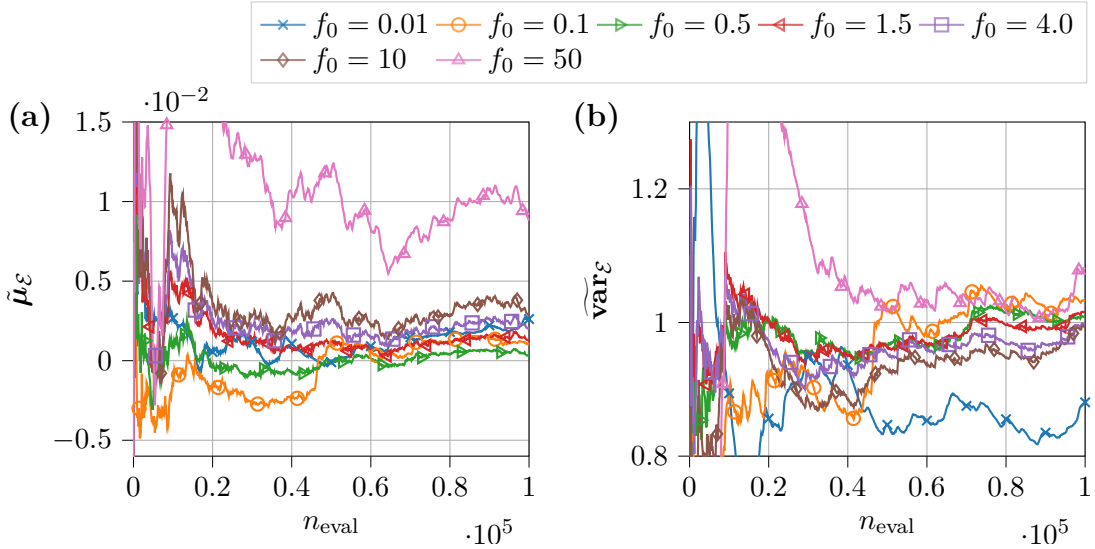


Figure C.4: Comparison of the convergence of the mean (a) and variance (b) of \mathcal{E} using the ISDE algorithm as a function of the number of function evaluations with different values for the damping factor f_0 .

Appendix D

Tables for the probabilistic predictions for gas elemental composition

This Appendix provides additional tables of the probabilistic predictions of the elemental gas composition on the heat shield surface at the wind side, lee side, and middle side. See the results in Section 9.3.1.

$\vartheta = n_t$			
	50.3 s		75 s
C	0.412	[0.279;0.461] (44.27 %)	0.564 [0.551;0.575] (4.29 %)
H	0.113	[0.072;0.127] (48.84 %)	0.104 [0.099;0.109] (8.92 %)
O	0.476	[0.413;0.650] (49.93 %)	0.332 [0.321;0.346] (7.52 %)
$\vartheta = 0.7n_t$			
	50.3 s		75 s
C	0.398	[0.261;0.467] (51.63 %)	0.568 [0.550;0.584] (5.95 %)
H	0.108	[0.057;0.129] (66.44 %)	75.000: 0.103 [0.097;0.111] (13.64 %)
O	0.495	[0.405;0.675] (54.60 %)	75.000: 0.329 [0.313;0.348] (10.53 %)

Table D.1: Probabilistic predictions for the gas elemental composition at the surface of the heat shield (wind side) for the two values of ϑ at $t = 50.3$ s (maximum uncertainty) and $t = 75$ s (around peak heating). The values are the median predictions with their 2.5–97.5 % probability intervals between squared brackets, and the ratio of the thickness of the probability interval and the median between parentheses.

$\vartheta = n_t$					
	50.7 s			75 s	
C	0.420	[0.282;0.470]	(45.02 %)	0.575	[0.563;0.588] (4.37 %)
H	0.072	[0.056;0.119]	(87.37 %)	0.102	[0.098;0.107] (9.15 %)
O	0.465	[0.399;0.648]	(53.40 %)	0.322	[0.311;0.336] (7.87 %)
$\vartheta = 0.7n_t$					
	50.7 s			75 s	
C	0.287	[0.233;0.447]	(74.23 %)	0.579	[0.561;0.596] (6.09 %)
H	0.071	[0.047;0.124]	(108.85 %)	0.102	[0.096;0.109] (12.99 %)
O	0.481	[0.394;0.669]	(57.22 %)	0.319	[0.302;0.338] (11.44 %)

Table D.2: Probabilistic predictions for the gas elemental composition at the surface of the heat shield (lee side) for the two values of ϑ at $t = 50.4$ s (maximum uncertainty) and $t = 75$ s (around peak heating). The values are the median predictions with their 2.5–97.5 % probability intervals between squared brackets, and the ratio of the thickness of the probability interval and the median between parentheses.

$\vartheta = n_t$					
	50.6 s			75 s	
C	0.347	[0.259;0.456]	(56.65 %)	0.565	[0.551;0.576] (4.43 %)
H	0.092	[0.061;0.126]	(71.09 %)	0.099	[0.094;0.103] (9.03 %)
O	0.560	[0.418;0.675]	(45.94 %)	0.330	[0.318;0.345] (8.09 %)
$\vartheta = 0.7n_t$					
	50.6 s			75 s	
C	0.327	[0.243;0.464]	(67.46 %)	0.569	[0.550;0.585] (6.07 %)
H	0.084	[0.050;0.128]	(93.65 %)	0.104	[0.098;0.112] (13.71 %)
O	0.593	[0.408;0.692]	(47.96 %)	0.327	[0.310;0.346] (11.01 %)

Table D.3: Probabilistic predictions for the gas elemental composition at the surface of the heat shield (middle side) for the two values of ϑ at $t = 50.6$ s (maximum uncertainty) and $t = 75$ s (around peak heating). The values are the median predictions with their 2.5–97.5 % probability intervals between squared brackets, and the ratio of the thickness of the probability interval and the median between parentheses.

Bibliography

- Abramowitz, M. and Stegun, I. (1972). *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*. Applied mathematics series. U.S. Department of Commerce, National Bureau of Standards, URL <https://books.google.be/books?id=Cxsty7Np9sUC>.
- Adams, B., Bauman, L., Bohnhoff, W., Dalbey, K., Ebeida, M., Eddy, J., Eldred, M., Hough, P., Hu, K., Jakeman, J., Stephens, J., Swiler, L., Vigil, D., , and Wildey, T. (2018). *Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.10 User's Manual*, sandia technical report sand2014-4633, july 2014. updated may 2019 (version 6.10) edition.
- Anderson, J. D. (1990). *Fundamentals of Aerodynamics*. McGraw-Hill, 2nd edition.
- Anderson, J. D. (2006). *Hypersonic and High Temperature Gas Dynamics*. AIAA Education Series, 2nd edition. 3, 4, 25, 28
- Andrieu, C. and Moulines, É. (2006). On the ergodicity properties of some adaptive MCMC algorithms. *The Annals of Applied Probability*, 16(3), 1462–1505, <https://doi.org/10.1214/105051606000000286>.
- Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing*, 18, 343–374, <https://doi.org/10.1007/s11222-008-9110-y>.
- Anzalone, R., Barr, B. W., Upadhyay, R. R., and Ezekoye, O. A. (2016). Use of a quasi-steady ablation model for design sensitivity with uncertainty propagation. *Journal of Thermal Science and Engineering Applications*, 9(1), <https://doi.org/10.1115/1.4034595>.
- Arnst, M., Álvarez, B. A., Ponthot, J.-P., and Boman, R. (2017). Itô-SDE MCMC method for Bayesian characterization of errors associated with data limitations in stochastic expansion methods for uncertainty quantification. *Journal of Computational Physics*, 349, 59–79, <https://doi.org/http://dx.doi.org/10.1016/j.jcp.2017.08.005>.
- Arnst, M. and Ponthot, J.-P. (2014). An overview of nonintrusive characterization, propagation, and sensitivity analysis of uncertainties in computational mechanics. *International Journal for Uncertainty Quantification*, 4(5), 387–421, <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2014006990>. 9, 79
- Arnst, M. and Soize, C. (2019). Identification and sampling of Bayesian posteriors of high-dimensional symmetric positive-definite matrices for data-driven updating of computational models. *Computer Methods in Applied Mechanics and Engineering*, 352, 300–323, <https://doi.org/10.1016/j.cma.2019.04.025>.

- Askey, R. and Wilson, J. (1985). Some basic hypergeometric orthogonal polynomials that generalize Jacobi polynomials. *Memoirs of the American Mathematical Society*, 54(319), <https://doi.org/10.1090/memo/0319>.
- Atchadé, Y. F. (2006). An adaptive version for the Metropolis adjusted Langevin algorithm with a truncated drift. *Methodology and Computing in Applied Probability*, 8(2), 235–254, <https://doi.org/10.1007/s11009-006-8550-0>.
- Atchadé, Y. F. and Rosenthal, J. S. (2005). On adaptive Markov chain Monte Carlo algorithms. *Bernoulli*, 11(5), 815–828, <https://doi.org/10.3150/bj/1130077595>.
- Bányai, T., Torres, E., Magin, T. E., Kashkovsky, A. V., Vashchenkov, P. V., Ivanov, M. S., and Rambaud, P. (2013). Evaluation of the aerodynamic properties of the Intermediate Experimental Vehicle in the rarefied and transitional regime. In P. Reijasse, D. Knight, M. Ivanov, and I. Lipatov (Eds.), *Progress in Flight Physics*: EDP Sciences. <https://doi.org/10.1051/eucass/201305425>.
- Baudin, M., Dutfoy, A., Iooss, B., and Popelin, A.-L. (2015). OpenTURNS: An industrial software for uncertainty quantification in simulation. In *Handbook of Uncertainty Quantification*, (pp. 1–38). Springer International Publishing, https://doi.org/10.1007/978-3-319-11259-6_64-1.
- Beck, J. L. and Katafygiotis, L. S. (1998). Updating models and their uncertainties. I: Bayesian statistical framework. *Journal of Engineering Mechanics*, 124(4), 455–461, [https://doi.org/10.1061/\(asce\)0733-9399\(1998\)124:4\(455\)](https://doi.org/10.1061/(asce)0733-9399(1998)124:4(455)).
- Bennett, N. D., Croke, B. F., Guariso, G., Guillaume, J. H., Hamilton, S. H., Jakeman, A. J., Marsili-Libelli, S., Newham, L. T., Norton, J. P., Perrin, C., Pierce, S. A., Robson, B., Seppelt, R., Voinov, A. A., Fath, B. D., and Andreassian, V. (2013). Characterising performance of environmental models. *Environmental Modelling & Software*, 40, 1–20, <https://doi.org/10.1016/j.envsoft.2012.09.011>.
- Bernard, R., Pichon, T., and Valverde, J. (2019). From IXV to Space Rider CMC thermal protection system evolutions. In *Proceedings of the 8th European Conference for Aeronautics and Space Sciences* Madrid, Spain. <https://doi.org/10.13009/EUCASS2019-991.6>
- Beskos, A., Jasra, A., Law, K., Tempone, R., and Zhou, Y. (2017). Multilevel sequential Monte Carlo samplers. *Stochastic Processes and their Applications*, 127(5), 1417–1440, <https://doi.org/10.1016/j.spa.2016.08.004>.
- Bessire, B. K., Lahankar, S. A., and Minton, T. K. (2014). Pyrolysis of phenolic impregnated carbon ablator (PICA). *ACS Applied Materials and Interfaces*, 7(3), 1383–1395, <https://doi.org/10.1021/am507816f>. 7, 38
- Bessire, B. K. and Minton, T. K. (2017). Decomposition of Phenolic Impregnated Carbon Ablator (PICA) as a function of temperature and heating rate. *ACS Applied Materials & Interfaces*, 9(25), 21422–21437, <https://doi.org/10.1021/acsami.7b03919>. xvii, 7, 11, 12, 36, 39, 40, 44, 50, 117, 118, 123, 130, 133, 136, 137, 139, 141, 142, 144, 145, 173
- Bhatia, A. and Roy, S. (2010). Modeling the motion of pyrolysis gas through charring ablating material using discontinuous galerkin finite elements. In *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2010-982>.
- Bianchi, D., Turchi, A., Nasuti, F., and Onofri, M. (2012). Coupled CFD analysis of thermochemical erosion and unsteady heat conduction in solid rocket nozzles. In *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit* (pp. 1–15). Atlanta, GA: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2012-4318>. 10, 24

- Bianchi, D., Turchi, A., Nasuti, F., and Onofri, M. (2013). Chemical erosion of carbon-phenolic rocket nozzles with finite-rate surface chemistry. *Journal of Propulsion and Power*, 29(5), 1220–1230, <https://doi.org/10.2514/1.b34791>. 10, 24
- Billingsley (2012). *Probability and Measure*. John Wiley & Sons.
- Bird, G. A. (1994). *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*. Oxford: Oxford University Press.
- Blondeau, J. (2013). *Investigation of pulverised biomass combustion: detailed modelling of particle pyrolysis and experimental analysis of ash deposition*. PhD thesis, Université catholique de Louvain.
- Blondeau, J. and Jeanmart, H. (2012). Biomass pyrolysis at high temperatures: Prediction of gaseous species yields from an anisotropic particle. *Biomass and Bioenergy*, 41, 107–121, <https://doi.org/10.1016/j.biombioe.2012.02.016>.
- Borner, A., Panerai, F., and Mansour, N. N. (2017). High temperature permeability of fibrous materials using direct simulation Monte Carlo. *International Journal of Heat and Mass Transfer*, 106, 1318–1326, <https://doi.org/10.1016/j.ijheatmasstransfer.2016.10.113>.
- Bosco, A. (2019). Bayesian inference for the identification of model parameters in atmospheric entry problems. Master’s thesis, University of Liège.
- Bose, D., White, T., Mahzari, M., and Edquist, K. (2014). Reconstruction of aerothermal environment and heat shield response of Mars Science Laboratory. *Journal of Spacecraft and Rockets*, 51(4), 1174–1184, <https://doi.org/10.2514/1.a32783>.
- Bousquet, P., Yamade, T., Munk, M., and Weihs, H. (2015). *Future robotic and human missions with opportunities for European technologies*. Technical report, International RE-Entry demoNstrator Action (IRENA) Project. 3
- Brandis, A. M., Saunders, D. A., Johnston, C. O., Cruden, B. A., and White, T. R. (2020). Radiative heating on the after-body of martian entry vehicles. *Journal of Thermophysics and Heat Transfer*, 34(1), 66–77, <https://doi.org/10.2514/1.t5613>. 8
- Brooks, S., Gelman, A., Jones, G. L., and Meng, X.-L. (2011). *Handbook of Markov Chain Monte Carlo*. London: Chapman & Hall/CRC.
- Brown, M. E., Maciejewski, M., Vyazovkin, S., Nomen, R., Sempere, J., Burnham, A., Opfermann, J., Strey, R., Anderson, H. L., Kemmler, A., Keuleers, R., Janssens, J., Desseyn, H. O., Li, C.-R., Tang, T. B., Roduit, B., Málek, J., and Mitsuhashi, T. (2000). Computational aspects of kinetic analysis. Part A: The ICTAC kinetics project-data, methods and results. *Thermochimica Acta*, 355, 125–143.
- Bruns, M. C. (2015). Inferring and propagating kinetic parameter uncertainty for condensed phase burning models. *Fire Technology*, 52(1), 93–120, <https://doi.org/10.1007/s10694-015-0457-2>. 12, 35, 117, 118
- Burnham, K. P. and Anderson, D. R. (2002). *Model selection and multimodel inference : a practical information-theoretic approach*. New-York: Springer, 2nd ed. edition.
- Burrage, K., Lenane, I., and Lythe, G. (2007). Numerical methods for second-order stochastic differential equations. *SIAM Journal on Scientific Computing*, 29(1), 245–264, <https://doi.org/10.1137/050646032>.
- Bussing, T. R. A. and Eberhardt, S. (1989). Chemistry associated with hypersonic vehicles. *Journal of Thermophysics and Heat Transfer*, 3(3), 245–253, <https://doi.org/10.2514/3.28772>.

- Butler, T., Jakeman, J., and Wildey, T. (2018). Combining push-forward measures and bayes rule to construct consistent solutions to stochastic inverse problems. *SIAM Journal on Scientific Computing*, 40(2), A984–A1011, <https://doi.org/10.1137/16m1087229>. 201
- Cao, Y., Li, S., Petzold, L., and Serban, R. (2003). Adjoint sensitivity analysis for differential-algebraic equations: The adjoint DAE system and its numerical solution. *SIAM Journal on Scientific Computing*, 24(3), 1076–1089, <https://doi.org/10.1137/S1064827501380630>.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), <https://doi.org/10.18637/jss.v076.i01>.
- Carton de Wiart, C. (2014). *Towards a discontinuous Galerkin solver for scale-resolving simulations of moderate Reynolds number flows, and application to industrial cases*. PhD thesis, Université Catholique de Louvain.
- Chen, M. (2001). Two-dimensional implicit thermal response and ablation program for charring materials. *Journal of Spacecraft and Rockets*, 38(4), 473–481, <https://doi.org/10.2514/2.3724>.
- Chen, Y.-K. and Gökçen, T. (2014). Implicit coupling approach for simulation of charring carbon ablators. *Journal of Spacecraft and Rockets*, 51(3), 779–788, <https://doi.org/10.2514/1.a32753>. 10, 24
- Chen, Y.-K. and Milos, F. S. (1999). Ablation and thermal response program for spacecraft heatshield analysis. *Journal of Spacecraft and Rockets*, 36(3), 475–483. 7, 24, 151
- Chen, Y.-K. and Milos, F. S. (2005). Navier-Stokes solutions with finite rate ablation for planetary mission Earth reentries. *Journal of Spacecraft and Rockets*, 42(6), 961–970, <https://doi.org/10.2514/1.12248>.
- Chen, Y.-K. and Milos, F. S. (2013). Effects of nonequilibrium chemistry and darcy—forchheimer pyrolysis flow for charring ablator. *Journal of Spacecraft and Rockets*, 50(2), 256–269, <https://doi.org/10.2514/1.a32289>.
- Cheung, S. H., Miki, K., Prudencio, E., and Simmons, C. (2016). Uncertainty quantification and robust predictive system analysis for high temperature kinetics of HCN/O₂/Ar mixture. *Chemical Physics*, 475, 136–152, <https://doi.org/10.1016/j.chemphys.2016.05.026>. 12
- Clayton, J. L. (2001). Thermal/pyrolysis gas flow analysis of carbon phenolic material. In *The Tenth Thermal and Fluids Analysis Workshop*, number 2001-211141 in NASA/CP.
- Coheur, J. (2016). Development of a unified model for flow-material interaction applied to porous ablators. Master’s thesis, von Karman Institute for Fluid Dynamics, URL <http://hdl.handle.net/2268/263536>. 13, 58, 59, 63, 69
- Coheur, J., Arnst, M., Chatelain, P., and Magin, T. (2018). A first step towards the Bayesian inference of uncertain kinetic parameters of pyrolysis decomposition laws. In *9th VKI PhD Symposium* (pp. 1–9). Rhodes-Saint-Genèse, Belgium.
- Coheur, J., Arnst, M., Chatelain, P., and Magin, T. (2019). Bayesian parameter inference for PICA devolatilization pyrolysis at high heating rates. In *10th VKI PhD Symposium* (pp. 1–12). Rhodes-Saint-Genèse, Belgium.
- Coheur, J., Arnst, M., Chatelain, P., and Magin, T. (2020). Gradient-informed sampling of Bayesian posteriors in probabilistic inversion of pyrolysis models. In *11th VKI PhD Symposium* (pp. 1–11). Rhodes-Saint-Genèse, Belgium.

- Coheur, J., Magin, T., Chatelain, P., and Arnst, M. (2021a). Bayesian identification of pyrolysis model parameters for thermal protection materials using an adaptive gradient-informed sampling algorithm with application to a mars atmospheric entry. *In preparation*. 15
- Coheur, J., Torres-Herrador, F., Chatelain, P., Mansour, N. N., Magin, T. E., and Arnst, M. (2021b). Analytical solution for multi-component pyrolysis simulations of thermal protection materials. *Journal of Materials Science*, 56(11), 6845–6860, <https://doi.org/10.1007/s10853-020-05727-8>. 14, 35, 119, 132
- Coheur, J., Turchi, A., Schrooyen, P., and Magin, T. (2017). Development of a unified model for flow-material interaction applied to porous charring ablators. In *47th AIAA Thermophysics Conference* (pp. 1–13). Denver, CO, US: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2017-3684>.
- Conesa, J., Marcilla, A., Caballero, J., and Font, R. (2001). Comments on the validity and utility of the different methods for kinetic analysis of thermogravimetric data. *Journal of Analytical and Applied Pyrolysis*, 58-59, 617–633, [https://doi.org/10.1016/s0165-2370\(00\)00130-3](https://doi.org/10.1016/s0165-2370(00)00130-3).
- Constantine, P. G., Eldred, M. S., and Phipps, E. T. (2012). Sparse pseudospectral approximation method. *Computer Methods in Applied Mechanics and Engineering*, 229-232, 1–12, <https://doi.org/10.1016/j.cma.2012.03.019>.
- Copeland, S., Mahzari, M., Cozmuta, I., and Alonso, J. (2012). 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. <https://doi.org/10.2514/6.2012-1763>. 12, 170, 177
- Covington, M. A., Heinemann, J. M., Goldstein, H. E., Chen, Y.-K., Terrazas-Salinas, I., Balboni, J. A., Olejniczak, J., and Martinez, E. R. (2008). Performance of a low density ablative heat shield material. *Journal of Spacecraft and Rockets*, 45(2), 237–247, <https://doi.org/10.2514/1.12403>.
- Cozmuta, I., White, T., Santos, J., Laub, B., and Mahzari, M. (2011). Proposed analysis process for Mars Science Laboratory heat shield sensor plug flight data. In *42nd AIAA Thermophysics Conference*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2011-3957>.
- Curtiss, C. F. and Hirschfelder, J. O. (1952). Integration of stiff equations. *Proceedings of the National Academy of Sciences*, 38(3), 235–243, <https://doi.org/10.1073/pnas.38.3.235>.
- Dania, J. (2021). *Machine Learning Metods for System Identification in Fluid Flows*. Technical report, von Karman Institute & EPF École d’Ingénieur(e)s.
- Debusschere, B., Sargsyan, K., Safta, C., and Chowdhary, K. (2017). Uncertainty quantification toolkit (UQTK). In *Handbook of Uncertainty Quantification*. Springer International Publishing.
- Di Blasi, C. (2008). Modeling chemical and physical processes of wood and biomass pyrolysis. *Progress in Energy and Combustion Science*, 34(1), 47–90, <https://doi.org/10.1016/j.pecs.2006.12.001>.
- Dick, J., Kuo, F. Y., and Sloan, I. H. (2013). High-dimensional integration: The quasi-Monte Carlo way. *Acta Numerica*, 22, 133–288, <https://doi.org/10.1017/s0962492913000044>.

- Dodwell, T. J., Ketelsen, C., Scheichl, R., and Teckentrup, A. L. (2015). A hierarchical multi-level Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1), 1075–1108, <https://doi.org/10.1137/130915005>.
- Duan, Q. Y., Gupta, V. K., and Sorooshian, S. (1993). Shuffled complex evolution approach for effective and efficient global minimization. *Journal of Optimization Theory and Applications*, 76(3), 501–521, <https://doi.org/10/dtf26t>.
- Duane, S., Kennedy, A., Pendleton, B. J., and Roweth, D. (1987). Hybrid Monte Carlo. *Physics Letters B*, 195(2), 216–222, [https://doi.org/10.1016/0370-2693\(87\)91197-x](https://doi.org/10.1016/0370-2693(87)91197-x).
- Duffa, G. (2013). *Ablative thermal protection systems modeling*. Reston, Virginia: American Institute of Aeronautics and Astronautics, Inc. 4, 7, 25
- Dullien, F. (1991). *Porous media: fluid transport and pore structure*. San Diego: Elsevier, <https://doi.org/10.1016/c2009-0-26184-8>.
- Duzel, U. (2020). *Development of Universal Solver for High Enthalpy Flows Through Ablative Materials*. PhD thesis, University of Kentucky, <https://doi.org/10.13023/ETD.2020.417>.
- Edquist, K., Dyakonov, A., Wright, M., and Tang, C. (2009). Aerothermodynamic design of the Mars Science Laboratory heatshield. In *41st AIAA Thermophysics Conference*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2009-4075>.
- Edquist, K. T., Hollis, B. R., Johnston, C. O., Bose, D., White, T. R., and Mahzari, M. (2014). Mars Science Laboratory heat shield aerothermodynamics: Design and reconstruction. *Journal of Spacecraft and Rockets*, 51(4), 1106–1124, <https://doi.org/10.2514/1.a32749>.
- Engl, H. W., Hanke, M., and Neubauer, A. (1996). *Regularization of Inverse Problems*. Mathematics and Its Applications. Springer Netherlands.
- Epstein, I. J. (1963). Conditions for a matrix to commute with its integral. *Proceedings of the American Mathematical Society*, 14(2), 266, <https://doi.org/10.1090/s0002-9939-1963-0151670-4>. 195
- Erb, R. B., Greenshields, D. H., Chauvin, L. T., Pavlosky, J. E., and Statham, C. L. (1970). Apollo thermal-protection system development. *Journal of Spacecraft and Rockets*, 7(6), 727–734, <https://doi.org/10.2514/3.30027>. 6
- Ferziger, J. H. and Kaper, H. G. (1972). *Mathematical Theory of Transport Processes in Gases*. Amsterdam: North-Holland Publishing Company.
- Flynn, J. H. (1997). The ‘temperature integral’ — its use and abuse. *Thermochimica Acta*, 300(1-2), 83–92, [https://doi.org/10.1016/s0040-6031\(97\)00046-4](https://doi.org/10.1016/s0040-6031(97)00046-4).
- Galagali, N. and Marzouk, Y. M. (2015). Bayesian inference of chemical kinetic models from proposed reactions. *Chemical Engineering Science*, 123, 170–190, <https://doi.org/10.1016/j.ces.2014.10.030>.
- Galwey, A. K. and Brown, M. E. (1999). *Thermal Decomposition of Ionic Solids*. Amsterdam: Elsevier.
- Galwey, A. K. and Brown, M. E. (2002). Application of the arrhenius equation to solid state kinetics: can this be justified? *Thermochimica Acta*, 386(1), 91–98, [https://doi.org/10.1016/S0040-6031\(01\)00769-9](https://doi.org/10.1016/S0040-6031(01)00769-9).

- Gazarik, M., Wright, M., Little, A., Cheatwood, F. M., Herath, J., Munk, M., Novak, F., and Martinez, E. (2008). Overview of the MEDLI project. *IEEE Aerospace Conference, IEEE Paper 2008-1510, March 2008.*, <https://doi.org/10.1109/aero.2008.4526285>.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2014). *Bayesian Data Analysis*. London: Chapman & Hall/CRC, third edition. 12, 87, 88, 91, 96, 117, 176, 199
- Gelman, A. G., Roberts, G. O., and Gilks, W. R. (1996). Efficient Metropolis jumping rules. In J. M. Bernardo, J. O. Berger, A. F. David, and A. F. M. Smith (Eds.), *Bayesian Statistics V*, (pp. 599–608). Oxford University Press.
- Geuzaine, C. and Remacle, J.-F. (2009). Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11), 1309–1331, <https://doi.org/10.1002/nme.2579>.
- Ghanem, R., Higdon, D., and Owhadi, H., Eds. (2017). *Handbook of Uncertainty Quantification*. Springer International Publishing, <https://doi.org/10.1007/978-3-319-12385-1>. 9, 102
- Ghanem, R. G. and Spanos, P. D. (1991). *Stochastic Finite Elements: A Spectral Approach*. Springer New York, <https://doi.org/10.1007/978-1-4612-3094-6>.
- Giles, M. B. (2015). Multilevel Monte Carlo methods. *Acta Numerica*, 24, 259–328, <https://doi.org/10.1017/s096249291500001x>.
- Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. (1996). *Markov Chain Monte Carlo in Practice*. London: Chapman & Hall/CRC.
- Girolami, M. and Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2), 123–214, <https://doi.org/10.1111/j.1467-9868.2010.00765.x>.
- Gülhan, A., Thiele, T., Siebe, F., Kronen, R., and Schleutker, T. (2019). Aerothermal measurements from the ExoMars Schiaparelli capsule entry. *Journal of Spacecraft and Rockets*, 56(1), 68–81, <https://doi.org/10.2514/1.a34228>.
- Goldstein, H. E. (1969). Pyrolysis kinetics of nylon 6-6, phenolic resin, and their composites. *Journal of Macromolecular Science: Pt. A - Chemistry*, 3(4), 649–673. xix, 7, 36, 39, 43, 44, 54, 61, 153, 171, 174, 175, 176
- Griewank, A. and Walther, A. (2008). *Evaluating Derivatives*. Society for Industrial and Applied Mathematics, <https://doi.org/10.1137/1.9780898717761>.
- Guilleminot, J. and Soize, C. (2014). Itô SDE-based generator for a class of non-Gaussian vector-valued random fields in uncertainty quantification. *SIAM Journal on Scientific Computing*, 36(6), A2763–A2786, <https://doi.org/10.1137/130948586>.
- Haario, H., Saksman, E., and Tamminen, J. (2001). An adaptive Metropolis algorithm. *Bernoulli*, 7(2), 223–242, <https://doi.org/10.2307/3318737>.
- Hairer, E., Lubich, C., and Wanner, G. (2006). *Geometric Numerical Integration*. Springer-Verlag, <https://doi.org/10.1007/3-540-30666-8>.
- Hairer, E. and Wanner, G. (1999). Stiff differential equations solved by radau methods. *Journal of Computational and Applied Mathematics*, 111(1-2), 93–111, [https://doi.org/10.1016/s0377-0427\(99\)00134-x](https://doi.org/10.1016/s0377-0427(99)00134-x).

- Harris, R., Stewart, M., and Koenig, W. (2018). Thermal protection systems technology transfer from Apollo and Space Shuttle to the Orion program. In *2018 AIAA SPACE and Astronautics Forum and Exposition*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2018-5134>. 6
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1), 97–109, <https://doi.org/10.1093/biomet/57.1.97>.
- Helber, B. (2016). *Material response characterization of low-density ablators in atmospheric entry plasmas*. PhD thesis, Vrije Universiteit Brussel & von Karman Institute for Fluid Dynamics, Belgium.
- Helber, B., Chazot, O., Hubin, A., and Magin, T. E. (2016a). Emission spectroscopic boundary layer investigation during ablative material testing in Plasmatron. *Journal of Visualized Experiments*, 112, <https://doi.org/10.3791/53742>. 7
- Helber, B., Turchi, A., and Magin, T. E. (2017). Determination of active nitridation reaction efficiency of graphite in inductively coupled plasma flows. *Carbon*, 125, 582–594, <https://doi.org/10.1016/j.carbon.2017.09.081>. 6
- Helber, B., Turchi, A., Scoggins, J. B., Hubin, A., and Magin, T. E. (2016b). Experimental investigation of ablation and pyrolysis processes of carbon-phenolic ablators in atmospheric entry plasmas. *International Journal of Heat and Mass Transfer*, 100, 810–824, <https://doi.org/10.1016/j.ijheatmasstransfer.2016.04.072>. 7, 23, 58, 189
- Heng, J., Jasra, A., Law, K. J. H., and Tarakanov, A. (2021). On unbiased estimation for discretized models. *arXiv preprint arXiv:2102.12230*.
- Herrador, F. T., Coheur, J., Blondeau, J., Meurisse, J., Panerai, F., Lachaud, J., Magin, T., and Mansour, N. N. (2019). Comparison between traditional and competitive reaction models for the pyrolysis of high temperature aerospace materials. In *AIAA Aviation 2019 Forum*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2019-3361>.
- Hillewaert, K. (2013). *Development of the discontinuous Galerkin method for large scale/high-resolution CFD and acoustics in industrial geometries*. PhD thesis, Université Catholique de Louvain, Belgium.
- Hoffman, M. D. and Gelman, A. (2014). The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15, 1593–1623.
- Hogg, R., McKean, J., and Craig, A. (2005). *Introduction to Mathematical Statistics*.
- Hosder, S. and Bettis, B. (2012). Uncertainty and sensitivity analysis for reentry flows with inherent and model-form uncertainties. *Journal of Spacecraft and Rockets*, 49(2), 193–206, <https://doi.org/10.2514/1.A32102>.
- Howe, J. T. (1989). *Hypervelocity Atmospheric Flight: Real Gas Flow Fields*. Technical memorandum 101055, NASA Ames Research Center, Moffett Field, CA, USA. 6
- Hu, D., Chen, M., Huang, Y., Wei, S., and Zhong, X. (2020). Evaluation on isothermal pyrolysis characteristics of typical technical solid wastes. *Thermochimica Acta*, 688, 178604, <https://doi.org/10.1016/j.tca.2020.178604>.
- Iain D. Boyd, T. E. S. (2017). *Nonequilibrium Gas Dynamics and Molecular Simulation*. Cambridge University Press, URL https://www.ebook.de/de/product/26338776/iain_d_boyd_thomas_e_schwartzentruber_nonequilibrium_gas_dynamics_and_molecular_simulation.html.

- Jackson, W. M. and Conley, R. T. (1964). High temperature oxidative degradation of phenol-formaldehyde polycondensates. *Journal of Applied Polymer Science*, 8(5), 2163–2193, <https://doi.org/10.1002/app.1964.070080516>.
- Jakeman, J. D., Franzelin, F., Narayan, A., Eldred, M., and Plfüger, D. (2019). Polynomial chaos expansions for dependent random variables. *Computer Methods in Applied Mechanics and Engineering*, 351, 643–666, <https://doi.org/10.1016/j.cma.2019.03.049>.
- Jasra, A., Law, K. J. H., and Lu, D. (2021). Unbiased estimation of the gradient of the log-likelihood in inverse problems. *Statistics and Computing*, 31(3), <https://doi.org/10.1007/s11222-021-09994-6>.
- Joffre, E., Derz, U., Perkinson, M.-C., Huesing, J., Beyer, F., and Perez, J.-M. S. (2018). Mars Sample Return: Mission analysis for an ESA Earth return orbiter. In *7th International Conference on Astrodynamics Tools and Techniques*. 3
- Kaipio, J. and Somersalo, E. (2005). *Statistical and Computational Inverse Problems*. New York: Springer. 12, 78
- Kendall, R. M., Barlett, E. P., Rindal, R. A., and Moyer, C. B. (1968). *An Analysis of the Coupled Chemically Reacting Boundary Layer and Charring Ablator: Part I*. CR-1060, NASA. 7, 24
- Kennedy, M. C. and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3), 425–464, <https://doi.org/10.1111/1467-9868.00294>.
- Khalil, M. and Najm, H. N. (2018). Probabilistic inference of reaction rate parameters from summary statistics. *Combustion Theory and Modelling*, 22(4), 635–665, <https://doi.org/10.1080/13647830.2017.1370557>. 12
- Koga, N. (1994). A review of the mutual dependence of Arrhenius parameters evaluated by the thermoanalytical study of solid-state reactions: the kinetic compensation effect. *Thermochimica Acta*, 244, 1–20, [https://doi.org/10.1016/0040-6031\(94\)80202-5](https://doi.org/10.1016/0040-6031(94)80202-5).
- Kolaitis, D. I. and Founti, M. A. (2013). Development of a solid reaction kinetics gypsum dehydration model appropriate for CFD simulation of gypsum plasterboard wall assemblies exposed to fire. *Fire Safety Journal*, 58, 151–159, <https://doi.org/10.1016/j.firesaf.2013.01.029>.
- Kontinos, D. and Stackpoole, M. (2008). Post-flight analysis of the Stardust Sample Return capsule Earth entry. In *46th AIAA Aerospace Sciences Meeting and Exhibit*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2008-1197>. 7, 148
- Kumar, R., Carroll, C., Hartikainen, A., and Martin, O. (2019). ArviZ a unified library for exploratory analysis of bayesian models in python. *Journal of Open Source Software*, 4(33), 1143, <https://doi.org/10.21105/joss.01143>.
- Lachaud, J., Cozmuta, I., and Mansour, N. N. (2010). Multiscale approach to ablation modeling of phenolic impregnated carbon ablators. *Journal of Spacecraft and Rockets*, 47(6), 910–921, <https://doi.org/10.2514/1.42681>. 8, 22, 24, 61, 62
- Lachaud, J., Magin, T. E., Cozmuta, I., and Mansour, N. N. (2011a). A short review of ablative-material response models and simulation tools. In L. Ouwehand (Ed.), *7th European Symposium on Aerothermodynamics*, SP-692 (pp. 91–98). Brugge, Belgium: ESTEC-ESA.

- Lachaud, J., Mansour, N., Ceballos, A., Pejakovic, D., Zhang, L., and Marschall, J. (2011b). Validation of a volume-averaged fiber-scale model for the oxidation of a carbon-fiber preform. In *42nd AIAA Thermophysics Conference* Honolulu, HA: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2011-3640>.
- Lachaud, J. and Mansour, N. N. (2014). Porous-material analysis toolbox based on OpenFOAM and applications. *Journal of Thermophysics and Heat Transfer*, 28(2), 191–202, <https://doi.org/10.2514/1.t4262>.
- Lachaud, J., Martin, A., Cozmuta, I., and Laub, B. (2011c). Ablation Workshop test case. In *4th Ablation Workshop* (pp. 1–4). Albuquerque, New Mexico.
- Lachaud, J., Martin, A., Van Eekelen, T., and Cozmuta, I. (2012). Ablation test-case series #2. In *5th Ablation Workshop* (pp. 10–18). Lexington, Kentucky.
- Lachaud, J., Scoggins, J., Magin, T., Meyer, M., and Mansour, N. (2017). A generic local thermal equilibrium model for porous reactive materials submitted to high temperatures. *International Journal of Heat and Mass Transfer*, 108, 1406–1417, <https://doi.org/10.1016/j.ijheatmasstransfer.2016.11.067>. 10, 32, 35, 42, 43, 151, 152, 171
- Lachaud, J., van Eekelen, T., Scoggins, J. B., Magin, T. E., and Mansour, N. N. (2015). Detailed chemical equilibrium model for porous ablative materials. *International Journal of Heat and Mass Transfer*, 90, 1034–1045, <https://doi.org/10.1016/j.ijheatmasstransfer.2015.05.106>.
- Lacroix, M. (2020). Propagation of uncertainties in pyrolysis kinetic parameters using polynomial chaos methods. Master’s thesis, University of Liege.
- Lani, A. (2008). *An object oriented and high performance platform for aerothermodynamics simulation*. Phd thesis, Université Libre de Bruxelles & Von Karman Institute for Fluid Dynamics, Belgium.
- Lautenberger, C. (2007). *A Generalized Pyrolysis Model for Combustible Solids*. PhD thesis, UC Berkeley.
- Lautenberger, C. and Fernandez-Pello, A. (2011). Optimization algorithms for material pyrolysis property estimation. *Fire Safety Science*, 10, 751–764, <https://doi.org/10.3801/iafss.fss.10-751>. 12, 35, 50
- Law, K. (2014). Proposals which speed up function-space MCMC. *Journal of Computational and Applied Mathematics*, 262, 127–138, <https://doi.org/10.1016/j.cam.2013.07.026>.
- Lawson, J., Stackpoole, M., and Shklover, V. (2010). Examination of scanning electron microscope and computed tomography images of pica.
- Le Maître, O. and Knio, O. (2010). *Spectral Methods For Uncertainty Quantification: With Applications To Computational Fluid Dynamics*. Springer. 9, 97, 98, 145
- Li, W., Huang, H., Tian, Y., and Zhao, Z. (2015). Nonlinear analysis on thermal behavior of charring materials with surface ablation. *International Journal of Heat and Mass Transfer*, 84, 245–252, <https://doi.org/10.1016/j.ijheatmasstransfer.2015.01.004>.
- Lin, W.-S. (2007). Quasi-steady solutions for the ablation of charring materials. *International Journal of Heat and Mass Transfer*, 50(5-6), 1196–1201, <https://doi.org/10.1016/j.ijheatmasstransfer.2006.11.011>.
- Link, W. A. and Barker, R. J. (2006). Model weights and the foundations of multimodel inference. *Ecology*, 87(10), 2626–2635, [https://doi.org/10.1890/0012-9658\(2006\)87\[2626:mwatfo\]2.0.co;2](https://doi.org/10.1890/0012-9658(2006)87[2626:mwatfo]2.0.co;2).

- Loh, W.-L. (1996). On Latin hypercube sampling. *The Annals of Statistics*, 24(5), <https://doi.org/10.1214/aos/1069362310>.
- Lunn, D., Spiegelhalter, D., Thomas, A., and Best, N. (2009). The BUGS project: Evolution, critique and future directions. *Statistics in Medicine*, 28(25), 3049–3067, <https://doi.org/10.1002/sim.3680>.
- Lv, Y. and Ihme, M. (2014). Discontinuous Galerkin method for multicomponent chemically reacting flows and combustion. 270, 105–137, <https://doi.org/10.1016/j.jcp.2014.03.029>.
- L’vov, B. (2007). *Thermal Decomposition of Solids and Melts*. Springer Netherlands, english edition, <https://doi.org/10.1007/978-1-4020-5672-7>.
- MacKay, D. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press.
- Magin, T. E. (2010–2011). Physical gas dynamics. Lecture notes. von Karman Institute for Fluid Dynamics.
- Magin, T. E., Caillault, L., Bourdon, A., and Laux, C. O. (2006). Nonequilibrium radiative heat flux modeling for the Huygens entry probe. *J. Geophys. Res.*, 111(E07S12), <https://doi.org/10.1029/2005je002616>. 5
- Mahzari, M., Braun, R. D., White, T. R., and Bose, D. (2015). Inverse estimation of the Mars Science Laboratory entry aeroheating and heatshield response. *Journal of Spacecraft and Rockets*, 52(4), 1203–1216, <https://doi.org/10.2514/1.a33053>.
- Mahzari, M., Cozmuta, I., Clark, I., and Braun, R. (2011). An inverse parameter estimation methodology for the analysis of aeroheating and thermal protection system experimental data. In *42nd AIAA Thermophysics Conference: American Institute of Aeronautics and Astronautics*. <https://doi.org/10.2514/6.2011-4027>.
- Mahzari, M., White, T., Braun, R., and Bose, D. (2013). Preliminary analysis of the Mars Science Laboratory’s entry aerothermodynamic environment and thermal protection system performance. <https://doi.org/10.2514/6.2013-185>.
- Málek, J. and Criado, J. (1992). Empirical kinetic models in thermal analysis. *Thermochimica Acta*, 203, 25–30.
- Mannella, R. (2006). Numerical stochastic integration for quasi-symplectic flows. *SIAM Journal on Scientific Computing*, 27(6), 2121–2139, <https://doi.org/10.1137/040620965>.
- Marelli, S. and Sudret, B. (2014). UQLab: A framework for uncertainty quantification in Matlab. In *The 2nd International Conference on Vulnerability and Risk Analysis and Management: American Society of Civil Engineers*. <https://doi.org/10.1061/9780784413609.257>.
- Marshall, T. and Roberts, G. (2011). An adaptive approach to langevin MCMC. *Statistics and Computing*, 22(5), 1041–1057, <https://doi.org/10.1007/s11222-011-9276-6>.
- Martin, A. (2013). Volume averaged modeling of the oxidation of porous carbon fiber material. In *44th AIAA Thermophysics Conference* (pp. 1–14). San Diego, CA: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2013-2636>. 10
- Martin, A. and Boyd, I. (2015). Strongly coupled computation of material response and nonequilibrium flow for hypersonic ablation. *Journal of Spacecraft and Rockets*, 52(1), 89–104, <https://doi.org/10.2514/1.a32847>. 10, 24, 35

- Martin, A. and Boyd, I. D. (2010). Non-Darcian behavior of pyrolysis gas in a thermal protection system. *Journal of Thermophysics and Heat Transfer*, 24(1), 60–68, <https://doi.org/10.2514/1.44103>.
- Martin, A., Scalabrin, L. C., and Boyd, I. D. (2012). High performance modeling of atmospheric re-entry vehicles. *Journal of Physics: Conference Series*, 341, 012002, <https://doi.org/10.1088/1742-6596/341/1/012002>. 10, 24
- McBride, B. J., Heimel, S., Ehlers, J. G., and Gordon, S. (1963). Thermodynamic properties to 6000 K for 210 substances involving the first 18 elements. URL <https://www.osti.gov/biblio/4634010>.
- McNeil, C. F. and Peter, A. G. (1996). User’s manual for the Langley Aerothermodynamic Upwind Relaxation Algorithm (laura).
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6), 1087–1092, <https://doi.org/10.1063/1.1699114>.
- Meurisse, J., Lachaud, J., Panerai, F., Tang, C., and Mansour, N. N. (2018). Multidimensional material response simulations of a full-scale tiled ablative heatshield. *Aerospace Science and Technology*, 76, 497–511, <https://doi.org/10.1016/j.ast.2018.01.013>. 10, 24, 33, 152, 171, 173
- Miller, R. S. and Bellan, J. (1997). A generalized biomass pyrolysis model based on superimposed cellulose, hemicellulose and lignin kinetics. *Combustion Science and Technology*, 126(1-6), 97–137, <https://doi.org/10.1080/00102209708935670>.
- Milos, F. S. (1997). Galileo probe heat shield ablation experiment. *Journal of Spacecraft and Rockets*, 34(6), 705–713, <https://doi.org/10.2514/2.3293>. 3
- Milos, F. S. and Chen, Y.-K. (2009). Two-dimensional ablation, thermal response, and sizing program for pyrolyzing ablators. *Journal of Spacecraft and Rockets*, 46(6), 1089–1099, <https://doi.org/10.2514/1.36575>.
- Milos, F. S. and Chen, Y.-K. (2010). Ablation and thermal response property model validation for phenolic impregnated carbon ablator. *Journal of Spacecraft and Rockets*, 47(5), 786–805, <https://doi.org/10.2514/1.42949>.
- Mishra, S., Schwab, C., and Sukys, J. (2012). Multi-level Monte Carlo finite volume methods for nonlinear systems of conservation laws in multi-dimensions. *Journal of Computational Physics*, 231, 3365–3388.
- Najm, H., Berry, R., Safta, C., Sargsyan, K., and Debusschere, B. (2014). Data-free inference of uncertain parameters in chemical models. *International Journal for Uncertainty Quantification*, 4(2), 111–132, <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2013005679>. 12, 117
- Natali, M., Puri, I., Rallini, M., Kenny, J., and Torre, L. (2016). Ablation modeling of state of the art EPDM based elastomeric heat shielding materials for solid rocket motors. *Computational Materials Science*, 111, 460–480, <https://doi.org/10.1016/j.commatsci.2015.09.050>. 7
- Navarro, M., Maître, O. P. L., Hoteit, I., George, D. L., Mandli, K. T., and Knio, O. M. (2018). Surrogate-based parameter inference in debris flow model. *Computational Geosciences*, 22(6), 1447–1463, <https://doi.org/10.1007/s10596-018-9765-1>. 200

- Neal, R. M. (2010). MCMC using Hamiltonian dynamics. In G. J. S. Brooks, A. Gelman and X.-L. Meng (Eds.), *Handbook of Markov Chain Monte Carlo*, (pp. 113–162). Boca Raton: Chapman and Hall–CRC.
- Nemytskii, V. and Stepanov, V. (1960). *Qualitative theory of differential equations*. Princeton University Press. 195
- Nguyen, Q., Ngo, T., Tran, P., Mendis, P., Zobec, M., and Aye, L. (2016). Fire performance of prefabricated modular units using organoclay/glass fibre reinforced polymer composite. *Construction and Building Materials*, 129, 204–215, <https://doi.org/10.1016/j.conbuildmat.2016.10.100>.
- Nocedal, J. and Wright, S. (1999). *Numerical optimization*. New York: Springer.
- Olejnik, A., Gosz, K., and Piszczczyk, Ł. (2020). Kinetics of cross-linking processes of fast-curing polyurethane system. *Thermochimica Acta*, 683, 178435, <https://doi.org/10.1016/j.tca.2019.178435>.
- Olynick, D., Chen, Y.-K., and Tauber, M. E. (1999). Aerothermodynamics of the Stardust Sample Return capsule. *Journal of Spacecraft and Rockets*, 36, 442–462, <https://doi.org/10.2514/2.3466>. 10
- Ouchi, K. and Honda, H. (1959). Pyrolysis of coal 1. thermal cracking of phenol-formaldehyde resins taken as coal models. *Fuel*, (38), 429–443.
- Palaninathan, R. and Bindu, S. (2005). Modeling of mechanical ablation in thermal protection systems. *Journal of Spacecraft and Rockets*, 42(6), 971–979, <https://doi.org/10.2514/1.10710>.
- Panerai, F. (2012). *Aerothermochemistry Characterization of Thermal Protection Systems*. PhD thesis, Università Degli Studi Di Perugia. 6
- Panerai, F., Ferguson, J. C., Lachaud, J., Martin, A., Gasch, M. J., and Mansour, N. N. (2017). Micro-tomography based analysis of thermal conductivity, diffusivity and oxidation behavior of rigid and flexible fibrous insulators. *International Journal of Heat and Mass Transfer*, 108, 801–811, <https://doi.org/10.1016/j.ijheatmasstransfer.2016.12.048>. 10, 31
- Panerai, F., Martin, A., Mansour, N. N., Sepka, S. A., and Lachaud, J. (2014). Flow-tube oxidation experiments on the carbon preform of a phenolic-impregnated carbon ablator. *Journal of Thermophysics and Heat Transfer*, 28(2), 181–190, <https://doi.org/10.2514/1.t4265>. 7, 8, 22, 23
- Panerai, F., White, J. D., Cochell, T. J., Schroeder, O. M., Mansour, N. N., Wright, M. J., and Martin, A. (2016). Experimental measurements of the permeability of fibrous carbon at high-temperature. *International Journal of Heat and Mass Transfer*, 101, 267–273, <https://doi.org/10.1016/j.ijheatmasstransfer.2016.05.016>. 7, 63
- Park, C., Jaffe, R. L., and Partridge, H. (2001). Chemical-kinetic parameters of hyperbolic Earth entry. *Journal of Thermophysics and Heat Transfer*, 15(1), 76–90, <https://doi.org/10.2514/2.6582>.
- Parno, M., Davis, A., Seelinger, L., and Marzouk, Y. (2014). Mit uncertainty quantification (muq) library.
- Patil, A., Huard, D., and Fonnesbeck, C. (2010). PyMC: Bayesian stochastic modelling inPython. *Journal of Statistical Software*, 35(4), <https://doi.org/10.18637/jss.v035.i04>.

- Paulson, J. A., Buehler, E. A., and Mesbah, A. (2017). Arbitrary polynomial chaos for uncertainty propagation of correlated random variables in dynamic systems. *IFAC Paper Online*, 50, 3548–3553, <https://doi.org/10.1016/j.ifacol.2017.08.954>.
- Pichon, T., Barreteau, R., Soyris, P., Foucault, A., Parenteau, J., Prel, Y., and Guedron, S. (2006). CMC thermal protection system for future reusable launch vehicles: generic shingle technological maturation and tests. In *57th International Astronautical Congress: American Institute of Aeronautics and Astronautics*. <https://doi.org/10.2514/6.iac-06-c2.4.03>. 6
- Pisaroni, M., Nobile, F., and Leyland, P. (2017). A continuation multi-level Monte Carlo (c-MLMC) method for uncertainty quantification in compressible inviscid aerodynamics. *Computer Methods in Applied Mechanics and Engineering*, 326, 20–50, <https://doi.org/10.1016/j.cma.2017.07.030>.
- Pochet, F., Hillewaert, K., Geuzaine, P., Remacle, J.-F., and Marchandise, É. (2013). A 3d strongly coupled implicit discontinuous galerkin level set-based method for modeling two-phase flows. *Computers & Fluids*, 87, 144–155, <https://doi.org/10.1016/j.compfluid.2013.04.010>.
- Pomerantsev, A. L., Kutsenova, A. V., and Rodionova, O. Y. (2017). Kinetic analysis of non-isothermal solid-state reactions: multi-stage modeling without assumptions in the reaction mechanism. *Physical Chemistry Chemical Physics*, 19(5), 3606–3615, <https://doi.org/10.1039/c6cp07529k>.
- Prakash, R., Burkhart, P. D., Chen, A., Comeaux, K. A., Guernsey, C. S., Kipp, D. M., Lorenzoni, L. V., Mendeck, G. F., Powell, R. W., Rivellini, T. P., Martin, A. M. S., Sell, S. W., Steltzner, A. D., and Way, D. W. (2008). Mars Science Laboratory entry, descent, and landing system overview. In *2008 IEEE Aerospace Conference: IEEE*. <https://doi.org/10.1109/AERO.2008.4526283>. 3, 4
- Prudencio, E. E. and Schulz, K. W. (2012). The parallel c++ statistical library ‘QUESO’: Quantification of uncertainty for estimation, simulation and optimization. In *Euro-Par 2011: Parallel Processing Workshops*, (pp. 398–407). Springer, https://doi.org/10.1007/978-3-642-29737-3_44.
- Raftery, A. E. and Lewis, S. M. (1996). Implementing MCMC. In W. R. Gilks, S. Richardson, and D. J. Spiegelhalter (Eds.), *Markov Chain Monte Carlo in Practice*, chapter 7, (pp. 115–130). London: Chapman & Hall/CRC.
- Reimer, T., Zuber, C., Rieser, J., and Rothermel, T. (2018). Determination of the mechanical properties of the lightweight ablative material Zuram. In *Ceramic Transactions Series*, (pp. 311–326). John Wiley & Sons, Inc., <https://doi.org/10.1002/9781119423829.ch28>. 7, 39
- Rein, G., Lautenberger, C., Fernandez-Pello, A. C., Torero, J. L., and Urban, D. L. (2006). Application of genetic algorithms and thermogravimetry to determine the kinetics of polyurethane foam in smoldering combustion. *Combustion and Flame*, 146(1-2), 95–108, <https://doi.org/10.1016/j.combustflame.2006.04.013>.
- Ritter, H., Bayle, O., Mignot, Y., Boulier, E., Portela, P., Builly, J.-M., and Sharda, R. (2011). Ongoing European developments on entry heatshields and TPS materials. In *8th International Planetary Probe Workshop* (pp. 1–8). Portsmouth, Virginia, USA. 7
- Rivier, M., Lachaud, J., and Congedo, P. (2019). Ablative thermal protection system under uncertainties including pyrolysis gas composition. *Aerospace Science and Technology*, 84, 1059–1069, <https://doi.org/10.1016/j.ast.2018.11.048>. 10, 12, 169, 170, 187
- Rivière, B. (2008). *Discontinuous Galerkin Methods for Solving Elliptic and Parabolic Equations*. SIAM, <https://doi.org/10.1137/1.9780898717440>.

- Robert, C. P. and Casella, G. (1999). *Monte Carlo Statistical Methods*. Springer New-York, <https://doi.org/10.1007/978-1-4757-3071-5>.
- Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer New-York, 2nd edition, <https://doi.org/10.1007/978-1-4757-4145-2>.
- Roberts, G. O. (2002). Langevin diffusions and Metropolis–Hastings algorithms. *Methodology and Computing in Applied Probability*, 4(4), 337–357, <https://doi.org/10.1023/a:1023562417138>.
- Roberts, G. O. and Rosenthal, J. S. (2009). Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18(2), 349–367, <https://doi.org/10.1198/jcgs.2009.06134>.
- Rodionova, O. E. and Pomerantsev, A. L. (2005). Estimating the parameters of the Arrhenius equation. *Kinetics and Catalysis*, 46(3), 305–308, <https://doi.org/10.1007/s10975-005-0077-9>.
- Rosenthal, J. S. (2010). Optimal proposal distributions and adaptive MCMC. In G. J. S. Brooks, A. Gelman and X.-L. Meng (Eds.), *Handbook of Markov Chain Monte Carlo*. Boca Raton: Chapman and Hall–CRC. 90, 122, 203
- Rothermel, T., Zuber, C. and Herdrich, G., and Walpot, L. (2014). A light-weight ablative material for research purposes. In *6th Ablation Workshop* Urbana-Champaign, IL.
- Salvatier, J., Wiecki, T. V., and Fonnesbeck, C. (2016). Probabilistic programming in python using PyMC3. *PeerJ Computer Science*, 2, e55, <https://doi.org/10.7717/peerj-cs.55>.
- Sanson, F., Panerai, F., Magin, T. E., and Congedo, P. M. (2018). Robust reconstruction of the catalytic properties of thermal protection materials from sparse high-enthalpy facility experimental data. *Experimental Thermal and Fluid Science*, 96, 482–492, <https://doi.org/10.1016/j.expthermflusci.2018.03.028>. 10
- Schrooyen, P. (2015). *Numerical Simulation of Aerothermal Flows Through Ablative Thermal Protection Systems*. PhD thesis, Université Catholique de Louvain & von Karman Institute for Fluid Dynamics. 8, 13, 21, 22, 23, 24, 28, 29, 30, 33, 57, 63, 67
- Schrooyen, P., Coheur, J., Turchi, A., Hillewaert, K., Chatelain, P., and Magin, T. (2017). Numerical simulation of a non-charring ablator in high enthalpy flows by means of a unified flow-material solver. In *47th AIAA Thermophysics Conference* (pp. 1–15). Denver, CO, US: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2017-3352>.
- Schrooyen, P., Hillewaert, K., Magin, T. E., and Chatelain, P. (2016). Fully implicit discontinuous Galerkin solver to study surface and volume ablation competition in atmospheric entry flows. *International Journal of Heat and Mass Transfer*, 103, 108–124, <https://doi.org/10.1016/j.ijheatmasstransfer.2016.07.022>. 10, 24, 32, 35, 58
- Schrooyen, P., Turchi, A., Hillewaert, K., Chatelain, P., and Magin, T. E. (2018). Two-way coupled simulations of stagnation-point ablation with transient material response. *International Journal of Thermal Sciences*, 134, 639–652, <https://doi.org/10.1016/j.ijthermalsci.2018.08.014>.
- Scoggins, J. B. (2017). *Development of numerical methods and study of coupled flow, radiation, and ablation phenomena for atmospheric entry*. PhD thesis, Université Paris-Saclay and von Karman Institute for Fluid Dynamics. 8, 25, 28
- Scoggins, J. B., Leroy, V., Bellas-Chatzigeorgis, G., Dias, B., and Magin, T. E. (2020). Mutation++: MULTicomponent thermodynamic and transport properties for IONized gases in c++. *SoftwareX*, 12, 100575, <https://doi.org/10.1016/j.softx.2020.100575>.

- Scoggins, J. B. and Magin, T. E. (2015). Gibbs function continuation for linearly constrained multiphase equilibria. *Combustion and Flame*, 162(12), 4514–4522, <https://doi.org/10.1016/j.combustflame.2015.08.027>.
- Seedhouse, E. (2016). Dragon design, development, and test. In *SpaceX's Dragon: America's Next Generation Spacecraft*, (pp. 23–44). Springer International Publishing, https://doi.org/10.1007/978-3-319-21515-0_3. 8
- Šesták, J. and Berggren, G. (1971). Study of the kinetics of the mechanism of solid-state reactions at increasing temperatures. *Thermochimica Acta*, 3(1), 1–12, [https://doi.org/10.1016/0040-6031\(71\)85051-7](https://doi.org/10.1016/0040-6031(71)85051-7).
- Sharpe, R. and Wright, M. D. (2009). Nasa materials research for extreme conditions. 7
- Smith, R. C. (2014). *Uncertainty Quantification: Theory, Implementation, and Applications*. Philadelphia: Society for Industrial and Applied Mathematics. 9, 12, 77, 78, 79, 85, 88, 91, 107
- Smolyak, S. A. (1963). Quadrature and interpolation formulas for tensor products of certain classes of functions. *Dokl. Akad. Nauk SSSR*, 148(5), 1042–1045. English Russian Translation: Soviet Mathematics Doklady, 4, 240–243 (1963).
- Soize, C. (2008). Construction of probability distributions in high dimension using the maximum entropy principle: Applications to stochastic processes, random fields and random matrices. *International Journal for Numerical Methods in Engineering*, 76(10), 1583–1611, <https://doi.org/DOI:10.1002/nme.2385>.
- Soize, C. (2015). Polynomial chaos expansion of a multimodal random vector. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1), 34–60, <https://doi.org/10.1137/140968495>.
- Soize, C. (2017). *Uncertainty Quantification*. Springer International Publishing, <https://doi.org/10.1007/978-3-319-54339-0>.
- Soize, C. and Ghanem, R. (2016). Data-driven probability concentration and sampling on manifold. *Journal of Computational Physics*, 321, 242–258, <https://doi.org/10.1016/j.jcp.2016.05.044>.
- Stackpoole, M., Sepka, S., Cozmuta, I., and Kontinos, D. (2008). Post-flight evaluation of Stardust Sample Return capsule forebody heatshield material. In *46th AIAA Aerospace Sciences Meeting and Exhibit*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2008-1202>. 7, 22, 148
- Stokes, E. H. (1995). Kinetics of pyrolysis mass loss from cured phenolic resin. *Journal of Thermophysics and Heat Transfer*, 9(2), 352–358, <https://doi.org/10.2514/3.667>.
- Strong, M. and Oakley, J. E. (2014). When is a model good enough? Deriving the expected value of model improvement via specifying internal model discrepancies. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1), 106–125, <https://doi.org/10.1137/120889563>.
- Stuart, A. M. (2010). Inverse problems: A Bayesian perspective. *Acta Numerica*, 19, 451–559, <https://doi.org/10.1017/S0962492910000061>. 12, 78
- Sun, N.-Z. and Sun, A. (2015). *Model Calibration and Parameter Estimation: For Environmental and Water Resource Systems*. New York: Springer.
- Sykes, G. F. (1967). *Decomposition Characteristics of a Char-Forming Phenolic Polymer Used for Ablative Composites*. Technical report, NASA TN D-3810, National Aeronautics and Space Administration, Washington. xvii, xix, 7, 37, 38, 64, 65, 172

- Tanaka, H. (1995). Thermal analysis and kinetics of solid state reactions. *Thermochimica Acta*, 267, 29–44, [https://doi.org/10.1016/0040-6031\(95\)02464-6](https://doi.org/10.1016/0040-6031(95)02464-6).
- Tang, K., Congedo, P. M., and Abgrall, R. (2015). Sensitivity analysis using anchored ANOVA expansion and high-order moments computation. *International Journal for Numerical Methods in Engineering*, 102(9), 1554–1584, <https://doi.org/10.1002/nme.4856>.
- Tarantola, A. (2005). *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial and Applied Mathematics, <https://doi.org/10.1137/1.9780898717921.12>, 78, 87
- Torre, L., Kenny, J. M., and Maffezzoli, A. M. (1998a). Degradation behaviour of a composite material for thermal protection systems part I—experimental characterization. *Journal of Materials Science*, 33(12), 3137–3143, <https://doi.org/10.1023/A:1004399923891>.
- Torre, L., Kenny, J. M., and Maffezzoli, A. M. (1998b). Degradation behaviour of a composite material for thermal protection systems Part II process simulation. *Journal of Materials Science*, 33(12), 3145–3149, <https://doi.org/10.1023/A:1004352007961>.
- Torres-Herrador, F., Coheur, J., Panerai, F., Magin, T. E., Arnst, M., Mansour, N. N., and Blondeau, J. (2019a). Competitive kinetic model for the pyrolysis of the phenolic impregnated carbon ablator. *Aerospace Science and Technology*, 100, 105784, <https://doi.org/10.1016/j.ast.2020.105784>. 13, 14, 15, 35, 41, 48, 115, 117, 133, 138, 181
- Torres-Herrador, F., Eschenbacher, A., Coheur, J., Blondeau, J., Magin, T. E., and Geem, K. M. V. (2021). Decomposition of carbon/phenolic composites for aerospace heatshields: Detailed speciation of phenolic resin pyrolysis products. *Aerospace Science and Technology*, (pp. 107079)., <https://doi.org/10.1016/j.ast.2021.107079>. 11
- Torres-Herrador, F., Leroy, V., Helber, B., Contat-Rodrigo, L., Lachaud, J., and Magin, T. (2020). Multicomponent pyrolysis model for thermogravimetric analysis of phenolic ablators and lignocellulosic biomass. *AIAA Journal*, 58(9), 4081–4089, <https://doi.org/10.2514/1.j059423>.
- Torres-Herrador, F., Meurisse, J. B., Panerai, F., Blondeau, J., Lachaud, J., Bessire, B. K., Magin, T. E., and Mansour, N. N. (2019b). A high heating rate pyrolysis model for the phenolic impregnated carbon ablator (PICA) based on mass spectroscopy experiments. *Journal of Analytical and Applied Pyrolysis*, 141, 104625–104635, <https://doi.org/10.1016/j.jaap.2019.05.014>. xix, 12, 35, 36, 43, 50, 51, 52, 119, 131, 187
- Tran, H., Johnson, C., Hsu, M.-T., Chem, H., Dill, H., Chen-Johnson, A., Tran, H., Johnson, C., Hsu, M.-T., Chem, H., Dill, H., and Chen-Johnson, A. (1997a). Qualification of the forebody heatshield of the Stardust’s sample return capsule. In *32nd Thermophysics Conference*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.1997-2482>. 10
- Tran, H. K., Johnson, C. E., Rasky, D. J., Hui, F. C. L., Hsu, M.-T., Chen, T., Chen, Y. K., Paragas, D., and Kobayashi, L. (1997b). *Phenolic Impregnated Carbon Ablators (PICA) as Thermal Protection Systems for Discovery Missions*. Technical Memorandum TM 110440, NASA Ames Research Center. 6
- Trick, K., Saliba, T., and Sandhu, S. (1997). A kinetic model of the pyrolysis of phenolic resin in a carbon/phenolic composite. *Carbon*, 35(3), 393–401, [https://doi.org/10.1016/s0008-6223\(97\)89610-8](https://doi.org/10.1016/s0008-6223(97)89610-8). 7, 22, 39, 40, 43

- Trick, K. A. and Saliba, T. E. (1995). Mechanisms of the pyrolysis of phenolic resin in a carbon/phenolic composite. *Carbon*, 33(11), 1509–1515, [https://doi.org/10.1016/0008-6223\(95\)00092-r](https://doi.org/10.1016/0008-6223(95)00092-r). 7, 22, 39, 40, 44
- Trumble, K. A., Cozmuta, I., Sepka, S., Jenniskens, P., and Winter, M. (2010). Postflight aerothermal analysis of the Stardust Sample Return capsule. *Journal of Spacecraft and Rockets*, 47(5), 765–774, <https://doi.org/10.2514/1.41514>. 7
- Turchi, A., Bianchi, D., Nasuti, F., and Onofri, M. (2013). A numerical approach for the study of the gas–surface interaction in carbon–phenolic solid rocket nozzles. *Aerospace Science and Technology*, 27(1), 25–31, <https://doi.org/10.1016/j.ast.2012.06.003>.
- Turchi, A., Congedo, P. M., Helber, B., and Magin, T. E. (2017a). Thermochemical ablation modeling forward uncertainty analysis—part II: Application to plasma wind-tunnel testing. *International Journal of Thermal Sciences*, 118, 510–517, <https://doi.org/10.1016/j.ijthermalsci.2017.04.005>.
- Turchi, A., Congedo, P. M., and Magin, T. E. (2017b). Thermochemical ablation modeling forward uncertainty analysis—part i: Numerical methods and effect of model parameters. *International Journal of Thermal Sciences*, 118, 497–509, <https://doi.org/10.1016/j.ijthermalsci.2017.04.004>.
- Urzay, J., Kseib, N., Constantine, P. G., Davidson, D. F., and Iaccarino, G. (2012). Uncertainty-quantifying models for chemical-kinetic rates. *Center for Turbulence Research Annual Briefs*. 12, 120
- Valdés-Parada, F. J. and Lasseux, D. (2021). A novel one-domain approach for modeling flow in a fluid-porous system including inertia and slip effects. *Physics of Fluids*, 33(2), 022106, <https://doi.org/10.1063/5.0036812>.
- van Eekelen, T., Martin, A., Lachaud, J., and Cozmuta, I. (2014). Ablation test-case series #3: numerical simulation of ablative-material response: code and model comparisons. In *6th Ablation Workshop* (pp. 19–38). Urbana-Champaign, IL.
- Vasavada, A. R., Chen, A., Barnes, J. R., Burkhart, P. D., Cantor, B. A., Dwyer-Cianciolo, A. M., Fergason, R. L., Hinson, D. P., Justh, H. L., Kass, D. M., Lewis, S. R., Mischna, M. A., Murphy, J. R., Rafkin, S. C. R., Tyler, D., and Withers, P. G. (2012). Assessment of environments for Mars Science Laboratory entry, descent, and surface operations. *Space Science Reviews*, 170(1-4), 793–835, <https://doi.org/10.1007/s11214-012-9911-3>. 3, 11
- Villa, U., Petra, N., and Ghattas, O. (2021). hIPPYlib. *ACM Transactions on Mathematical Software*, 47(2), 1–34, <https://doi.org/10.1145/3428447>.
- Vincenti, W. G. and Kruger, C. H. (1965). *Introduction to Physical Gas Dynamics*. New York: John Wiley & Sons.
- Vyazovkin, S., Burnham, A. K., Criado, J. M., Pérez-Maqueda, L. A., Popescu, C., and Sbirrazzuoli, N. (2011). ICTAC kinetics committee recommendations for performing kinetic computations on thermal analysis data. *Thermochimica Acta*, 520(1-2), 1–19, <https://doi.org/10.1016/j.tca.2011.03.034>.
- Vyazovkin, S. and Wight, C. A. (1997). Kinetics in solids. *Annual Review of Physical Chemistry*, 48(1), 125–149, <https://doi.org/10.1146/annurev.physchem.48.1.125>.
- Way, D. W., Powell, R. W., Chen, A., Steltzner, A. D., Martin, A. M. S., Burkhart, P. D., and Mendeck, G. F. (2007). Mars Science Laboratory: Entry, descent, and landing system performance. In *2007 IEEE Aerospace Conference: IEEE*. <https://doi.org/10.1109/AERO.2007.352821>. 3

- Weng, H. and Martin, A. (2014). Multidimensional modeling of pyrolysis gas transport inside charring ablative materials. *Journal of Thermophysics and Heat Transfer*, 28(4), 583–597, <https://doi.org/10.2514/1.t4434>.
- Weng, H. and Martin, A. (2017). Development of a universal solver and its application to ablation problems. In *47th AIAA Thermophysics Conference* (pp. 1–16). Denver, CO, US: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2017-3355>.
- Whitaker, S. (1999). *The Method of Volume Averaging*, volume 13 of *Theory and Applications of Transport in Porous Media*. Springer Netherlands, 1st edition.
- White, C., Scanlon, T. J., and Brown, R. E. (2016). Permeability of ablative materials under rarefied gas conditions. *Journal of Spacecraft and Rockets*, 53(1), 134–142, <https://doi.org/https://doi.org/10.2514/1.A33279>.
- White, J. E., Catallo, W. J., and Legendre, B. L. (2011). Biomass pyrolysis kinetics: A comparative critical review with relevant agricultural residue case studies. *Journal of Analytical and Applied Pyrolysis*, 91(1), 1–33, <https://doi.org/10.1016/j.jaap.2011.01.004>.
- White, T. R., Mahzari, M., Bose, D., and Santos, J. A. (2013). Post-flight analysis of Mars Science Laboratory’s entry aerothermal environment and thermal protection system response. In *44th AIAA Thermophysics Conference*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2013-2779>. 11, 171
- Wiener, N. (1938). The homogeneous chaos. *American Journal of Mathematics*, 60(4), 897–936, <https://doi.org/10.2307/2371268>.
- Willcockson, W. H. (1999). Stardust Sample Return capsule design experience. *Journal of Spacecraft and Rockets*, 36, 470–474, <https://doi.org/10.2514/2.3468>. 7
- Williams, S. D. and Curry, D. M. (1992). *Thermal protection materials: Thermophysical property data*. Technical report, National Aeronautics and Space Administration. 5
- Witteveen, J. A. and Bijl, H. (2006). Modeling arbitrary uncertainties using Gram-Schmidt polynomial chaos. In *44th AIAA Aerospace Sciences Meeting and Exhibit* (pp. 1–17).: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2006-896>.
- Wong, H.-W., Peck, J., Assif, J., Lachaud, J., and N., M. (2015a). Quantitative determination of species production from the pyrolysis of the phenolic impregnated carbon ablator (PICA). In *Science and Tech. Forum and Expo*. Kissimmee, FL.
- Wong, H.-W., Peck, J., Assif, J., Panerai, F., Lachaud, J., and Mansour, N. N. (2016). Detailed analysis of species production from the pyrolysis of the phenolic impregnated carbon ablator. *Journal of Analytical and Applied Pyrolysis*, 122, 258–267, <https://doi.org/10.1016/j.jaap.2016.09.016>. 7, 11, 39, 44, 46, 133, 139
- Wong, H.-W., Peck, J., Bonomi, R. E., Assif, J., Panerai, F., Reinisch, G., Lachaud, J., and Mansour, N. N. (2015b). Quantitative determination of species production from phenol-formaldehyde resin pyrolysis. *Polymer Degradation and Stability*, 112, 122–131, <https://doi.org/10.1016/j.polymdegradstab.2014.12.020>. xvii, 7, 37, 39, 40, 44, 46, 195
- Wright, M., Cozmuta, I., Laub, B., Chen, Y.-K., and Wilcoxson, W. H. (2011). Defining ablative thermal protection system margins for planetary entry vehicles. In *42nd AIAA Thermophysics Conference*: American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2011-3757>. 7
- Wright, M., White, T., and Mangini, N. L. (2009). Data Parallel Line Relaxation (DPLR) code user manual: Acadia - version 4.01.1.

- Wright, M. J., Beck, R. A. S., Edquist, K. T., Driver, D., Sepka, S. A., Slimko, E. M., and Willcockson, W. H. (2014). Sizing and margins assessment of Mars Science Laboratory aeroshell thermal protection system. *Journal of Spacecraft and Rockets*, 51(4), 1125–1138, <https://doi.org/10.2514/1.a32579>. 8, 10, 147, 171
- Wright, M. J., Bose, D., and Chen, Y.-K. (2007). Probabilistic modeling of aerothermal and thermal protection material response uncertainties. *AIAA Journal*, 45(2), 399–410, <https://doi.org/10.2514/1.26018>. 12
- Xiu (2007). Efficient collocational approach for parametric uncertainty analysis. *Communications in computational physics*, 2(2), 293–309.
- Xiu, D. (2010). *Numerical Methods for Stochastic Computations*. Princeton University Press, <https://doi.org/10.2307/j.ctv7h0skv>.
- Xiu, D. and Karniadakis, G. E. (2002). The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24(2), 619–644, <https://doi.org/https://doi.org/10.1137/S1064827501387826>.
- Yuen, A., Chen, T., Wang, C., Wei, W., Kabir, I., Vargas, J., Chan, Q., Kook, S., and Yeoh, G. (2020). Utilising genetic algorithm to optimise pyrolysis kinetics for fire modelling and characterisation of chitosan/graphene oxide polyurethane composites. *Composites Part B: Engineering*, 182, 107619, <https://doi.org/10.1016/j.compositesb.2019.107619>.
- Yuen, A. C. Y., Chen, T. B. Y., Yeoh, G. H., Yang, W., Cheung, S. C.-P., Cook, M., Yu, B., Chan, Q. N., and Yip, H. L. (2018). Establishing pyrolysis kinetics for the modelling of the flammability and burning characteristics of solid combustible materials. *Journal of Fire Sciences*, 36(6), 494–517, <https://doi.org/10.1177/0734904118800907>.
- Zhang, W., Zhang, J., Ding, Y., He, Q., Lu, K., and Chen, H. (2021). Pyrolysis kinetics and reaction mechanism of expandable polystyrene by multiple kinetics methods. *Journal of Cleaner Production*, 285, 125042, <https://doi.org/10.1016/j.jclepro.2020.125042>.