POLITECNICO DI TORINO

DEPARTMENT OF MECHANICAL AND AEROSPACE ENGINEERING (DIMEAS)

Master's Degree Thesis in Aerospace Engineering

Novel framework for Bayesian inference of material properties of ablative thermal protection systems

Terzo Riccardo (Matr. 295903)

Supervisor: Dr. Fabrizio Bisetti - The University of Texas at Austin

Supervisor: Prof. Domenic D'Ambrosio - Politecnico di Torino

A.A. 2023-2024

I confirm that this thesis type is my own work and I have documented all sources and material used.

Author: Terzo Riccardo

Acknowledgments

I would like to express my sincere gratitude to Dr. Fabrizio Bisetti of The University of Texas at Austin, who welcomed me into his research group, the Reactive Flow Modeling Laboratory (RFML), allowing me to work in a collaborative, dynamic, and supportive environment. Additionally, I would like to thank Professor Domenic D'Ambrosio from Politecnico di Torino for his support during this journey and for his readiness to assist me in the preparation of this work. Their passion and relentless pursuit of perfection, down to the smallest detail, are qualities that I will carry with me throughout my whole professional life.

Riccardo Terzo

Abstract

The work performed in this thesis focuses on assessing the PICA thermal conductivity ratio and the cold wall stagnation point heat flux using Bayesian Inference. Within the inference framework, which has been implemented in *MATLAB* using *UQLab* [1], a surrogate model based on solutions coming from the McDougall solver [2] has been introduced. Before building this surrogate by running the solver and storing its solutions, a mesh convergence study, as well as a time step convergence study and a validation campaign using available experimental data have been conducted.

By examining the uncertainty associated with the two considered parameters, the inference framework allowed the definition of the posterior distributions of α, the thermal conductivity ratio, and q_0, the cold wall stagnation point heat flux. The posterior distributions of the inferred parameters derived using some thermocouples (TCs) were then used to estimate the temperature evolution in other thermocouples. This was done to validate the previously calculated posteriors through a forward propagation analysis.

To determine which thermocouples to use when inferring the two parameters, the Kullback Leibler Divergence was employed. The KL measured the information gain from the prior to the posterior distribution for each set of TCs, and both parameters. Using the Kullback Leibler Divergence, the most informative combinations of thermocouples were identified, achieving the highest information gain about both the inferred parameters.

Finally, a study about the correlation factor between the two parameters was performed. This study indicated the possibility of performing a sequential inference analysis, and the results were compared to those obtained when inferring both parameters within the same simulation.

Contents

List of figures

List of tables

Chapter 1

Introduction

As the quest for space exploration and hypersonic travel advances, the need for protective materials that can withstand extreme atmospheric conditions becomes paramount.
These materials are essential for ensuring vehicle safety without increasing weight, and so they are subjected to ongoing research. In this context, their mechanical and thermal properties are often either declared nominally, with inherent uncertainties, or require costly experiments for accurate assessment. The integration of statistical tools, such as Bayesian inference, with experimental outcomes, could be instrumental in precisely defining these properties while accounting for the associated uncertainties. Such integration could play a vital role in reducing design phase costs and enhancing the reliability of safety measures.

As part of the vast family of protective materials, ablative materials are essential during atmospheric re-entry, as they are designed to absorb and dissipate extreme heat generated during this phase. Re-entry vehicles face one of the most hostile environments known to man: the intense heat generated due to rapid compression and friction with Earth's atmosphere at high speeds. This heating can lead to temperatures that exceed thousands of degrees Celsius, necessitating a protective shield to prevent the vehicle's failure.

When exposed to high temperatures, ablative materials undergo a process of pyrolysis and oxidation. This complex phenomenon is influenced by many factors such as material composition, heating rate, aerodynamics, and thermochemical reactions. These elements make it exceedingly challenging to model and predict, so experiments, such as the ones

performed by Milos and Chen [4], are often needed to assess the ablative material response to high-enthalpy flows.

However, the computational model developed by McDougall [2], which will be the core of the inference framework, achieved results that closely mirror experimental outcomes. This model allows the calculation of the temperature distribution inside a sample of ablative material, subjected to a heating condition.

This thesis delves into a specific class of ablative materials known as PICA (Phenolic Impregnated Carbon Ablator), even though the work performed can be seen as a valid tool for a broad range of materials. PICAs are lightweight ablative composites developed by NASA Ames Research Center (ARC) and first used in the Stardust sample return capsule, which is the fastest man-made re-entry vehicle to date. On that occasion, the capsule entered Earth's atmosphere with a speed of 12.5 km/s, leading the ablative material to reach a temperature of more than 2900 °C (Fig. 1.1). PICAs were later successfully used in the Mars 2001 [7] and Mars Science Laboratory mission [8].

PICA stands out for its low thermal conductivity and high heat capacity, which enables it to manage the extreme heat fluxes encountered during atmospheric re-entry. PICA is made of a highly porous carbon fiberform structure infused with a commercial phenolic resin. When exposed to high enthalpy flows, the phenolic resin breaks down, releasing gaseous products and leaving behind char [9], which insulates and shields the underlying structure. The porous nature of the fiberform's structure, which is prone to oxidation or sublimation, is essential, as it allows the gases generated from the decomposition of the resin to effectively leave the surface.

When compared to other ablative materials, PICA is renowned for its low density and effective ablative capacity. Traditional materials such as reinforced carbon-carbon (RCC) or silica-based ablators may offer higher strength at extreme temperatures but are often heavier and less efficient at dissipating heat. PICA's lightweight nature combined with its ablation efficiency makes it an excellent choice for missions requiring a reduction in launch mass without compromising on thermal protection.

The specific makeup of PICA varies with the fibrous preform selected and the quantity and variety of phenolic resin incorporated. In recent years, PICA ablators first developed by

NASA, have been subjected to improvements and changes. SpaceX produced its own PICA, called PICA-X, employed in the Dragon capsule [10]. Subsequently, the same company developed another version called PICA-3, first employed in the flight demonstration of the Crew Dragon Capsule in 2019 and then put into regular service on that same spacecraft.

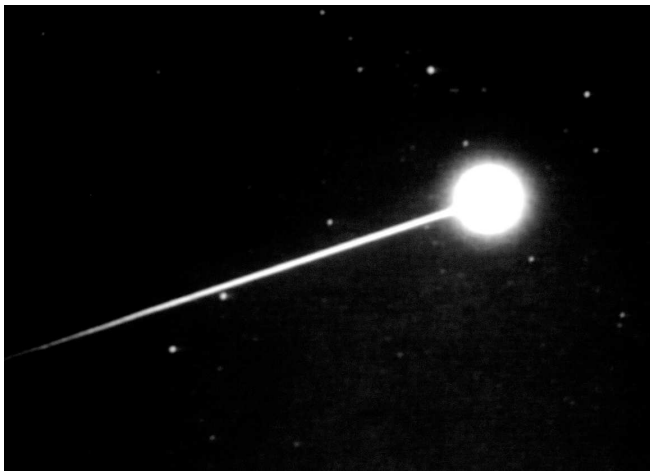

Figure 1.1: Dr. Mike Taylor (Utah State University) got a spectacular video sequence of the Stardust capsule re-entry onboard the DC-8 aircraft using a Xybion intensified camera. Credits: Dr. Mike Taylor. [3]

Chapter 2

Research Objectives

As previously discussed in the Introduction of this thesis, this work aims to highlight the great potential of combining computational models and statistical tools to assess material properties. The key objectives of this thesis can be summarized as follows:

1. Harnessing the powerful tool of Bayesian inference to determine thermal properties of PICA materials, using the experiments conducted by Milos and Chen [4] as a guideline;

2. Highlight the broad potential of using inference methods to estimate material parameters of PICA and, by extension, materials in general. This could help guide the design phase and allow for a more precise and informed material development process, saving costs and time;

3. By estimating PICA material properties through Bayesian inference, this work also provides critical insights into the most informative thermocouple placements used in the experiments of Milos and Chen [4], needed for capturing specific material properties. These insights pave the way for more efficient and cost-effective experimental designs and can be extrapolated to enhance testing methodologies for a broader range of materials.

Chapter 3

Thermal Model

In this Chapter, the computational model developed by McDougall [2] is briefly introduced. This model has been the core of the inference framework later employed for the inference analyses.

The objective of this computational model is to precisely simulate the heat transfer in an ablative material, replicating the experiments conducted by F.S. Milos and Y.-K. Chen [4]. These experiments were performed in the Aerodynamic Heating Facility (AHF) at NASA Ames Research Center (ARC). They focused on the investigation of the ablative material response to different heating conditions, obtained using arc-jets and represented by the value of the stagnation point cold wall heat flux, q_0. Heat flux at the stagnation point was recorded with a slug calorimeter, which was designed to match the external shape of the TPS samples under examination. This device was exposed to the arc-jet stream for about 3 seconds [4]. Based on the previous experience with NASA arc jet facilities, an uncertainty between ± 10 and 15 % is expected on heat flux measurements [4]. Aiming to reproduce these experiments, McDougall model [2] aims to calculate the temperature evolution over time at specific locations inside the ablative material sample, where thermocouples are placed in the real experimental setup.

The ablative material sample used in the experiments is called "iso-q". This type of sample offers a significant advantage over the traditional flat-faced geometry by providing a heat flux distribution that remains relatively constant across most of the front face.

In this sense, a typical computational fluid dynamics (CFD) heat flux calculation for a 10.16 cm-diameter sample is provided in Section 3.6. These particular samples also retain the same surface profile as they ablate, as also visualized in Fig. 3.1.

In the experiments, thermocouples have been placed inside the sample, along its axis and radial coordinates, as presented in Figure 3.1. Among all the thermocouples utilized, those aligned with the axis of the sample (i.e. along its "z" coordinate axis, y = 0) were inserted using press-fit plugs. In contrast, the off-axis thermocouples (i.e. thermocouples placed along the "y" coordinate) were installed with the help of curved plugs. Specifically, within the on-axis thermocouples, Type R was selected for those positioned nearest to the surface. Meanwhile, the thermocouples situated deeper inside, along with all radial thermocouples, were of Type K.

It is crucial to note that the aim of this computational model is not to develop an exhaustive model covering all aspects of pyrolysis and ablation. Instead, it targets a specific scenario where minimal pyrolysis takes place within the material. Thus, even though comparisons will be drawn from the entire spectrum of available experimental data, the model seeks to represent the initial few seconds of arc-jet testing, a period when temperature measurements from thermocouples are particularly sensitive to heat transfer only.

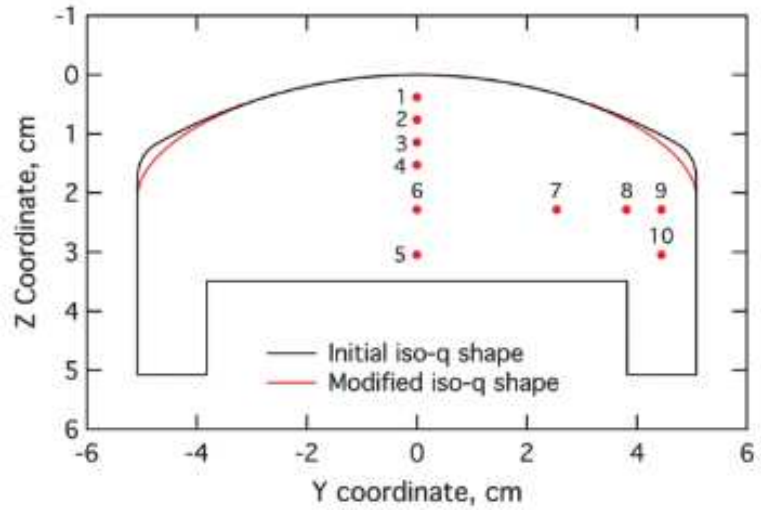

Figure 3.1: Coordinate system of the sample and thermocouples locations. Adapted from [4].

3.1 Assumptions

McDougall's modelization of the phenomena aims to represent three main aspects: heat conduction, radiation at exposed surface, and energy exchange through the chemical reactions of pyrolysis.

The heat conduction in PICA is related to an orthotropic thermal conductivity tensor, due to the fact that PICAs are often manufactured in layers [6] [11]. Given that, heat conduction will occur at different rates along the two principal directions (parallel 'y' and perpendicular 'z', as represented in Figure 3.1), due to distinct thermal conductivity values for each. This leads to a different spatial heat propagation speed in the two directions, and as a result, to an uneven spatial temperature distribution inside the sample. For what concerns the energy exchange related to pyrolysis, this is taken into account in the energy equation (Eq. 3.1) by the source term. Finally, the radiation at the exposed surface is taken into account by using the Stefan-Boltzmann formula for radiation and taken into consideration via the boundary condition presented in Section 3.3. Absorption is negligible as boundary layer interactions are not considered in the model.

Some processes, which can be seen as negligible when considering the first few seconds of exposure and a low heating condition, have not been taken into consideration. These include convection of pyrolysis gases, boundary layer interactions, and surface recession. For what concerns surface recession, it is known from the work conducted by Milos and Chen in [6], that a TPS (Thermal Protection System) sample tested for 200 seconds, with a radius of 10.16 cm, has had a recession of 1.3 cm at the centerline. Given that we will consider only a few seconds of exposure (around 40 seconds), recession can be seen as negligible. For what concerns not including the boundary layer interactions, this means that the model does not account for the transfer of heat and mass between the material and its surroundings, as well as the aerodynamic drag experienced by the material during ablation. A final approximation pertains to the fact that specific heat and thermal conductivity are considered uniform across the sample and constant over the considered time period, even though it is known they should change during pyrolysis.

In order to reduce the impact of the neglected processes, along with a short time scale, a reference case with a low cold wall heat flux has been used. A reference case is essential when studying the model's capability to replicate experimental data. In particular, among all the heat fluxes samples have been subjected to, a condition of cold wall heat flux of 169 W/cm^2 has been chosen as the reference case.

3.1.1 Initial conditions

A uniform initial temperature distribution along the sample has been chosen as the initial condition and set to 300 K. This temperature is later referred to as T_0. However, this assumption has not been confirmed by experimental data, which highlighted an initial nonuniform temperature distribution. This discrepancy could stem from the sample's sidewall heating before the beginning of the experiment. In particular, it is probably related to radiative heating on the sample's surface before its placement in the arc-jet, coupled with a brief transition phase where the arm holding the sample shifts to the arc-jet's core, as supposed by McDougall in [2]. These phenomena have a greater impact on the thermocouples placed on the side, such as TC9 and TC10 (Fig. 3.1). However, it has been verified that this initial condition, even if not entirely correct, does not lead to great errors when comparing model and experimental results.

3.2 Energy equation

McDougall's model is based on the equation for unsteady heat conduction with a source term:

$$\rho_s C_p \frac{\partial T}{\partial t} = \nabla \cdot (K \cdot \nabla T) + (h_g - \bar{h}) \frac{\partial \rho_s}{\partial t}, \tag{3.1}$$

where ρ_s is the density of the solid, C_p is the specific heat, T is the temperature, K is the thermal conductivity tensor and h_g is the gaseous specific enthalpy, which is a function of pressure and temperature. In Eq. 3.1, the quantity $\bar{h}$ is referred to the partial heat of charring introduced by Milos and Chen in [7] and defined as in Eq. 3.2.

$$\bar{h} = \frac{\rho_c h_c + \rho_v h_v}{\rho_c + \rho_v}, \tag{3.2}$$

ρ_c is the final (char) density in PICA and ρ_v is the initial density of virgin PICA, both treated as constants.

The energy equation is based on three important terms: an unsteady term, proportional to the temporal derivative of T, a term that accounts for heat conduction, related to the conductivity tensor K, and a source term due to pyrolysis.

It is important to notice that, within the source term, $\frac{\partial \rho_s}{\partial t}$ is always negative, as the material is pyrolyzing, while the sign of $h_g - \bar{h}$ depends on whether the pyrolysis reaction is exothermic or endothermic, which ultimately depends on temperature. In fact, although pyrolysis is typically an endothermic reaction, at lower temperatures the species formed during this process can subsequently undergo exothermic reactions. Pyrolysis can indeed generate radicals and other chemical species that, upon interacting with one another, can release heat. In certain situations, the amount of heat released can exceed the one absorbed, resulting in a net exothermic effect. This is well represented in Fig. 3.2, where the specific enthalpy of pyrolysis ($h_g - \bar{h}$) from the TACOT model and its evolution with temperature has been plotted. Specific enthalpy values from Porous Material Analysis Toolbox based on OpenFOAM (PATO) [12] [13] were used for values of char, virgin, and gaseous specific enthalpies for the Theoretical Ablative Composite for Open Testing (TACOT) [14] [15], which is a low-density porous carbon and phenolic material, similar to PICA but open-source.

The thermal conductivity in a PICA material is represented by the tensor K and it is orthotropic. It is possible to define the tensor K using only two components: $k_{\parallel}$ and $k_{\perp}$. These, respectively represent thermal conductivity in the radial (i.e. "in plane", y or r) direction and thermal conductivity in the axial (i.e. "through the thickness" or z) direction, as defined in Figure 3.1. α is the ratio between the thermal conductivity in the radial and axial direction. Tensor K is represented by Eq. 3.3.

$$K = \begin{bmatrix} k_{\parallel} & 0 \\ 0 & k_{\perp} \end{bmatrix} = k_{\perp} \begin{bmatrix} \alpha & 0 \\ 0 & 1 \end{bmatrix}. \tag{3.3}$$

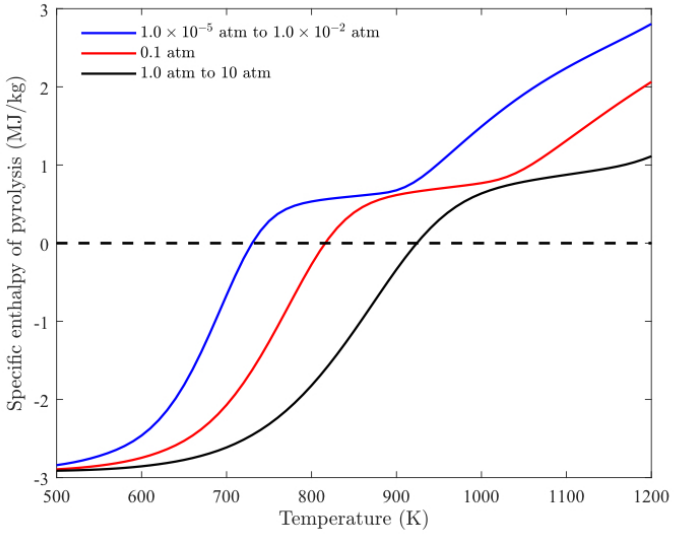

Figure 3.2: TACOT specific enthalpy of pyrolysis, $h_g - \bar{h}$. Adapted from [2].

3.3 Boundary conditions

The first boundary condition pertains to the energy balance at the exposed surface which, neglecting the contributions of pyrolysis gas and char, is taken to be:

$$K\nabla T \cdot \bar{n} + \sigma\varepsilon T_w^4 = q, \tag{3.4}$$

where T_w is the surface temperature and σ is the Stefan-Boltzmann constant. Emissivity, denoted by ε, is assumed to be 0.89 for a surface that is completely charred from [16] and [17]. q is the cold wall heat flux along the sample's surface, which representation is available from a CFD simulation performed by Milos and Chen in Figure 3.3. For a 10.16cm-diameter sample, it can be seen how the heat flux, along a surface curvilinear coordinate[1], varies only 10 % on the front face, while varying between 10 and 20 % of the

[1]Measures the distance along the surface from the stagnation point, which corresponds to the point of coordinates (0,1) in Fig. 3.3

stagnation value on the cylindrical side.

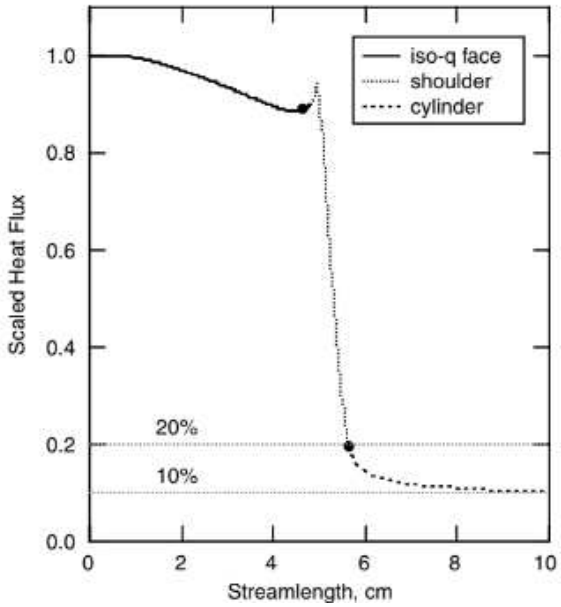

Figure 3.3: Heat flux (scaled by q_0) spatial distribution as a function of a curvilinear coordinate. Adapted from [4].

The bottom surface temperature of the material is then assumed to be constant and equal to the initial temperature, giving a Dirichlet condition:

$$T = T_0. \tag{3.5}$$

A Neumann condition is finally applied on the axis of the sample:

$$\left.\frac{\partial T}{\partial r}\right|_{r=0} = 0, \tag{3.6}$$

representing the "symmetry" of the temperature field across the sample's rotational axis, corresponding to the z-axis.

3.4 Nondimensional equations

A method to nondimensionalize the thermal model's equations has then been employed by McDougall [2]. This allows for solving the problem in nondimensional terms, managing less variables, and making the model computationally lighter. The approach is based on the definition of a reference temperature, time and enthalpy, defined as follows:

$$T_{ref} = \frac{q_0 R}{k_\perp}, \tag{3.7}$$

$$t_{ref} = \frac{\rho_v C_{p,0} R^2}{k_\perp}, \tag{3.8}$$

$$h_{ref} = C_{p,0} T_{ref}, \tag{3.9}$$

where R is the sample radius and $C_{p,0}$ is the initial specific heat of the solid.
Using these reference quantities, the following dimensionless variables are defined.

$$\tilde{r} = \frac{r}{R}, \tag{3.10}$$

$$\tilde{z} = \frac{z}{R}, \tag{3.11}$$

$$\tilde{T} = \frac{T - T_0}{T_{ref}}, \tag{3.12}$$

$$\tilde{t} = \frac{t}{t_{ref}}, \tag{3.13}$$

$$\tilde{K} = \frac{K}{k_\perp}, \tag{3.14}$$

$$\tilde{q} = \frac{q}{q_0}, \tag{3.15}$$

$$\tilde{\rho}_s = \frac{\rho_s}{\rho_v}, \tag{3.16}$$

$$\tilde{h} = \frac{h}{h_{ref}}. \tag{3.17}$$

By doing so, it is possible to re-write the energy balance of Eq. 3.1 in a nondimensional form as follows:

$$\frac{\partial \tilde{T}}{\partial \tilde{t}} = \tilde{\nabla} \cdot (\tilde{K} \cdot \tilde{\nabla} \tilde{T}) + (\tilde{h_g} - \tilde{\tilde{h}}) \frac{\partial \tilde{\rho}_s}{\partial \tilde{t}}. \tag{3.18}$$

3.5 Material model

In a pyrolyzing material, heat coming from its exterior leads to its decomposition, releasing gases that move toward the surface. What remains is predominantly carbon char and carbon fibers, intended as part of the carbon matrix not charred.

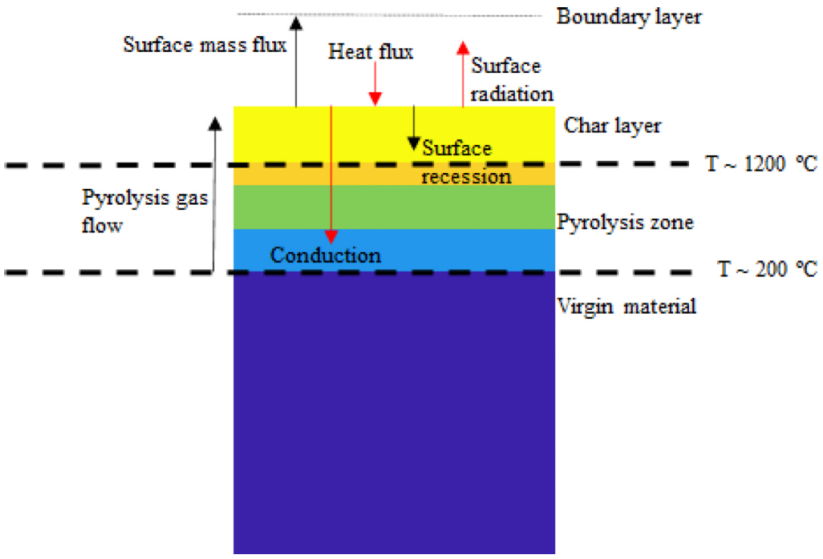

Figure 3.4: Cross-section of a pyrolyzing material in an external flow. Adapted from [2].

It is possible to conceptualize the ablative substance as an amalgamation of virgin material, gases from pyrolysis, carbon fibers, and char, as seen in Figure 3.4. In order to model the response of PICA material to external heating, two different chemical models were used by McDougall [2] and then compared. First, the competitive model developed by Torres Herrador, et. al. [5], and then the second based on the Aerotherm's Charring Material Thermal Response and Ablation Program (CMA) and Fully Implicit Ablation Thermal Response (FIAT) code [18]. For the purpose of this work, the competitive model is chosen to be part of the solver, as comparison studies performed by McDougall [2] did not highlight great differences in using one or another.

3.5.1 Competitive model

The competitive model enables the material to undergo decomposition through various reactions in a branching fashion, as visualized in Figure 3.5. This indicates that the original solid substance breaks down into various intermediary entities that can further decompose into the end products.

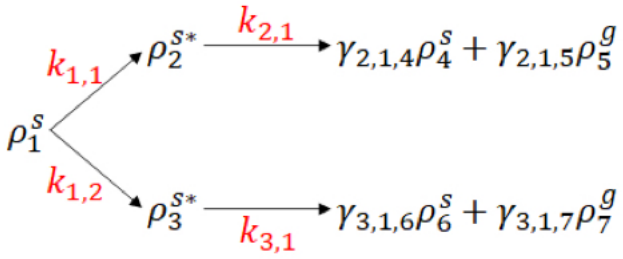

Figure 3.5: Competitive reaction mechanism. Adapted from [5].

Every element in the reaction comes with a corresponding production rate (the positive term in Eq. 3.19) as the intermediary and final entities emerge from preceding reactions. Simultaneously, there is a rate of destruction (negative term in Eq. 3.19) for each component as it undergoes decomposition via multiple competing pathways leading to the final products. This modelization with two branches represents a "low heating rate" and a "high heating rate" decomposition pathway, as proposed by Torres-Herrador, et. al. [5]. In their work, they proposed a model for PICA ablation, according to which subjecting the material to low

heating rates enhances the formation of ρ_2^{s*}. Conversely, at high heating rates, the dominant pathway leads to the production of ρ_3^{s*}. However, the overall rate of density change of component i is then given by:

$$\frac{\partial \rho_i}{\partial t} = -\sum_{j=1}^{n_i^P} k_{i,j}\rho_i + \sum_{j=1}^{n}\sum_{l=1}^{n_j^P} k_{j,l}\gamma_{j,l,i}\rho_j, \tag{3.19}$$

where n_i^P is the number of competitive reactions by which component i decays into sub-species j, $\gamma_{j,l,i}$ is the mass ratio of species j in the products and n is the number of reactants. Conservation of mass gives $\gamma_{j,l,i} = (1 - \gamma_{j,l,i-1})$. As it can be seen by the competitive reaction mechanism implemented, the initial solid, ρ_1^s, decomposes into two intermediate solids, ρ_2^{s*} and ρ_3^{s*}, without producing pyrolysis gases. These two further decompose into solid species ρ_4^s and ρ_6^s (char) along with gaseous species ρ_5^g and ρ_7^g (pyrolysis gases), respectively. The rate at which component i decomposes into j via reaction $P_{i,j}$ is modeled using a first-order Arrhenius equation:

$$k_i = A_i e^{(\frac{-\varepsilon_i}{\bar{R}T})}, \tag{3.20}$$

where A_i is the pre-exponential factor, $\bar{R}$ is the universal gas constant and ε_i is the activation energy for the reaction. The density in the solid material is then:

$$\rho_s = \sum_{i=1}^{n_s} \varepsilon_i \rho_i, \tag{3.21}$$

where ε_i is the volume fraction of component i. Rate coefficients, activation energies, and mass coefficients were utilized as specified in [2]. The ODE system solved for the rate of change of component densities can then be written in matrix form as in Eq. 3.22.

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho_1^s\\ \rho_2^{s*}\\ \rho_3^{s*}\\ \rho_4^s\\ \rho_5^g\\ \rho_6^s\\ \rho_7^g\end{pmatrix} = \begin{bmatrix}-(\mathrm{k}_{1,1}+k_{1,2}) & 0 & 0 & 0 & 0 & 0 & 0\\ \mathrm{k}_{1,1} & \text{-}\mathrm{k}_{1,2} & 0 & 0 & 0 & 0 & 0\\ \mathrm{k}_{1,2} & 0 & \text{-}\mathrm{k}_{3,1} & 0 & 0 & 0 & 0\\ 0 & \gamma_{2,1,4}k_{2,1} & 0 & 0 & 0 & 0 & 0\\ 0 & \gamma_{2,1,5}k_{2,1} & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & \gamma_{3,1,6}k_{3,1} & 0 & 0 & 0 & 0\\ 0 & 0 & \gamma_{3,1,7}k_{3,1} & 0 & 0 & 0 & 0\end{bmatrix}\begin{pmatrix}\rho_1^s\\ \rho_2^{s*}\\ \rho_3^{s*}\\ \rho_4^s\\ \rho_5^g\\ \rho_6^s\\ \rho_7^g\end{pmatrix}. \tag{3.22}$$

7 species have been considered as seen to be overall representative of the physical phenomenon. The reaction rate coefficients k_i have also been nondimensionalized using the reference quantities as follows:

$$\tilde{k}_i = (A_i T_{\mathrm{ref}}) \exp\left(-\frac{\tilde{\varepsilon}_i/R}{\tilde{T}(T_{\mathrm{ref}})+T_0}\right). \tag{3.23}$$

The dimensional system represented in Eq. 3.22 has then been nondimensionalized.

3.6 Implementation in MATLAB

Experiments were conducted on a PICA sample of which geometry is symmetrical with respect to the scaled axial coordinate axis, visualized in Figure 3.6.

The vertical axis corresponds to the axis of rotation and the shoulder radius is 12.5% of the sample radius R. All samples that were tested with a heat condition of q_0 = 169 W/cm^2, which will be our reference case, had a 5.08 cm radius.

The thermal model has then been implemented in *MATLAB* [2].
In order to address the mathematical problem, the *MATLAB*'s *PDE Toolbox* has been used to model conduction through the material. It incorporates the possibility of building a triangular elements mesh (visualized in Figure 3.7) for a given geometry and then solves partial differential equations with a finite element method.

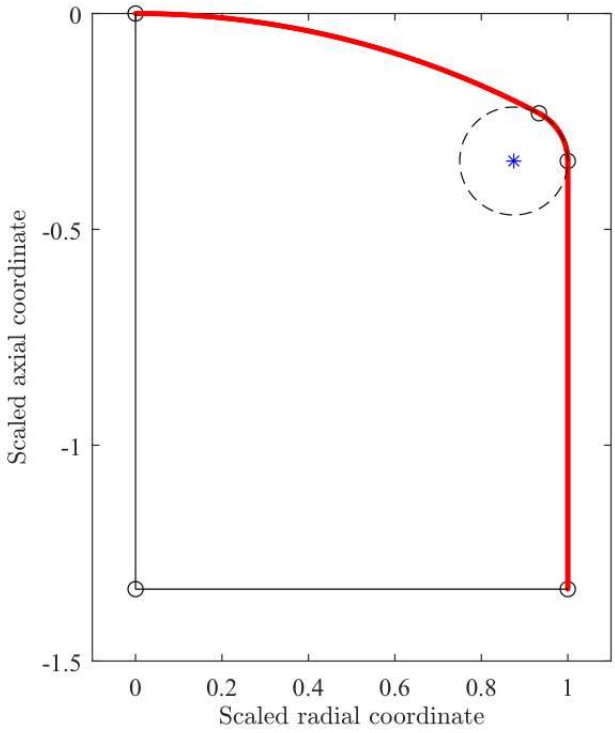

Figure 3.6: Iso-q geometry, with coordinates scaled by sample radius R. Adapted from [2].

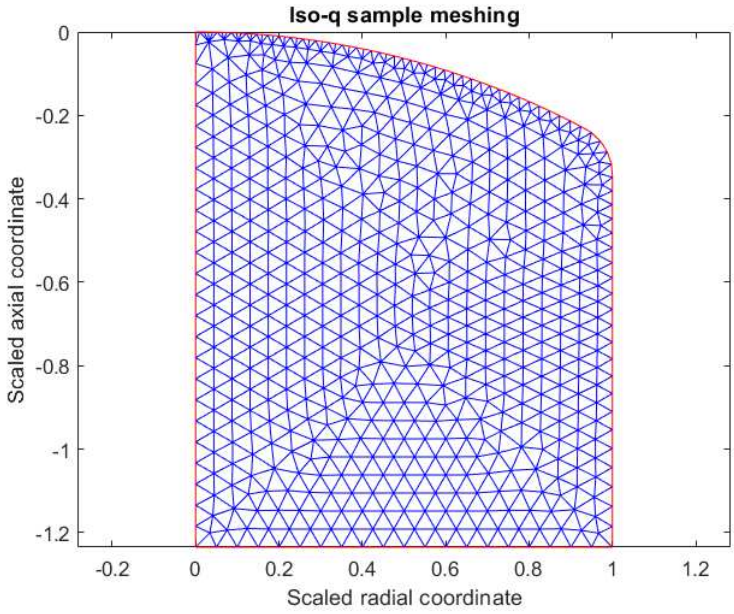

Figure 3.7: Triangular mesh for 2499 nodes and 1186 elements.

The finite element discretization originally used by McDougall has been changed with the purpose of building the inference model. The one that will be used can be visualized in Figure 3.7 and consists of 2499 nodes and 1186 elements, being coarser than the one previously used. A coarser mesh gives the possibility to run the analysis much faster but can affect negatively the accuracy of the model. Given that, a grid convergence study was performed, presented in Chapter 4, Section 4.1.

A second-order splitting method to improve computational efficiency was used. By doing so, for each time step, a PDE is solved for half a time step accounting solely for heat conduction (with appropriate boundary conditions applied):

$$\frac{\partial \tilde{T}}{\partial \tilde{t}} = \tilde{\nabla} \cdot (\tilde{K} \cdot \tilde{\nabla} \tilde{T}), \tag{3.24}$$

giving an intermediate temperature field, $\tilde{T}^* = \tilde{T}^*(\tilde{r}, \tilde{z}; \tilde{t}_{n+\frac{1}{2}})$.
Then, this is taken to be the initial condition to solve the system of ODEs, defined by the competitive model, for the seven phases and temperature, at each element over a full time step from $\tilde{t}_n$ to $\tilde{t}_{n+1}$, according to Eqs. 3.19 and 3.21. Finally, the temperature field obtained in this manner is advanced over a half time step accounting for conduction and boundary conditions, using again Eq. 3.24. The obtained solution has then been interpolated at thermocouple locations, visualized in Fig. 3.1.

Due to limited space for installation, it was not possible to position thermocouples at all ten designated locations within a single model. For our cold wall heat flux reference case, the sample has been tested by using two different types of thermocouple placements, called Placement A and Placement D. Placement A had thermocouples at locations 1 to 5, while Placement D at locations 5 to 10. Their coordinates are presented in Table 3.1.

TC Number	TC Placement	Axial depth (cm)	Radial distance (cm)
1	A	0.381	0.000
2	A	0.762	0.000
3	A	1.143	0.000
4	A	1.524	0.000
5	A and D	3.048	0.000
6	D	2.286	0.000
7	D	2.286	2.540
8	D	2.286	3.810
9	D	2.286	4.445
10	D	3.048	4.445

Table 3.1: Thermocouples coordinates from Milos and Chen [4].

3.7 Model sensitivity to pyrolysis

To assess the importance of particular physical processes, McDougall [2] sequentially introduced individual terms of the energy equation and compared them with experimental data. The key discovery is that conduction itself failed to offer a basis for comparison, even for brief periods and close to the surface. However, upon integrating radiation, the disparity between models incorporating pyrolysis and those without was negligible. Throughout these simulations, it has been shown that conductive and radiative effects dominate early on during the exposure to the flow and that pyrolysis begins to affect some thermocouples after 10 or 20 seconds of simulation. In particular, McDougall [2] found that pyrolysis begins to affect TC1 at around 10 seconds while beginning to affect TC2, TC9, and TC10 at around 20 seconds. This discovery yields attention to which thermocouples to use when testing the model's predictions against experimental data, as the model only accounts for a situation where very little pyrolysis is occurring in the material.

Chapter 4

Surrogate Model

The computational model presented in the previous Chapter allows the calculation of the temperature evolution over time at specific thermocouple locations for a PICA ablative material. The material is defined in the model through its virgin density, thermal conductivity tensor, and volumetric thermal capacity (product of virgin PICA density and initial specific heat of the solid PICA). However, through the nondimensionalizing process described in Section 3.4, it is possible to operate a nondimensional problem that allows the calculation of nondimensional temperatures in nondimensional times. Nondimensional temperatures are then transformed in dimensionals through the definition of T_{ref} via Eq. 3.12, while nondimensional times in dimensionals through t_{ref} and via Eq. 3.13. In particular, this nondimensionalization process allows for solving a problem that only relies on one variable, α, which is the ratio of thermal conductivities:

$$\alpha = \frac{k_{\parallel}}{k_{\perp}}. \tag{4.1}$$

To perform the inference process, which will require running the solver for a large number of iterations, it becomes necessary to build a surrogate model for a matter of computational efficiency. A surrogate is nothing less than a library in which model results (i.e. nondimensional temperatures in nondimensional times) are already stored and ready to be used by the inference framework, without the need to run the solver each time a new solution is required.

Given that α is the only variable of the nondimensional problem defined in *MATLAB*, the surrogate will consist of a library of solutions for every chosen value of α. By building the surrogate, solutions will be already available for different values of α and ready to be used by the framework.

However, as the solutions to be stored in the surrogate will be calculated through the computational model, it is necessary to first decide a mesh size and a time step for solving the differential equations involved, presented in Section 3.6. In fact, these will affect the accuracy of the solutions to be stored within the surrogate and so, of the inference process itself. To define a proper mesh size and time step, the solver has been run several times using different settings and comparing each time model and the experimental results, as explained in this Chapter in Section 4.1. By doing so, specific values have been chosen aiming to minimize the error between model and experimental results, as well as the computational time required by each analysis.

Another crucial aspect that needs to be addressed before building the surrogate is deciding for which values of α the solutions will be calculated and stored, as this will also have a role in the accuracy of the inference process, as detailed in Section 4.2.

Before defining the surrogate, the thermal model itself has been subject to changes. In fact, the solver has been dismembered into two major parts: a part that contains static[1] parameters of the analysis (such as sample geometry, heat flux spatial distribution, thermocouples placements, boundary, and initial conditions, etc.), and a part that contains the solver core, in which mesh and equations are included. By doing so, the computational time of each analysis has been reduced, given that the "static" part of the solver has been run only the first time, creating variables that remained unchanged throughout the following analysis.

[1] Parameters that are not subject to changes between consecutive analyses.

4.1 Mesh & Time step

The surrogate building process consists in nothing more than running the thermal solver for specific values of α and storing the solutions. In order to decide the mesh size and the time step to use for the solver, an optimization study has been conducted. This has been necessary to investigate how changing the mesh size and the PDE's time step affects the model accuracy when comparing its results with experiments. In particular, comparisons were always performed utilizing experimental data for a cold wall heat flux at a stagnation point of 169 W/cm^2, available from the work of Milos and Chen [4]. Analyses have been performed by using six different mesh sizes and five different time steps, as reported in Table 4.2 and 4.3. Only TCs from Placement D have been utilized.

Throughout this study, nominal values for the thermal model inputs have been used, as defined in [2] and reported in Table 4.1.

$k_\perp$, W/mK	q0, W/cm^2	rhoCp, J/m^{3}K	α
0.358	169	3920000	2

Table 4.1: Thermal analysis inputs nominal values.

Case	Nodes	Elements
1	2499	1186
2	3568	1715
3	4447	2150
4	5909	2872
5	8741	4276
6	13394	6585

Table 4.2: Mesh sizes tested.

For each mesh size tested (Case 1 to 6 of Table 4.2), 5 different nondimensional time steps (reported in Table 4.3) have been used, for a total of 30 simulations.

Nondimensional	Dimensional (s)
0.0014	4
0.0007	2
0.0004	1
0.0003	0.8
0.0002	0.5

Table 4.3: Timesteps used for analysis.

In this Table, dimensional time steps have been calculated applying Eq. 3.13 with nominal values available from Table 4.1.

So, for every case reported in Table 4.2, five simulations with five different time steps have been performed, where every analysis simulated 40 seconds of experiment. From each analysis, ten values of dimensional temperatures have been saved, in order to perform comparisons with experimental data. In particular, temperatures from specific thermocouples of Placement D at specific instants in time have been taken, as reported in Table 4.4. As an example, starting from the first values, the temperature at thermocouple 8 was taken at 4 seconds, while temperature in thermocouple 9 at 8 seconds, and so on.

Time (s)	4	8	12	16	20	24	28	32	36	40
TC	8	9	10	8	9	10	7	8	9	10

Table 4.4: Temperatures from TCs saved, at specific time instants.

4.1.1 Results

Saving specific temperatures for each analysis allowed for a relevant comparison with experimental data, as they are available at the same time instants as the ones in which model results are saved. To define their accuracy with respect to experimental data, two indicators were used: Mean Absolute Error (*MAE*) and Ecludian Norm (*L2*). These, have been calculated as follows:

- *MAE* (Mean Absolute Error) has been calculated as the absolute difference between experimental and model temperatures, at the same TC locations and at the same times indicated in Table 4.4. Then, of the 10 difference values available, the arithmetic mean has been considered;

- *Eucledian Norm (L2)* has been calculated by considering the ten available differences, of which each one has been squared and summed up. Then, the square root of this value has been taken.

Calculating *MAE* and *L2 Norm* for each one of the 30 analyses, allowed to withdraw an important consideration:

- From the most imprecise analysis (relative to 2499 nodes and 0.0014 time step) to the most precise (relative to 13394 nodes and 0.0002 time step), MAE only diminishes by 5.7 % while L2 norm by 2.5 %. Also, the run time of the solver[2] increases from 375 seconds to 6302 seconds, when going from the most imprecise to the most refined analysis.

It is then clear how refining the mesh and reducing the time step, within the intervals before specified, does not seem to produce a huge difference in terms of results precision when comparing model results to experimental values. Figs. 4.1, 4.2, 4.3 and 4.4 show the negligible variation of MAE and L2 norm within a specific mesh and time step case.

To confirm this hypothesis, another run of the model was performed. The objective was to perform a very accurate analysis, with a very refined mesh and very short time step, to verify how much model results would gain precision. This analysis was performed with 105482 nodes and 52615 elements, along with a timestep of 0.0002. This simulation showed a decrease of MAE and L2 norm, with respect to the most imprecise run of 2499 nodes and 0.0014 time step, of only 6 % and 2.65 %, respectively. This MAE and L2 norm results are plotted in Fig. 4.1 and 4.3 with a dashed line, as reference. In light of this, it can be said that to obtain a significant difference in terms of MAE and L2 Norm, an unreasonable number

[2]Run time only refers to the run time of the solver "core" (hence, without the static part), introduced in the introduction of this Chapter.

of mesh nodes/elements and a really short timestep would be needed, massively increasing the run time of the solver.

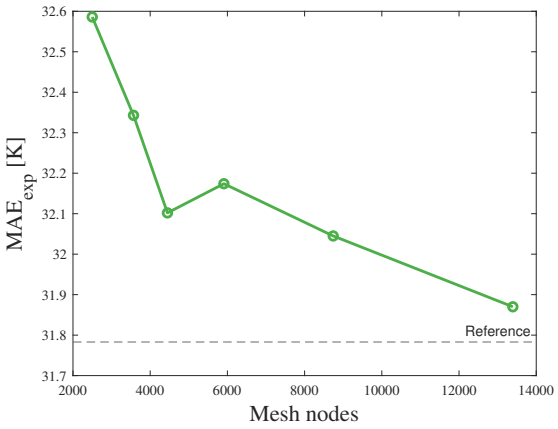

Figure 4.1: MAE variation with the number of mesh nodes, relative to a 0.5 s time step.

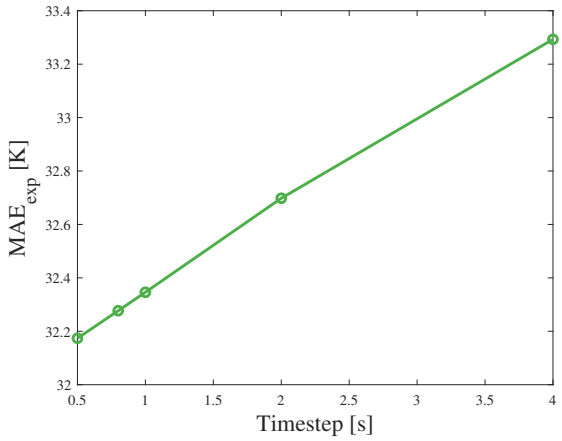

Figure 4.2: MAE variation with dimensional time step, relative to 5909 mesh nodes.

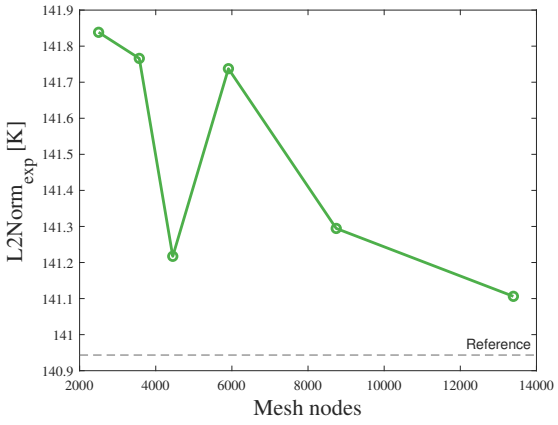

Figure 4.3: L2Norm variation with the number of mesh nodes, relative to a 4 s time step.

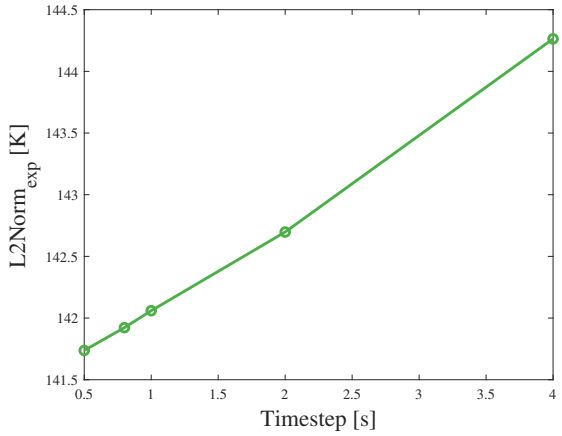

Figure 4.4: L2Norm variation with dimensional time step, relative to 5909 mesh nodes.

Taking this into consideration and for a matter of computational efficiency, a mesh size of 2499 nodes and 1186 elements, represented in Figure 3.7, has been selected and will be used for the thermal solver when building the surrogate model. Even though the mesh chosen was the coarsest of the ones tested, model results still appeared to be representative of experimental data. For what concerns the time step, the study did not show the necessity of using a specific value between 0.0014 (4 seconds) and 0.0002 (0.5 seconds), considering the negligible differences in terms of MAE and L2 norm showed when using either one of the time steps tested. However, a major dissimilarity concerns the run time, which increases

from 375 s to 1446 s when going from 0.0014 to 0.0002 and considering a 2499 nodes' mesh size. Nevertheless, the time step choice has been addressed by taking into account another crucial aspect that will be detailed in the next section.

4.2 α Discretization

Now that a mesh size has been decided, it is important to define for which values of α the surrogate will be developed. Considering the nominal value of α = 2 from Milos and Chen [4], a surrogate incapsulating model results for values of α between 0.1 and 5 has been built, as it enables to have a wide selection of results available to be used during the inference process. In particular, in the above-mentioned α range, the following intervals of α have been taken:

First Value	Increment	Last Value
0.1	0.1	1
1	0.05	3
3	0.1	5

Table 4.5: α values for which the solutions are determined.

In particular, as indicated in Table 4.5, solutions between 1 and 3 were obtained every $\alpha = 0.05$, instead of every $\alpha = 0.1$. This choice is led by the fact that a value of α comprised between 1 and 3 is taken to be more probable to be representative of the real material physics, so it has been decided to build a surrogate with more available solutions in that interval. In fact, having more results in that range also allows for increasing the accuracy of the process, as the query of the surrogate, which will be performed by the inference framework, will also comprehend an interpolation process when required. Specifically, when a solution will be desired for a value of α for which a solution is not immediately available in the surrogate, the code will perform a linear interpolation between the solutions available for the two values of α contiguous to the one needed by the framework. For example, if a solution is needed for $\alpha = 1.13$, it will be obtained by interpolating the solutions available for $\alpha = 1.10$ and $\alpha = 1.15$ (this applies for the whole 0.1-5 interval). In this sense, having

more solutions stored within the interval α = 1-3 will allow getting more precise results, when the interpolation process is involved.

As previously said, the solutions for α values in Table 4.5 have been obtained by running the solver with the mesh specifics described in Section 4.1.

The PDEs time step has been set as 0.5 seconds (0.0002 in nondimensional time). To understand why this time step has been chosen, let's recall that all the analyses presented in Section 4.1 simulated 40 seconds of arc-jet material's exposition. Nevertheless, the thermal model solves the problem in nondimensional times, which are calculated from the dimensional times divided by the reference time (see Eq. 3.9). As it can be seen from Eq. 3.13, in the time nondimensionalizing process, the quantities $\rho_v \cdot C_{p,0}$ (later always referred to as rhoCp) and $k_\perp$ are involved. The key fact is that these quantities could later be inferred and assume statistical distributions, more or less centered in their nominal value, hence not assuming a static value throughout the inference process. This means that the reference time could possibly change at every iteration of the inference process, as different values of rhoCp and $k_\perp$ are withdrawn from their respective distributions. It is then fundamental to be sure that the surrogate has solutions available even for the largest possible nondimensional time span, which is related to the smallest possible t_{ref} by Eq. 3.13. Following this idea, and by taking into consideration a maximum variation of rhoCp and $k_\perp$ of two times their nominal values, the smallest t_{ref} has been calculated:

$$t_{ref} = \frac{\rho_v C_{p,0} R^2}{4 \cdot k_\perp} = 706.4322\ s. \tag{4.2}$$

By using this value of reference time, the dimensional time interval from 0 to 40 seconds in which we want to obtain our solutions will now correspond to a very large nondimensional time interval. At this point, in order to maintain a high accuracy of model results with respect to experiments, it has been necessary to utilize the smallest dimensional time step between the ones presented in Table 4.3, of 0.5 seconds. By doing so, the solver solution will always be made of 81 temperatures, as the simulation is run for 40 seconds and temperatures are calculated every 0.5.

Chapter 5

Fundamental Concepts of Statistics

The objective of this Chapter is to give the reader a brief but yet comprehensive overview of some statistical concepts and tools that have been later used during the inference process, useful for the interpretation of its results. Along with these, the objective is also to give a clear representation of what inference is and how it works.

5.1 Bayesian Inference

> *"Bayesian inference is a statistical method that uses Bayes' theorem to update the probability of a hypothesis as new evidence or information is obtained."*
> *Analytics Vidhya*

Usually, the presence of the following fundamental aspects can be easily distinguished in every Bayesian inference setting:

- **Prior**: the prior probability, often simply called the "prior", reflects what is known or believed about a variable before considering the new evidence. It represents our initial belief/hypothesis about the distribution of a parameter;
- **Likelihood**: refers to the probability of observing the data given a model or hypothesis. It measures how well the model explains the observed data;
- **Posterior**: the posterior probability is the probability of the hypothesis given the observed data. It combines the likelihood and the prior to give the new parameter's distribution, after considering the evidence. It constitutes the actual output of the analysis.

So, Bayesian inference will be the crucial statistical tool that will help refine PICA material's properties. These, which are firstly declared as constant, will be re-assessed using the experimental data available and the computational model previously presented, in order to obtain a posterior distribution of the properties finally accounting for uncertainties. The three aspects previously presented are mathematically related through Bayes' Theorem.

5.1.1 Bayes' Theorem

Before delving into the actual Theorem formulation, an example from [19] will help understand how Bayes' Theorem works, and how powerful this tool is.

Example

You wake up one morning and you feel a little bit sick. No particular symptoms, just not at your 100%. So, you go to the doctor and she also does not know what is going on with you, so she suggests they run a battery of tests. After a week goes by, the results come back, turns out you tested positive for a very rare disease that affects only 0.1% of the population.

How certain is it that you have the disease?

The doctor replies saying that the accuracy of the test is to correctly identify 99% of the people that have the disease and only incorrectly identify 1% of people who don't.
At this point, most people would think they have a 99% possibility to actually have the disease, because that is the accuracy of the test. But that is.. incorrect. Let's use Bayes to get some perspectives.

Bayes' Theorem can help by giving the probability that some hypothesis (that you actually have the disease) is true, given an event (that you tested positive for the disease). This probability can be defined as $P(H|E)$, where H stands for Hypothesis and E for Event, whereas "|" is used to symbolize "given that". In order to calculate that, and so, to find out your actual probability of having the disease, you will need to take into account several aspects. Among these, two in particular have a key role:

1. Prior $P(H)$, defines how likely you thought it was that you had this disease before the event occurred, so, before you got the results;

2. Likelihood $P(E|H)$, which is the probability of the event, given that the hypothesis is true. This is the probability that you would test positive if you had the disease;

Along with them, you will need to define the probability of not having the disease, that can be written as $P(-H)$ and the probability of being falsely identified, $P(E|-H)$.
All these come together when defining Bayes' Theorem:

$$P(H|E) = \frac{P(H) \cdot P(E|H)}{P(H) \cdot P(E|H) + P(-H) \cdot P(E|-H)}. \tag{5.1}$$

which is essential to define the posterior probability $P(H|E)$.

Estimating the prior probability $P(H)$ can sometimes be tricky. For example, in this case, a reasonable starting point could be the frequency of the disease in the population: 0.1%. Along with this, the other probabilities can be applied to Equation 5.1:

- $P(E|H) = 0.99$, as there is a 99% chance to be tested positive, if you have the disease;
- $P(-H) = 0.999$, which is the probability of not having the disease among the population;
- $P(E|-H) = 0.01$, which is the probability of being falsely identified, if you do not have the disease.

Combining these data gives out a probability of having the disease which is: $P(H|E) = 9\%$. This, which is the ultimate output of the Bayes' Theorem, is called posterior probability. In fact, as the name suggests, it consists on an update from the prior probability (of having the disease inside the population) through evidence and data (the fact that you tested positive). It is also important to keep in mind that posteriors can always be updated when observing new evidence.

As an example, let's say that you then go to another doctor to be tested in another laboratory and the test comes back positive another time. This evidence would update again your probability of having the disease, which can be simply calculated through Eq. 5.1, except that this time the prior probability of having the disease would not be 0.001, but it would be the posterior probability coming from the previous analysis, 0.09. Applying Eq. 5.1, you will find out a new posterior probability of 91%. This use of the Bayes' Theorem is typical of an inference analysis, where the posterior probability keeps getting updated by adding new evidence and data.

In the context of material properties estimation, the Bayesian inference process can be applied in a manner analogous to the provided Bayes' theorem example. Let us consider a material property, named X, which is nominally stated as a constant. This represents our prior knowledge (the prior) about the material's behavior in the absence of additional experimental data. However, similar to a medical diagnosis scenario, we possess experimental data providing new evidence on the actual characteristics of the material.

Within the inference framework that will be defined in *MATLAB*, presented in the next Chapter, we will employ our computational model to simulate the experiments and compare its results against the actual experimental data. This comparison aligns with the use of the likelihood function in Bayes' theorem, where our computational model plays the role of the diagnostic test, and the experimental data represent the test result (the event E). Through this comparison, we are able to update our prior knowledge about property X, redefining it to more accurately reflect the experimentally observed behavior of the material. This updated result is known as the posterior probability of the property X's value.

Ultimately, just as Bayesian inference allows us to update the likelihood of a diagnosis based on test outcomes, similarly, we update our understanding of material property X. This approach enables us to incorporate experimental uncertainties and material variations into the model, providing a more accurate and reliable estimation of the material properties that will be used for further analysis and engineering applications.

5.1.2 Prior probabilistic distributions

In Bayesian inference, prior distributions play a pivotal role in encoding our pre-existing knowledge or beliefs about the parameters of the model before observing any data. They represent our initial understanding or assumptions about the parameters' values. In particular, two probability distributions will later be used during inference the inference process: uniform and log-normal.

Uniform distribution

Uniform distributions, also referred to as rectangulars, are a family of symmetrical distributions. This distribution is characterized by its simplicity and fundamental nature in representing a scenario where all outcomes are equally likely, and it is always bounded between two values, that we can call a and b. Mathematically, a uniform distribution can be defined in two forms: discrete and continuous.

In the case of a discrete uniform distribution, which is applicable when dealing with a finite set of equally likely outcomes, the probability mass function is given by:

$$P(X = x) = \frac{1}{N}, \tag{5.2}$$

defined for x = 1,2,3,...,N. The outcomes range from 1 to N. $P(X = x)$ represents the probability of the random variable X taking a particular value x.

In contrast, a continuous uniform distribution is used when the outcomes are spread across an interval. In this case, a probability density function (PDF) is defined, for an interval [a, b], as:

$$f(x) = \begin{cases} \frac{1}{b-a} & \text{for } a \leq x \leq b, \\ 0 & \text{for } x < a \text{ or } x > b. \end{cases} \tag{5.3}$$

This equation defines that the distribution is constant over the interval [a, b] and zero otherwise, as it can be also visualized in the next Figure.

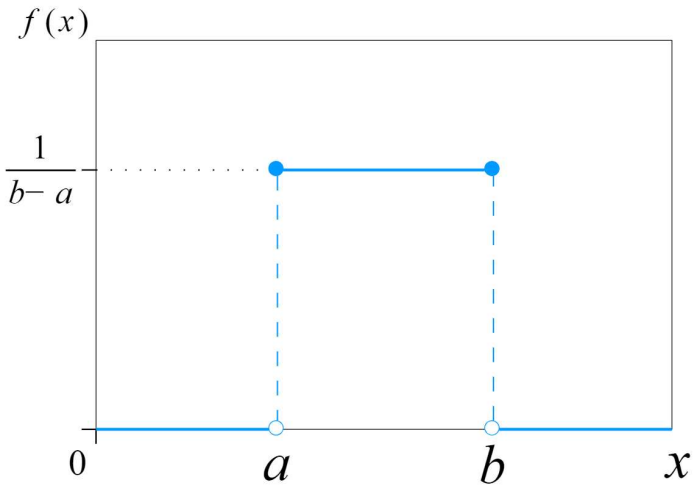

Figure 5.1: Continuous uniform distribution defined for an interval [a, b]. Adapted from Wikipedia.

In the discrete case, each outcome has an equal chance of occurring, and in the continuous case, every point in the interval [a,b] has the same density. This characteristic makes the uniform distribution a fundamental reference point in probability and statistics. In practical applications, the uniform distribution is often used to model scenarios where there is no preference or bias towards any outcome within a specified range.

Log-normal

The log-normal distribution is a key concept in probability and statistics, particularly useful in representing variables whose values are positively skewed and confined to the positive real line. The log-normal distribution is defined for a random variable whose natural logarithm (log) is normally distributed. So, if X is a random variable with a log-normal distribution, then $log(X)$ (the natural logarithm of X) follows a normal distribution.

Mathematically, if $Y = log(X)$ is normally distributed with mean μ and variance σ^2, then the probability density function (PDF) of X is given by:

$$f(x;\mu,\sigma) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln x - \mu)^2}{2\sigma^2}\right), \tag{5.4}$$

for $x \in [0,+\infty)$, μ real number and $\sigma > 0$. For a normal distribution of $Y = log(X)$, it is always true that the 68% of data will be inside one standard deviation σ from the mean value μ, and 95% will be inside 2σ. However, this exact relation does not apply to the log-normal distribution.

It is possible to determine the mean and the variance values of the log-normal distribution by knowing the values of mean and variance of the normal distribution of $log(X)$:

$$\text{Log-normal Mean}: \quad \mu_X = e^{\left(\mu + \frac{\sigma^2}{2}\right)}, \tag{5.5}$$

$$\text{Log-normal Variance}: \quad \sigma_X^2 = \left(e^{\sigma^2} - 1\right) e^{2\mu+\sigma^2}. \tag{5.6}$$

It is also possible to calculate the standard deviation and mean values of the underlying normal distribution by using the ones for the log-normal. The mean of the normal distribution can be calculated by knowing the median value (median) of the log-normal as:

$$\mu = \log(\text{median}). \tag{5.7}$$

The standard deviation σ can be computed by first defining the percentiles, for example the 2.5th and 97.5th percentiles[1], of the log-normal distribution. By doing that, z-scores[2] can be calculated:

$$z_{\text{lower}} = \Phi^{-1}(A) = \Phi^{-1}(0.025), \tag{5.8}$$

$$z_{\text{upper}} = \Phi^{-1}(B) = \Phi^{-1}(0.975), \tag{5.9}$$

where A and B are, respectively, the 2.5th and the 97.5th percentiles. In these expressions, Φ^{-1} represents the inverse cumulative distribution function (CDF)[3] of the normal distribution. Now, the standard deviation of the normal distribution σ can be computed as follow:

$$\sigma = \frac{\log(\text{upperbound}) - \mu}{z_{\text{upper}}}, \tag{5.10}$$

where *upperbound* is the parameter's value which defines the 97.5th percentile.

As it can be seen from Figure 5.2, log-normal distribution is characterized by its asymmetry and heavy right tail. Unlike the normal distribution, which is symmetrical and can take on negative values, the log-normal distribution is always positive and skewed to the right. It is also not uncommon to see log-normal distributions, as well as normal ones, to have right and left tails truncated. This is useful in applications where we know for sure that our parameter cannot assume values greater than or less than specific values, that can be used to limit the curve. In our context, the log-normal distribution, which will be applied to the variable α as detailed in Chapter 7, will be truncated from 0.1 to 5, which is the interval of α taken into consideration in the surrogate.

[1]The 2.5th percentile represents a value below which 2.5% of the data falls, while the 97.5th percentile represents a value below which 97.5% of the data lies, typically used in statistics to define the lower and upper bounds of a 95% confidence interval.

[2]z-scores are used as points of reference of the distribution, as they defined the number of standard deviations between the mean value and the value for which the z-score is calculated (for example, the STD will be at 1 z-score, while the mean at 0).

[3]The Cumulative Distribution Function (CDF) represents the probability that a random variable will take a value less than or equal to a specific value, essentially quantifying the cumulative probability up to that point in the distribution.

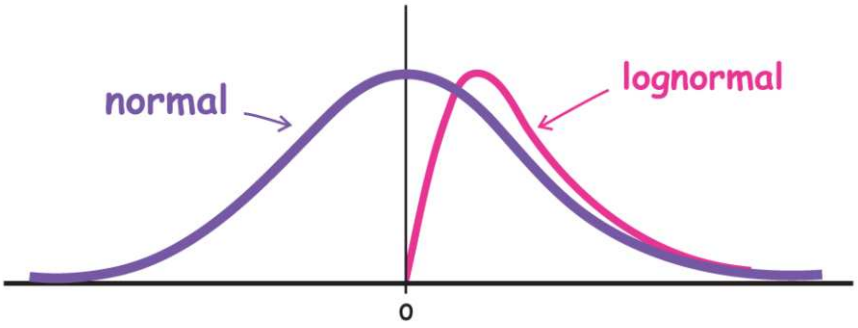

Figure 5.2: Visual difference between a log-normal and a normal distribution. Adapted from CalcWorkshop.

5.2 Markov Chain Monte Carlo & AIES

Moving from the foundational concepts of Bayesian inference, where we updated our beliefs about a material property with new data, we now delve into the practical tools that make such computational updates feasible and efficient when inference is performed within a software. This is where the Markov Chain Monte Carlo (MCMC) and, by extension, the Affine Invariant Ensemble Sampler (AIES) become pivotal.

In the context of our framework, MCMC will serve as the workhorse for Bayesian updating, facilitating the sampling from posterior distributions. The method does so by generating sequences of samples that eventually converge to the target posterior distribution, thus mimicking the Bayesian process of updating prior knowledge with new evidence. The process involves proposing new samples and deciding whether to accept or reject them based on their likelihood, given the observed data and prior information.

AIES, a specific implementation within the MCMC framework, excels by employing a group of "walkers", parallel chains of samples that simultaneously traverse the parameter space. This collaborative movement ensures a thorough and unbiased exploration, effectively sidestepping the pitfalls of local optima that single-chain MCMC methods might fall into.

In order to better understand these concepts, let's go back to the example concerning the estimation of our material property, X. Initially, we only have a rough estimate for X based on prior knowledge, which forms our prior distribution. The MCMC algorithm begins by taking an initial sample of X from this prior distribution.

It then inputs this sample into our computational model (which, in our framework, will be divided into the surrogate and the later dimensionalizing process) which simulates the material's behavior under experimental conditions. The output of this model is then compared to actual experimental observations. If the model's output closely matches the experimental data, the sample of X is considered to be a good estimate and is "accepted." If not, the algorithm may either discard the sample or accept it with a certain probability that depends on how closely the simulated results match the experimental data. The MCMC process iteratively continues this procedure, taking new samples of X, running the model, comparing with the experimental data, and accepting or rejecting each sample based on how well it predicts the observed results. Over many iterations, this procedure constructs a chain of accepted X values which forms the posterior distribution.

In applying the MCMC algorithm to our material property X, we will not actually rely on a single sampling chain but rather implement multiple chains that operate in parallel. These parallel chains are crucial as each one can independently explore different parts of the parameter space. While each chain follows the described mechanism (drawing samples from the prior distribution and testing them against the computational model and experimental data) the existence of multiple chains allows for cross-validation of the results. By employing several parallel chains, the posterior distribution of the parameter will be calculated using a combination of the accepted samples from all the chains. The presence of multiple parallel chains also increases computational efficiency. While a single chain might require an extended period to adequately explore the parameter space, multiple chains can do this work more swiftly and thoroughly.

The mean of the posterior distribution gives us the most probable estimate of X after considering the experimental evidence, while the variance provides insight into how certain we are about this estimate. From now on, X will not refer anymore to the material property, only used until now as an example, but will change its significance based on the context.

5.2.1 Gelman-Rubin statistics

The Gelman-Rubin statistic ($\hat{R}$) is a diagnostic measure of convergence for Markov Chain Monte Carlo (MCMC) simulations [20].

The statistic is crucial in the context of Bayesian analysis, as it provides a quantitative method to evaluate whether the chains have sufficiently explored the parameter space and if they have converged to the target posterior distribution.

This statistic is particularly useful because it does not require the samples to be independent and identically distributed (i.i.d.), a condition that is often not met in the initial phases of an MCMC simulation. In fact, in typical MCMC runs, the initial samples may be highly correlated and may not reflect the true underlying distribution as the chains are still 'burning in' and have not yet reached their stationary distribution. The Gelman-Rubin statistic circumvents this issue by comparing the variation within and between chains, allowing for an assessment of convergence even when the assumption of i.i.d. samples is violated.

To calculate the Gelman-Rubin statistic, one first computes the within-chain variance (W), which reflects the average variance of samples within each chain, and the between-chain variance (B), which measures the variance between the means of each chain, respectively reported in Eqs. 5.11 and 5.12:

$$W = \frac{1}{m}\sum_{j=1}^{m}\frac{1}{n-1}\sum_{i=1}^{n}(x_{i,j}-\bar{x}_j)^2, \tag{5.11}$$

$$B = \frac{n}{m-1}\sum_{j=1}^{m}(\bar{x}_j-\bar{x}_{ovll})^2, \tag{5.12}$$

where $x_{i,j}$ is the i-th sample in the j-th chain, $\bar{x}_j$ is the mean of the j-th chain, $\bar{x}_{ovll}$ is the overall mean of all samples (i.e. the mean of the m $\bar{x}_j$ values of chain means), m is the number of chains, and n is the number of iterations per chain.

The estimated variance of the target distribution ($\hat{\sigma}^2$) is then a weighted average of W and B:

$$\hat{\sigma}^2 = \left(1-\frac{1}{n}\right)W + \frac{1}{n}B. \tag{5.13}$$

Finally, the Gelman-Rubin statistic is calculated as:

$$\hat{R} = \sqrt{\frac{\hat{\sigma}^2}{W}}. \tag{5.14}$$

Values of $\hat{R}$ close to 1 indicate that convergence has likely been achieved, as the between-chain and within-chain variances are similar, suggesting that the chains are sampling from the same distribution. In fact, a Gelman-Rubin statistic near 1 implies that the different chains have become indistinguishable from one another in terms of their distributional properties. This homogeneity among chains is a strong indication that they are all exploring the same underlying posterior distribution and that the Markov Chain Monte Carlo process is generating consistent and reliable samples across all chains.

5.3 Kullback Leibler Divergence

> *"The Kullback-Leibler divergence metric (also known as relative entropy) is a statistical measurement from information theory that is commonly used to quantify the difference between one probability distribution from a reference probability distribution (for example, between distributions P and Q)." Towards Data Science*

The Kullback Leibler Divergence (later also named only KL) was introduced by Solomon Kullback and Richard Leibler in 1951 and has since become a fundamental concept in information theory, statistics, and machine learning.

In our context, where one of our objectives will be finding which thermocouples can be more informative about the definition of a certain parameter (for example, the material property X), the Kullback Leibler Divergence is a tool of paramount importance, as it defines how much information we can gain about that parameter using some thermocouples rather than others. This information gain is related to how much the posterior probability distribution of the inferred parameter narrows, with respect to the prior, during the inference process. It is important to say that, even though the KL could possibly be visualized as a distance between two distributions, it is not actually a metric, as it is not symmetric (i.e. the relative entropy from Q to P is not the same as the one from P to Q).

Mathematically, it can be defined for discrete probability distributions as well as for distributions of a continuous random variable. Taking two discrete probability distributions P and Q defined on the same sample space, X, the relative entropy from Q to P is defined as:

$$D_{KL}(P \parallel Q) = \sum_{x \in X} P(x) \log \left(\frac{P(x)}{Q(x)} \right). \tag{5.15}$$

For distributions P and Q of a random continuous variable, we have:

$$D_{KL}(P \parallel Q) = \int_{-\infty}^{\infty} p(x) \log \left(\frac{p(x)}{q(x)} \right) dx, \tag{5.16}$$

where p(x) and q(x) are probability densities of P and Q. In Bayesian Inference analysis, Q plays the role of prior distribution probability while P of the posterior distribution probability, so that the divergence measures the information gain by revising one's belief considering new evidence or data.

Chapter 6

Inference Framework

In this Chapter, all the parts that compose the inference framework and their interconnections will be discussed, from the definition of the prior distributions to the discrepancy model, along with the sampling mechanism.

With the goal of performing inference analyses using *MATLAB*, an open-source software called *UQLab* has been used. *UQLab* is an Uncertainty Quantification framework, developed at ETH Zurich (Switzerland), made of open-source scientific modules which are smoothly connected to perform uncertainty quantification [1].

UQLab's framework is constituted of different parts, represented in Figure 6.2. All these components play different roles and are overall complementary during the process. These, with the exception of the Seed and the Solver, which are mainly related to the software used, are usually shared between every inference framework. Between them, the forward model will be the core of our analysis and it will encapsulate the surrogate model presented in Chapter 4.

6.1 Prior distributions

As discussed in the previous Chapter, in Section 5.1, every inference analysis needs the definition of a prior distribution for the inferred parameters, which has to best reflect the user's knowledge about their distribution before new evidence is considered. Inside the *UQLab* framework, several prior distributions for the inferred parameters can be defined, such as uniform, log-normal, Gaussian, and Beta, each one reflecting different types and degrees of prior knowledge.

Specifically, our inference process in *UQLab* will employ the use of uniform and log-normal distributions. Within the software, these distributions can be set by defining the mean and the standard deviation of the underlying normal distribution, for what concerns the log-normal, and by its bounds, a and b, for what concerns the uniform.

6.2 Seed & Solver

In statistical analyses that involve random sampling, such as Markov Chain Monte Carlo (MCMC) methods, commonly used in Bayesian inference and used in *UQLab*, the Seed sets the starting point for the sequence of random samples generated by the algorithm. In particular, these random numbers will always be withdrawn from the specified prior distribution of the inferred parameter, and will constitute what is called a "chain".

In order to better understand this concept, it is important to know that in *UQLab* the MCMC solver can be configured to use various sampling algorithms, such as the Affine Invariant Ensemble Sampler (AIES), that will be used for our inference analyses. The AIES relies on multiple parallel chains for sampling. So, the "Seed" initializes the random number generator determining the starting points for these chains. In every iteration of the process, so for every value of the chain, the surrogate model and the later dimensionalization process will be called and resolved, allowing for comparison with experimental data.

After the first iteration, subsequent values in each chain are generated based on the prior distribution and evaluated using the likelihood function and the model's forward predictions. This iterative process involves accepting or rejecting proposed samples based on their probability (to give, as output, the experimental data), ensuring that the sampling explores

the parameter space effectively while converging toward the high-probability regions.

A key factor is that by fixing the Seed, it is possible to ensure the reproducibility of the results. In fact, as *UQLab* utilizes the Mersenne Twister algorithm for the generation of random numbers, employing the *rng(10, 'twister')* command in *MATLAB* before every analysis ensures that the values generated from the Seed and throughout the MCMC process, under the AIES sampler, will always be consistent. This is fundamental as it ensures that, in every analysis, the forward model will be evaluated with the exact same samples, allowing for reproducibility of results and comparisons between analyses.

After the posterior distribution is obtained, the posterior samples are always subjected to burn-in. In fact, in *UQLab*, the first half of generated samples are typically eliminated to guarantee the convergence of the sampling by eliminating poorly conditioned chains. This process ensures that the remaining samples accurately reflect the posterior distribution, providing reliable insights into the parameters being analyzed.

6.3 Forward model

In *UQLab*, the "Forward Model" plays a crucial role as it represents the mathematical or computational model that describes the system or process under investigation. Essentially, it is a set of equations or algorithms that simulate how the system behaves under various conditions. The forward model takes the parameters of the system, either stochastics or constants, producing predicted outcomes or responses. These predictions are then compared with the actual experimental or observational data.

The forward model links the parameters (and their uncertainties) to the observable data, thus bridging the gap between theoretical models and real-world observations. By integrating the forward model within the Bayesian framework, we can update our beliefs about the system's parameters based on how well the model's predictions match with the observed data.

For the purpose of the inference analyses regarding PICA, the forward model will include the surrogate and the following dimensionalization, previously discussed in Chapter 3 and 4. In particular, at each iteration of the inference process, a solution will be withdrawn

from the surrogate based on the α value used. This solution, will be then dimensionalized by using the definition of reference time and temperature. After that, six temperatures, related to six different times (from 15 to 40 seconds and every 5) will be taken to be compared with experimental values available at the same simulation times.

The process of retrieving adimensional results from the surrogate requires the definition of α, while the subsequentially dimensionalization involves the specification of inputs such as *rhoCp*, *q0* and $k_{\perp}$, as presented in Fig. 6.1. Depending on which ones of those are inferred, these inputs can either be defined as constant (and so, not "participating" in the inference process) or assuming a statistical distribution defined by their prior. By employing a prior distribution for a specific parameter, the forward model will accept as input the samples that have been withdrawn from its distribution by the solver. Otherwise, if the variable is set as constant, it will take as an input just its value, which will not change during the following iterations.

It is important to focus on the definition of the reference quantities, as they contain parameters that could later be inferred. In fact, a crucial aspect lies in the fact that t_{ref} is linked to rhoCp and $k_{\perp}$ through Eq. 3.13, so its definition could possibly change at each iteration when these inputs are not constants during the inference process. This would lead to obtaining dimensional temperatures referred to dimensional time intervals that will not match the ones at which experimental data are available so that comparison will no longer be logical. An interpolation is then required and is carried out during the dimensionalizing process. In this process, a dimensional time interval with a specific time step (for example, from 0 to 40 seconds, with 1 second of time step) can be defined by the user. On this, a spline interpolation of the dimensional temperatures retrieved from the surrogate will be performed, so that the final output of the forward model will consist of dimensional temperatures calculated with respect to the user-defined dimensional time interval, allowing for a reliable comparison with experimental values.

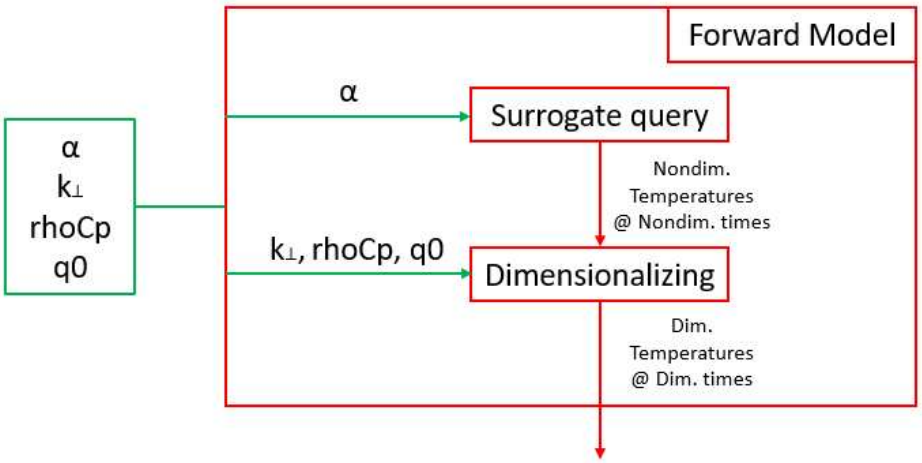

Figure 6.1: Forward model diagram.

6.4 Experimental data

In Bayesian inference, "Experimental Data" plays a pivotal role as it represents the real-world observations or measurements against which theoretical models are tested and refined. Within UQLab, these experimental results are used as a critical benchmark for evaluating the accuracy and validity of the computational model employed in the analysis, in our case, the forward model described by Figure 6.1.

In the case of our inference analyses, the set of experimental data available from the work of Milos and Chen [4] will be used. Depending on which thermocouples will be employed in the process, experimental data relative to them will be inserted into the framework.

Through the application of Bayes' Theorem (Chapter 5, Section 5.1.1) and by using experiments, the initial beliefs about the parameters will be updated, resulting in posterior distributions. Therefore, in *UQLab*, experimental data is not merely a set of observations but a fundamental component that drives the iterative process of model improvement and uncertainty quantification in Bayesian inference.

6.5 Discrepancy

In the Bayesian inference framework implemented in *UQLab*, the Discrepancy is used to account for the differences between the predictions of the computational model and the actual experimental observations. This discrepancy is crucial as it encompasses both the model's inherent limitations in accurately representing reality and the experimental measurement errors.

In our scenario where the model does not account for certain physical phenomena, such as pyrolysis or the recession of the sample, along with the ones explained in Chapter 3 Section 3.1, the discrepancy term helps to bridge this gap. Additionally, it includes thermocouple measurement errors, which in our case are estimated to be ± 5 %, as well as errors in the measurement of the cold wall heat flux, estimated as ± 10 and 15 % from [4]. In UQLab, the discrepancy model will be defined by the variance of its Gaussian (i.e. normal) distribution.

The Gaussian distribution is particularly useful for representing discrepancy allowing for a symmetrical, bell-shaped representation of errors. Furthermore, the variance of the discrepancy model itself can be treated as a parameter subject to inference. By assigning a prior distribution of the discrepancy, either uniform or Gaussian, it becomes an integral part of the Bayesian inference process. This approach allows for a more comprehensive and nuanced estimation, as it not only seeks to infer the model parameters but also the degree of uncertainty associated with the model's ability to replicate reality.

All the aspects previously discussed regarding the Bayesian inference process and their interrelationships are summarized in Figure 6.2, which concisely illustrates the flow of the analysis by following the incoming arrows to each block, indicating which parts of the inference framework interact with each other.

The Prior Distributions set the stage with initial parameter assumptions, while the Seed ensures consistent random number generation for reproducibility. These elements feed into the Forward Model which uses the parameters to simulate the system's behavior. The Solver is at the heart of the Bayesian update, considering the Forward Model outputs and the Experimental Data to refine the parameter estimates. It also accounts for the Discrepancy

representing the gap between model predictions and actual observations, encompassing both modeling errors and experimental uncertainties.

The goal of the Solver is to reduce this discrepancy, thereby fine-tuning the parameters to better reflect the observed data.

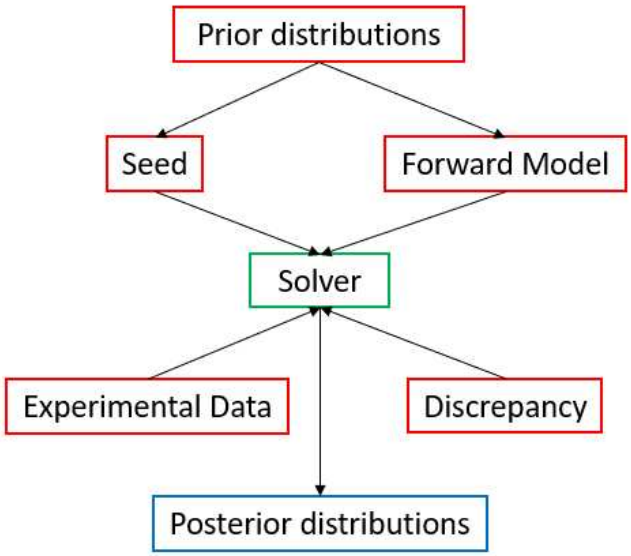

Figure 6.2: Inference framework functional diagram.

Chapter 7

Framework Settings

Inference has been made on two of the thermal analysis inputs represented in Figure 6.1. Namely, α, the ratio of thermal conductivities, and q_0, the cold wall heat flux at stagnation point. Throughout this process, the other two inputs of the model, rhoCp and $k_\perp$, have been set as constants, assuming nominal values reported in Table 4.1. In this context, it is crucial to remember that assuming $k_\perp$ constant but varying α, which are related by Eq. 4.1, means that a change of α is related to a change in $k_\parallel$.

All the parts presented in the previous Chapter will play their specific role inside the inference framework, aiming to obtain a posterior distribution for both the inferred parameters. To correctly set this parts, some studies have been first carried out in order to investigate the inference results dependency on the Sampler, part of the Solver, and Seed chosen. To perform these simulations, which results are presented in Sections 7.3 and 7.4, sets of single and double thermocouples have been used.

Along with these, a study assessed how many iterations are needed for the MCMC chains to converge to the posterior distribution. This is crucial as it indicates the number of iterations to set in order to get reliable posterior distributions. To investigate the chain convergence, visual comparisons of the posterior PDFs of α and q_0, obtained with a different number of iterations, have been made. Alongside this, the calculation of the Gelman-Rubin statistics (presented in Chapter 5, Section 5.2) allowed for further confirmation. Studying the convergence of the chains has been crucial to prove the reliability of the results but, also, to limit the computational time of the following analyses, which have been possible to run

using limited iterations. Before performing the real inference process on α and q_0, inference about the discrepancy model's variance has been performed, as presented in Section 7.5. After this process, which has been carried out using different sets of thermocouples, two values of variance have been chosen to be tested during the actual inference on α and q_0. Inference results obtained with both variances have then been compared in Chapter 8.

Studies have also investigated the relation between the number of parallel chains and iterations. It has been demonstrated that using a N_1 number of parallel chains, each with a N_2 number of iterations, gives the same results as using N_2 number of parallel chains and a N_1 number of iterations each.

N.B.: Analyses regarding MCMC chains convergence, inference dependency on Seed, and discrepancy have all been run using the AIES sampler. All analyses simulated 40 seconds of testing.

7.1 Prior distributions of α and q_0

Before performing all the subsequent analyses, it has been crucial to decide which priors to use for α and q_0. From now on, whenever one of the two parameters is inferred, its prior distribution will always be defined as the one presented in this Section. In particular, for α, given that it can only assume positive values, a log-normal distribution has been chosen. For what concerns q_0, a uniform distribution has been selected.

Log-Normal Distribution of α

The log-normal α distribution chosen is characterized by these values:

- Median value, α = 2;
- 2.5th percentile, α = 1.2;
- 97.5th percentile, α = 2.8.

By using the median value, the mean of the normal distribution (of the natural logarithm of α) has been calculated via Eq. 5.7, while its standard deviation (STD) has then been estimated by applying the above-mentioned percentiles to Eq. 5.10. Results are: $\mu = 0.6931$ (mean) and $\sigma = 0.1717$ (STD).

By employing these percentiles, the log-normal distribution will have 5% of α values below 1.2, while 97.5% below 2.8. This in turn means that the 95% of α will be between 1.2 and 2.8. This distribution is displayed in Fig.7.1, with a dashed line indicating the median.

Uniform Distribution of q_0

The q_0 uniform distribution bounds (a and b as introduced in Chapter 5) have been chosen as:

- Inferior bound, q_0 = 84.5 W/cm^2;
- Superior bound, q_0 = 338 W/cm^2,

which consist of half and two times the given nominal value of 169 W/cm^2, reported in Table 4.1. The probability density is constant between its bounds and its value can be estimated by Eq. 5.3.

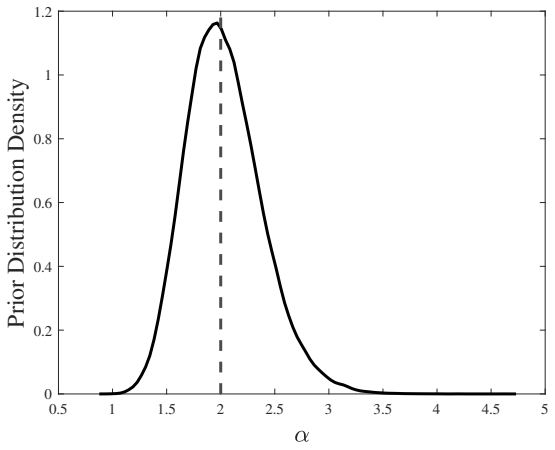

Figure 7.1: α prior.

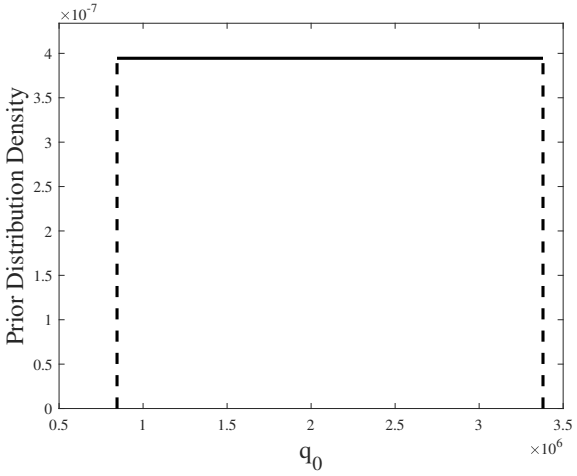

Figure 7.2: q_0 prior, bounded from a to b.

7.2 Markov Chains convergence

To study the convergence of the chains, analyses have been run by inferring only one parameter at a time and by changing the number of iterations for each simulation. Posterior distributions of both parameters obtained with different iterations have then been compared, as visualized in the next Figures. Analyses have been performed using a specific Seed for both parameters, which have then been later proven to be reliable by following analyses, presented in the next Section.

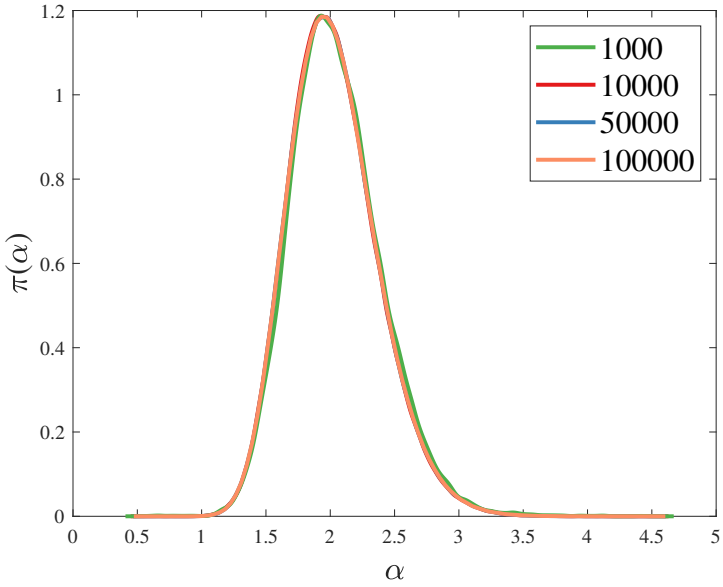

Figure 7.3: α posteriors, obtained with a different number of iterations. Seed 0.5 - 4.5.

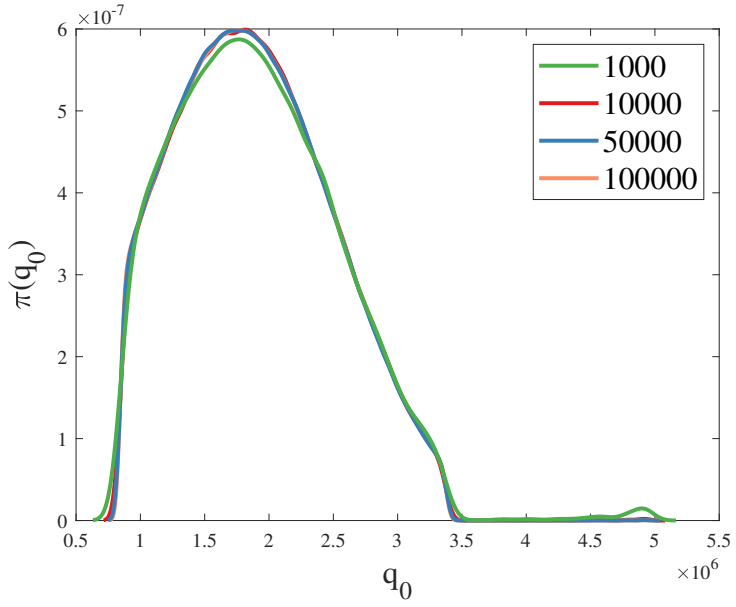

Figure 7.4: q_0 posteriors, obtained with a different number of iterations. Seed 84.5 - 500 W/cm^2.

As it can be visualized in Figure 7.3 and 7.4, the posterior PDFs of α and q_0, obtained with different numbers of iterations, stay consistent. By the way, the real confirmation is obtained by employing the Gelman-Rubin statistics. $\hat{R}$ values have been obtained for each case by applying Eqs. 5.11, 5.12 and 5.14. As it can be seen from Table 7.1 and 7.2, all $\hat{R}$ values are close to 1. This indicates that the Markov Chains have likely settled into their stationary distribution and that the sampling process is providing reliable estimates of the posterior distributions. These tests have been run using only TC6, whose placement can be seen in Fig. 8.1.

Iterations	$\hat{R}$
100000	1.0001
50000	1.0003
10000	1.0014
1000	1.0155

Table 7.1: G.-Rubin statistics for α PDFs.

Iterations	$\hat{R}$
100000	1.0001
50000	1.0002
10000	1.0014
1000	1.0184

Table 7.2: G.-Rubin statistics for q_0 PDFs.

7.3 Inference dependency on Seed

In order to investigate the relationship between the posterior PDFs obtained through inference and the Seed chosen, posterior distributions of α and q_0 obtained with different Seeds have been compared. Along with a visual comparison, the mean and standard deviation values related to each distribution have been taken into account to investigate possible differences. It is important to note that, in the context of these analyses, these two values have been taken as nothing else as indicators of the difference between PDFs calculated using different Seeds, as the real inference process that will allow the assessment of these values related to the posteriors of α and q_0 will be carried out later.

To perform this study, a discrepancy model's variance of 100 K has been employed, a value that has been justified by later analyses reported in Section 7.5.

The study has been performed using 3 different sets of thermocouples, employing axial and radials. These have been taken to be representative of every possible combination of TCs that could be possibly used for the inference process. In fact, thermocouples chosen for testing are: TC6, placed on the axis, TC8, placed in a radial position, and the combination of the two tested together. Their placements are represented in Fig. 8.1.

The posterior PDFs obtained using TC6, displayed in Figs. 7.5 and 7.6, have been obtained using 50000 iterations and 100 parallel chains. PDFs for sets TC8 and TC6-TC8 have been obtained using 10000 iterations and the same number of parallel chains, using the proven convergence. For the sake of brevity, given that they all demonstrated consistency as the ones for TC6, they are presented in Appendix [1].

7.3.1 Seed variaton of α

When testing the Seed variation on α, q_0 has been taken constant.
Considering the nominal value of $q_0 = 169\, W/cm^2$ from Table 4.1, three different Seeds (Case 1 to 3) of α have been used for testing, as reported in the next Table:

Case 1	0.5 - 4.5
Case 2	1 - 3
Case 3	1.5 - 2.5

Table 7.3: α Seeds tested.

By comparing the posterior density functions for all the Seeds tested, only a slight change in the PDF of α when employing different Seeds has been observed. This is also proven by the calculation of mean and standard deviation from each PDF, reported in the next Tables. It is then possible to draw the conclusion that using either one of the three different Seeds proposed in Table 7.3 does not have any implications on the result (i.e. the posterior) of the inference process of this variable.

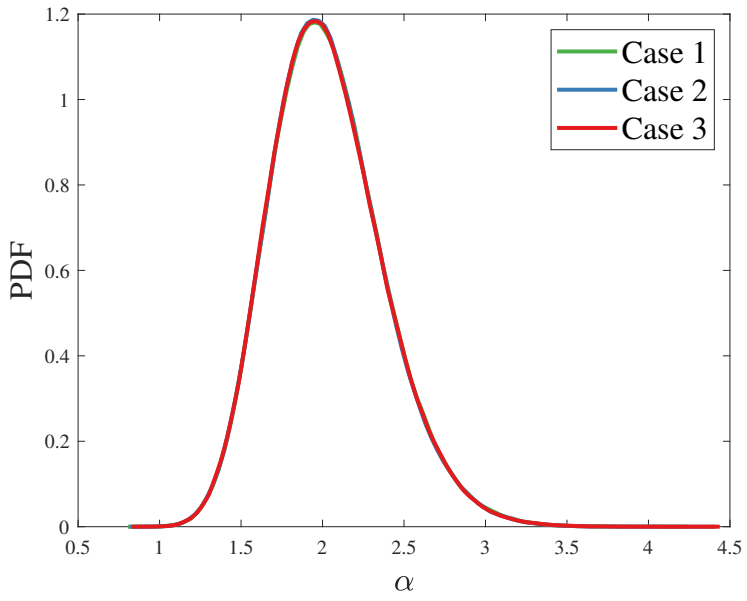

Figure 7.5: Posterior PDFs of α with different Seeds for TC6.

Case	Mean	STD
1	2.0281	0.35
2	2.0275	0.34
3	2.0284	0.34

Table 7.4: TC6.

Case	Mean	STD
1	1.9509	0.046
2	1.9508	0.046
3	1.9510	0.046

Table 7.5: TC8.

Case	Mean	STD
1	1.9514	0.046
2	1.9511	0.046
3	1.9505	0.046

Table 7.6: TC6-TC8.

7.3.2 Seed variation of q_0

Considering the nominal value of $\alpha = 2$ from Table 4.1, q_0 has been inferred considering three different Seeds (Case 1 to 3), as reported in the next Table:

Case 1	84.5 - 500
Case 2	84.5 - 400
Case 3	84.5 - 300

Table 7.7: q_0 Seeds tested, in W/cm^2.

As Figure 7.6 shows, the study confirmed that, as for α, no major differences are evident when using different Seeds for q_0, in terms of posteriors and their characteristics.

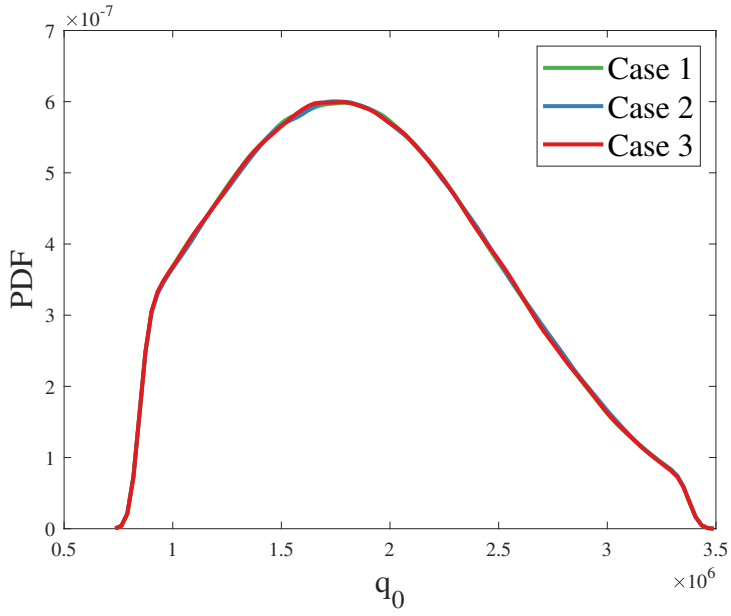

Figure 7.6: Posterior PDFs of q_0 with different Seeds for TC6.

Case	Mean	STD
1	189.47	59
2	189.61	59
3	190.08	59

Table 7.8: TC6, in W/cm^2.

Case	Mean	STD
1	166.71	2.5
2	166.75	2.5
3	166.69	2.5

Table 7.9: TC8, in W/cm^2.

Case	Mean	STD
1	166.71	2.5
2	166.70	2.5
3	166.76	2.5

Table 7.10: TC6-8, in W/cm^2.

7.4 Inference dependency on the Sampler

UQLab supports four different samplers that can be used and incorporated in the solver when running inference analysis: Metropolis Hastings (MH), Adaptive-Metropolis (AM), Hamiltonian Monte Carlo (HMC) and Affine Invariant Ensemble Sampler (AIES).
Even though a complete study about the sampling mechanisms and related theoretical aspects has not being the focus of this thesis, the relevancy of the sampler choice on the inference process has been investigated. To do so, the 3 different sets of thermocouples used in Section 7.3 have been utilized.

Analyses have been carried out with a discrepancy variance of 100 K and by only inferring the parameter α, choosing its Seed to be 0.5 - 4.5. By doing so, a posterior distribution for each sampler used has been derived, and their main values (such as mean and STD) have then been compared, as in the previous simulations.

As it can be depicted from Figure 9.5, for TC6 the posterior distribution functions of α, obtained with 100 parallel chains and 10000 iterations, stay consistent while using different Samplers. The same results have been obtained for TC8 and TC6-TC8, represented in Appendix [3]. Consistency is also confirmed when comparing STD and mean values of the different PDFs, reported in Tables 7.11, 7.12 and 7.13.

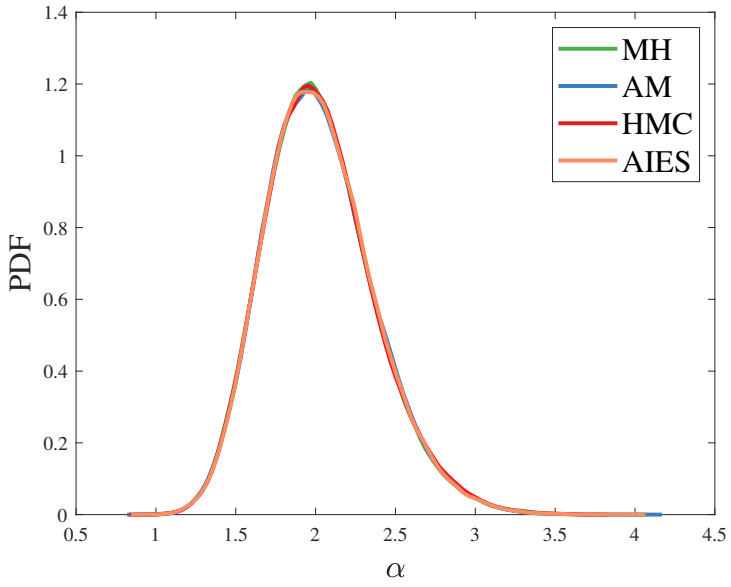

Figure 7.7: Posterior PDFs of α with different samplers for TC6.

Case	Mean	STD
AIES	2.0281	0.35
AM	2.0275	0.34
MH	2.0257	0.34
HMC	2.0283	0.35

Table 7.11: TC6.

Case	Mean	STD
AIES	1.9509	0.046
AM	1.9509	0.046
MH	1.9508	0.047
HMC	1.9508	0.046

Table 7.12: TC8.

Case	Mean	STD
AIES	1.9514	0.046
AM	1.9509	0.046
MH	1.9508	0.047
HMC	1.9509	0.047

Table 7.13: TC6-TC8.

Throughout this process, all samplers demonstrated similar computational times with the exception of HMC, which required at least 10/12 more times when compared to the others. Among these options, the AIES sampler has been chosen for the following inference analyses.

7.5 Discrepancy choice

Performing inference on α and q_0 also requires the definition of a discrepancy model. In order to choose an adequate model, which is defined in *UQLab* through the variance of its distribution, inference on this parameter has been performed. To do so, different sets of single, double, and triple thermocouples, as resumed in the next Tables, have been used. For each set, simulations have been performed with 100 parallel chains and 1000 iterations.

TC3	TC4	TC5	TC6	TC8

Table 7.14: Sets with single TCs tested.

TC5-TC7	TC5-TC8	TC6-TC7	TC6-TC8

Table 7.15: Sets with double TCs tested.

TC6-TC7-TC8	TC5-TC7-TC8

Table 7.16: Sets with triple TCs tested.

As the variance of the discrepancy model is now the inferred parameter, its prior distribution has to be defined, as required by the framework. In particular, a uniform distribution has been used, which bounds have been chosen to be 25 K and 6400 K. The prior uniform distribution has been taken to be uniform and fairly wide, reflecting our limited prior knowledge about this variable, which encapsulates several different aspects. Its Seed has also been set as going from 25 K to 6400 K. Throughout these analyses, only the variance has been inferred, and by doing so the thermal inputs have been all set as constant and equal to their nominal values reported in Table 4.1.

For every set described in Tables 7.14, 7.15 and 7.16, a Gaussian posterior distribution of the discrepancy has been obtained through inference and its mean and standard deviation values have been considered and reported in Table 7.17. The Table clearly shows how using sets with double and triple TCs, compared to sets with one thermocouple, gives a lower STD of the posterior, related to a narrower distribution and hence a bigger information gain. However, between the two and three TCs' sets, TC5-TC7 and TC5-TC8 showed higher STD values, indicating more inaccuracy.

Thermocouple Set	**Mean (K)**	**STD (K)**
TC3	849.29	810
TC4	583.74	650
TC5	380.04	460
TC6	77.25	120
TC8	177.22	350
TC5-TC7	122.04	74
TC5-TC8	144.21	95
TC6-TC7	41.44	20
TC6-TC8	49.92	28
TC5-TC7-TC8	92.34	39
TC6-TC7-TC8	44.02	17

Table 7.17: Mean and STD of the discrepancy model's variance distribution.

Upon examining Table 7.17, it is apparent how in every set containing the TC5 thermocouple, the standard deviation (STD) of the posterior is higher compared to sets with

the same number of thermocouples but without this specific thermocouple. This could be attributed to the fact that TC5, being the deepest axial thermocouple, may not provide comprehensive insights into the discrepancy. In fact, discrepancies serve to narrow the gap between experimental results and the model, which in our case primarily reflects the initial seconds of testing. Therefore, employing a deep thermocouple like TC5 might not yield relevant information since it may not have sufficient time to detect the heat flux through the sample. This is also confirmed by experimental data in [4]. In fact, it can be seen from Milos and Chen how the temperature in TC5 only increases by 30 K during the first 40 seconds of testing, when considering a low cold wall heat flux such as 169 W/cm^2. At the same time, experimental data related to higher cold wall heat fluxes, show a higher temperature increase in the same simulation interval in TC5. Finally, a higher temperature increase is also shown at 169 W/cm^2 but in thermocouples positioned closer to the sample's surface.

However, given the uncertainty regarding the discrepancy factor, which accounts for several aspects not entirely known, two variances have been selected for later use in the inference process, giving space for comparisons between results.

One variance has been taken to be 36 K, as the mean between the mean values of the STDs for the sets TC6-TC7, TC6-TC8, and TC6-TC7-TC8. Another value has been taken to be 100 K, as representative of mean values of the STDs obtained through sets TC5-TC7, TC5-TC8, and TC5-TC7-TC8.

Chapter 8

Inference Results

In this Chapter, the results of the inference process performed on α and q_0 will be presented. The process has been carried out using the prior distributions and the two discrepancy model's variances presented in Chapter 7, comparing the results obtained with both. The Seed for both parameters have been chosen between the ones previously investigated and the sampler as the AIES.

A first round of inference analyses, in which the parameters have both been inferred in the same simulation, have been performed with 100 parallel chains and 1000 iterations each. These analyses aimed to assess which thermocouple sets are more informative about the inferred parameters. To define the most informative sets, the information gain[1] has been calculated through the Kullback Leibler Divergence presented in Chapter 5. A high KL value indicates a high difference between the posterior and the prior distribution, hence a high information gain achieved through the inference process. However, unlike what has been displayed in Eq. 5.15, a small increment ε has been added to both P(x) and Q(x) in their ratio, to ensure not to obtain undefined values in *MATLAB* during the calculation. The KL has then been computed as visualized in the next equation:

$$D_{KL}(P \parallel Q) = \sum_{x \in X} P(x) \log \left(\frac{P(x) + \varepsilon}{Q(x) + \varepsilon} \right). \tag{8.1}$$

[1] In this context, "gain" refers to how much the prior distribution of the parameter has narrowed after the inference process has been performed

However, in order to make sure the increment ε had not influenced the outcomes, the Kullback Divergences have been computed with other values of ε, verifying their value was not changing. KL calculated with different ε increments demonstrated consistency and can be visualized in Appendix [5] and [6].

To compare KL values obtained using different thermocouple sets, every analysis performed employed the same prior distribution of α and q_0, as presented in Section 7.1. Simulations included sets of single, double, and triple thermocouples, presented in Section 8.1. The most informative sets have later been used for performing more in-depth inference analysis, presented in Section 8.2. These have been run with 100 parallel chains and 500000 iterations each, for a total of 50 million samples generated by the framework. To validate the results of these analyses, a forward propagation study, presented in Section 8.3, has been conducted. By doing so, it has been possible to test the reliability of the posterior distributions formerly calculated by utilizing them to guess the temperature evolution in other thermocouples.

Inference performed using different thermocouple sets also allowed for the investigation of the correlation factor between the two inferred parameters, as presented in Section 8.4. Correlation studies also highlighted the possibility of performing a sequential inference analysis (i.e. inferring one parameter first, using a specific set, and then inferring the other, using a second different set), which results have been compared to the ones obtained inferring both parameters within the same simulation.

Utilizing both discrepancy models' variances for the inference process did not highlight major differences. In particular, using the variance value of 36 K allowed for obtaining lower STDs of the posterior distributions and so higher values of KLs, as reported in Appendix [4]. This result is in line with expectations, as a declared lower discrepancy between model and experimental data necessarily leads to narrower posterior distributions and hence higher Kullback Divergences. However, a change in the variance did not lead to changes in terms of which thermocouple sets are more informative about the two variables, q_0 and α, which is by far the most important aspect.

8.1 Definition of the most informative sets

Analyses were performed using the Case 1 Seed for α and the Case 2 for q_0, respectively reported in Tables 7.3 and 7.7. Simulations have not been performed on sets including TC1, TC2, TC9 and TC10. In fact, as it has been demonstrated in Chapter 3, Section 3.7, these TCs are found to be not reliable throughout the whole 40 seconds interval of sample's exposition to the arc-jet simulated by the Model, failing to represent the physical reality after 10 or 20 seconds.

For every simulation performed, the mean and standard deviation values of the posterior distributions have been calculated and reported in the next Table. The Kullback Divergence, also reported in Table 8.1, has been calculated between the posterior and the prior of each variable's distribution.

Set	**Thermocouples**	α **Mean**	α **STD**	q_0 **Mean**	q_0 **STD**	α **KL**	q_0 **KL**
1	TC3	1.96	0.33	171.47	4.58	0.02	2.62
2	TC4	2.02	0.34	142.67	8.83	0.01	1.91
3	TC5	2.04	0.35	212.98	73.36	0.01	0.06
4	TC6	2.02	0.34	186.46	59.39	0.01	0.15
5	TC7	1.96	0.33	155.32	30.55	0.02	0.71
6	TC8	1.92	0.14	171.13	7.75	0.52	2.11
7	TC5-TC7	1.96	0.33	156.08	30.37	0.03	0.70
8	TC5-TC8	1.92	0.14	171.35	7.80	0.53	2.11
9	TC6-TC7	1.95	0.32	157.19	30.05	0.04	0.75
10	TC6-TC8	1.92	0.14	171.39	7.71	0.51	2.10
11	TC5-TC7-TC8	1.87	0.12	173.59	7.35	0.66	2.17
12	TC6-TC7-TC8	1.87	0.12	173.40	7.17	0.66	2.17
13	TC5-TC6-TC7-TC8	1.87	0.12	173.62	7.33	0.65	2.17

Table 8.1: Inference results for $\sigma^2 = 100$ K. q_0 values are reported in W/cm^2. Mean and STD are referred to posterior PDFs. KL have been calculated with an increment of $\varepsilon = 10^{-10}$.

As resumed in Table 8.1, the KL calculation clearly shows how using some sets instead of others allows for a higher information gain about the inferred parameters. This is the case of sets 6, 8, 10, 11, and 12, for which the KL value is significantly higher.

This first round of simulations has been particularly crucial in demonstrating two major general facts about how thermocouple placement impacts the information gain for a specific variable. These can be resumed as follows:

- **Information gain on α**: the presence of a radial thermocouple is instrumental for achieving a high information gain on α, as a variation of α is induced by a variation in $k_{\parallel}$, which is the thermal conductivity of the material in the radial (i.e. parallel) direction. This is clearly shown by the comparison between the α KL calculated using an axial thermocouple and a radial thermocouple, such as TC3 and TC8. In fact, as visualized in Table 8.1, the information gain is significantly higher using TC8, which is the radial one. The same consideration is valid when looking at TC8 and TC7, which comparison highlights a lower KL in TC7, being in a lower radial coordinate, and so closer to the axis (see thermocouple locations in Fig. 8.1);

- **Information gain on q_0**: information gain seems to be linked to the position of the thermocouple with respect to the outer wall and shoulder of the sample. In fact, by looking at Tabe 8.1, it can be seen that using thermocouples closer to the surface, such as TC3 or TC4, gives a higher value of KL when compared to using deeper TCs like TC5 or TC6. This can also be noticed by comparing the sets using TC7 or TC8 with the one using TC6, where the first two give more information gain, as they share the same axial coordinate with TC6 but they correspond to a higher radial coordinate, hence, being closer to the shoulder of the sample (see TCs locations Fig. 8.1).

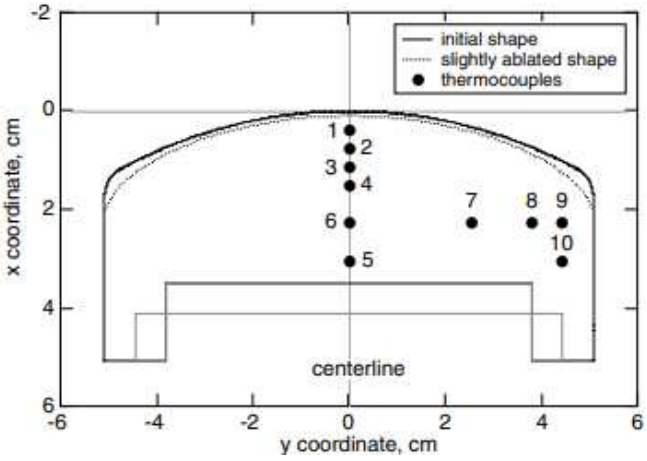

Figure 8.1: Thermocouple locations. Adapted from [6].

So, it is evident that thermocouple sets using a combination of radial and axial TCs will lead to a higher value of Kullback Divergence.

In this sense, set number 6, containing only thermocouple 8, seems really interesting, as it represents the intersection of these two concepts in one spot. This thermocouple is in fact close to the surface, resulting in a high information gain of q_0, and it is radial, essential for a high information gain on α. Moreover, this thermocouple does not suffer the inaccuracies of the Model linked to pyrolization. In fact, as demonstrated in Chapter 3 Section 3.7, pyrolization does not occur on TC8 throughout the 40 seconds of arc-jet exposition.

Finally, a simulation ran using a set with four thermocouples (set 13) showed, through the KL calculation, that using more than three thermocouples does not lead to a greater information gain. Cases with five or six thermocouples in one set, technically possible using all TCs within Placement A and D respectively, have not been investigated, as they would have included thermocouples that fail after 10 or 20 seconds, as discussed before.

8.2 Posterior distributions of α and q_0

After assessing the most informative sets through the analysis reported in the previous Section, three sets have been chosen to perform analyses with more iterations. In particular, simulations with 500000 iterations for each of the 100 parallel chains have been run for sets number 6, 8, and 12, containing respectively TC8, TC5-TC8, and TC6-TC7-TC8. By doing so, the parameters' posterior distributions have been obtained. In the following part, posterior distributions, along with their mean and standard deviation values, are reported, for both the inferred parameters.

Set 6: TC8

Parameter	Mean	STD
α	1.92	0.14
q_0	171.26	7.70

Table 8.2: Mean and STD values. q_0 values are reported in W/cm^2.

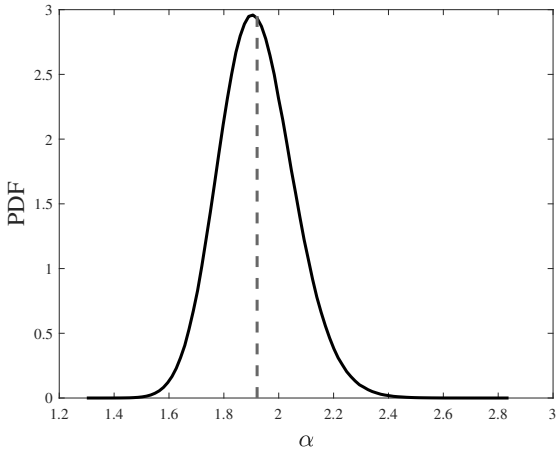

Figure 8.2: Posterior PDF of α, with dashed line indicating the mean value.

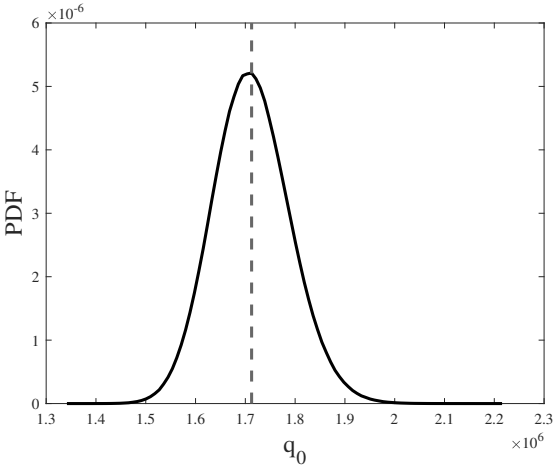

Figure 8.3: Posterior PDF of q_0, with dashed line indicating the mean value. Values in W/cm^2.

Set 8: TC5-TC8

Parameter	Mean	STD
α	1.92	0.14
q_0	171.26	7.70

Table 8.3: Mean and STD values. q_0 values are reported in W/cm^2.

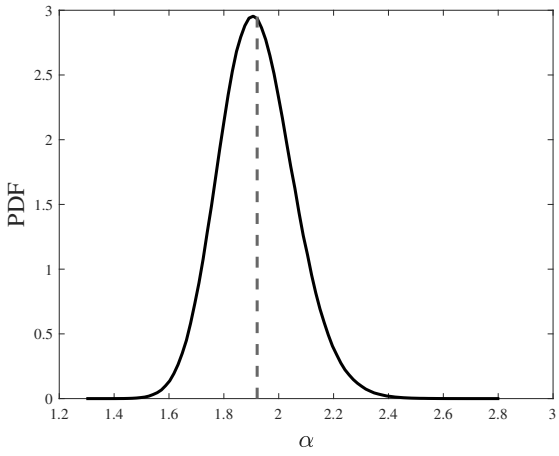

Figure 8.4: Posterior PDF of α, with dashed line indicating the mean value.

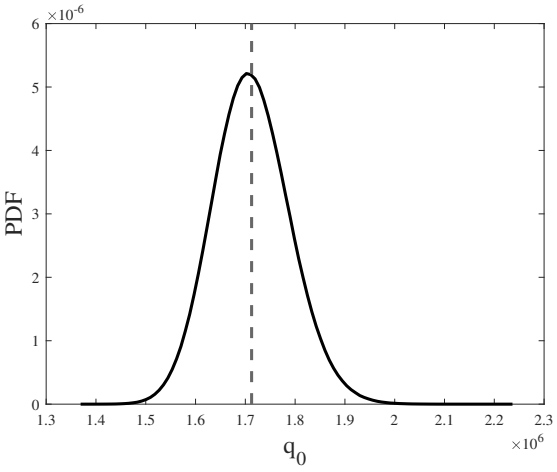

Figure 8.5: Posterior PDF of q_0, with dashed line indicating the mean value. Values in W/cm^2.

Set 12: TC6-TC7-TC8

Parameter	Mean	STD
α	1.87	0.12
q_0	173.70	7.30

Table 8.4: Mean and STD values. q_0 values are reported in W/cm^2.

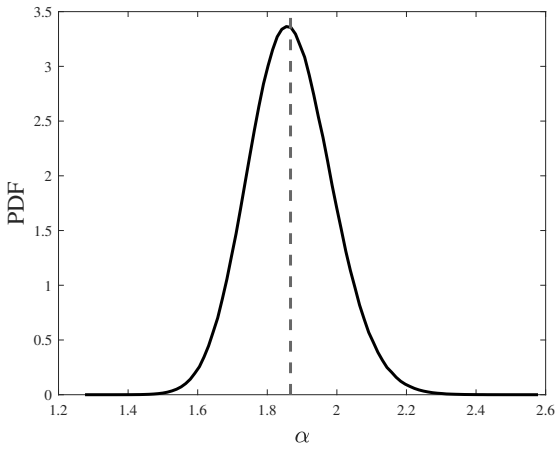

Figure 8.6: Posterior PDF of α, with dashed line indicating the mean value.

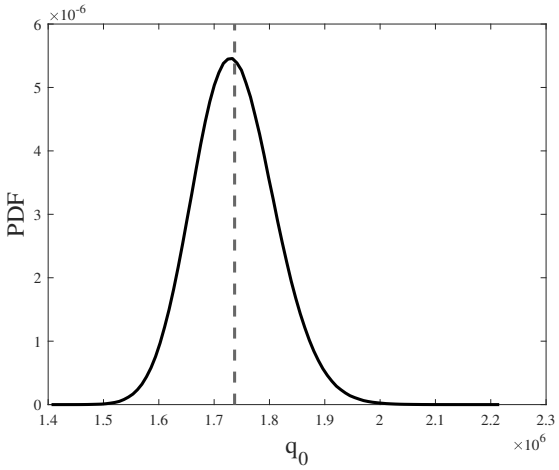

Figure 8.7: Posterior PDF of q_0, with dashed line indicating the mean value. Values in W/cm^2.

As can be noticed from the posterior distributions, the application of Bayesian inference to the experimental data has yielded insights that challenge the nominal values initially declared for α and q_0, reported in Table 4.1. The results obtained through these three different sets did not highlight great dissimilarities. By the way, set 12, containing three TCS, allowed for obtaining lower values of STD for both distributions, of α and q_0. The estimations obtained with all these sets have later been proven to be strongly reliable by the forward propagation analyses.

However, while the results of the Bayesian inference provide a promising indication, the current inference model incorporates aspects that warrant further refinement. One such aspect is the choice of the discrepancy model, which plays a critical role in capturing the divergence between the model predictions and observed data, and that incorporates aspects that could benefit from further studies.

8.3 Forward propagation

In order to validate the inferential outcomes of the most informative sets presented in the previous Section, a comprehensive forward propagation analysis has been conducted. For seek of brevity, only forward propagations obtained using TC8 (set 6) and TC6-TC7-TC8 (set 12) are displayed, as the ones for TC5-TC8 showed very similar results. For both sets 6 and 12, the temperature estimation for TC5 and TC9, is displayed in Appendix [9], [10], [11], and [12].

A forward propagation analysis utilizes couples (α, q_0) withdrawn from the posterior PDFs obtained using a specific set, as inputs for the Forward Model. The Forward Model returns as a result the temperature over time calculated in selected thermocouples, which have not been part of the inference analysis from which the posteriors have been previously calculated. For every sample (α, q_0), a vector of temperatures over time, in a specific TC location, is calculated and given as output. In particular, six values of temperature inside the vector are always given, with respect to six different instants in time. By obtaining a vector of temperatures for each sample, with all the samples tested, a temperature probability distribution will be obtained for each one of the six instants in time. From each distribution, the mean value has then been taken and compared to the experimental data available for that instant in time, to define the prediction's accuracy. In this way, it is possible to investigate the accuracy of the previously calculated posterior PDF from which the samples have been withdrawn.

In the next Figures, the mean temperature of each one of the six distributions has been taken and plotted together with experimental data. The orange area represents the STD from the mean values calculated by the Model. The absolute value of these differences, along with the STD of every distribution of temperature, are resumed in the next Tables.

Set 6: TC8

As shown in Figs. 8.8. 8.9, 8.10, 8.11, and 8.12 the posteriors of α and q_0 obtained with TC8 seem to predict temperatures in other thermocouples with good precision, given that experimental results are also expected to have a 5 % of uncertainty.

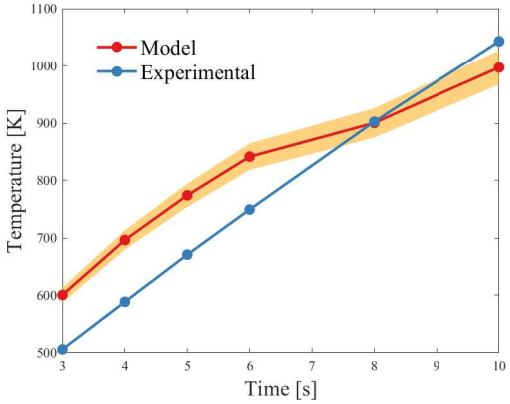

Figure 8.8: TC1 temperature evolution, predicted using TC8 posterior.

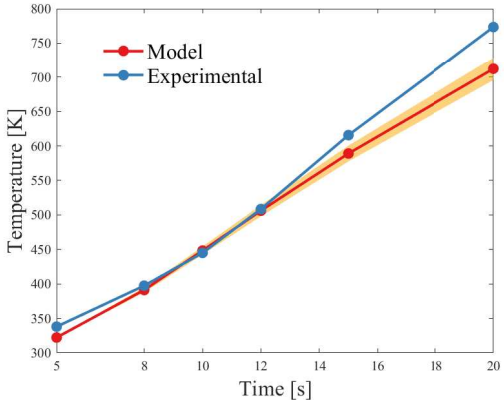

Figure 8.9: TC2 temperature evolution, predicted using TC8 posterior.

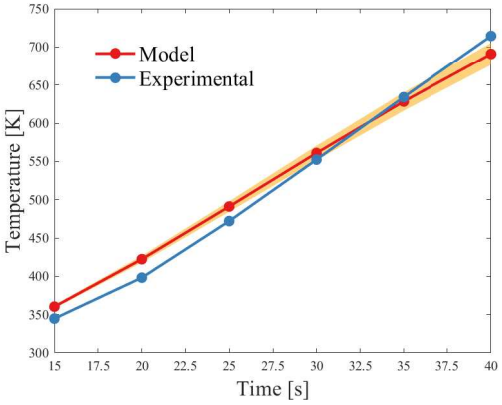

Figure 8.10: TC3 temperature evolution, predicted using TC8 posterior.

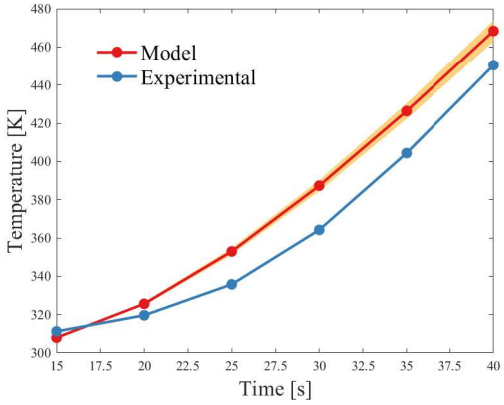

Figure 8.11: TC4 temperature evolution, predicted using TC8 posterior.

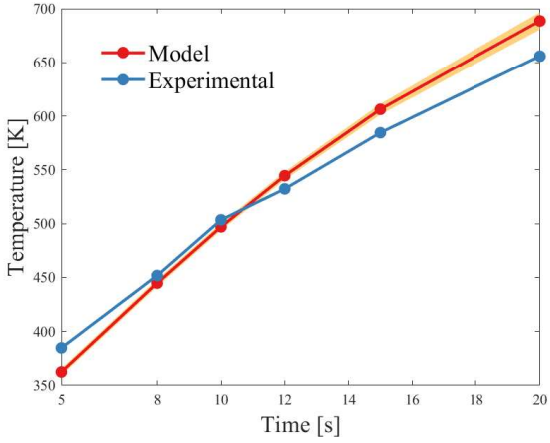

Figure 8.12: TC10 temperature evolution, predicted using TC8 posterior.

Higher inaccuracy is seen in predicting the temperature evolution in TC1, as well as

around 20 seconds in TC2. The inaccuracy about TC1 could be related to the impact of the initial condition, which from the Model is taken to be 300 K uniformly, disagreeing with experimental data. Also, it could be related to the fact that this thermocouple is the first to experience pyrolysis, which is not taken into account by the Model. The inaccuracy in TC2 is most probably linked to the fact that from 20 seconds of simulation, the computational Model starts to fail in representing the physics, as thoroughly discussed in Chapter 3. In fact, as can be noticed from the next Tables, TC1, TC2, and TC10 predictions were made for a time interval of 10 and 20 seconds respectively, in order to not exceed the time interval in which the Model predictions are reliable. The STD value reported refers to the standard deviations of the temperature distributions for each instant in time.

Time (s)	Mean - Exp (K)	STD (K)
3	94.68	12.90
4	107.63	16.95
5	103.17	20.28
6	92.04	23.17
8	2.23	25.68
10	45.67	29.80

Table 8.5: TC1, differences and STD.

Time (s)	Mean - Exp (K)	STD (K)
5	15.75	0.90
8	6.18	3.60
10	3.27	5.86
12	2.34	8.16
15	26.66	11.47
20	61.35	16.44

Table 8.6: TC2, differences and STD.

Time (s)	Mean - Exp (K)	STD (K)
15	15.57	2.18
20	24.13	4.43
25	19.03	6.93
30	8.66	9.50
35	5.77	12.03
40	23.48	14.45

Table 8.7: TC3, differences and STD.

Time (s)	Mean - Exp (K)	STD (K)
15	3.20	0.26
20	5.99	0.83
25	17.19	1.71
30	23.11	2.83
35	21.99	4.10
40	17.96	5.46

Table 8.8: TC4, differences and STD.

Time (s)	Mean - Exp (K)	STD (K)
5	22.28	3.25
8	7.02	3.37
10	6.47	3.32
12	12.23	3.67
15	21.81	5.02
20	32.68	8.04

Table 8.9: TC10, differences and STD.

Set 12: TC6-TC7-TC8

Similar results are obtained when using the posterior of TC6-TC7-TC8 for the forward propagation of the temperature in thermocouples 2, 3, 4, and 10, as displayed in the next Figures and confirmed from the next Tables. However, even though these estimation stay consistent with the previous results, in this case, a major dissimilarity with respect to experimental data is found in TC1.

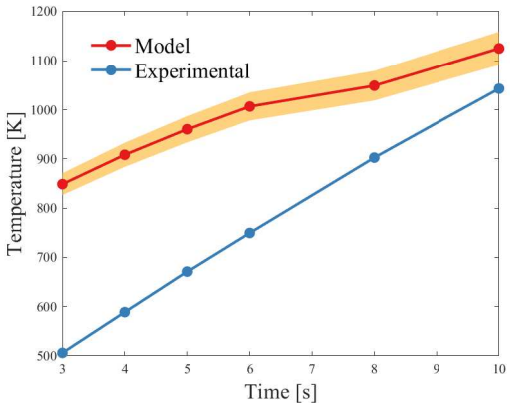

Figure 8.13: TC1 temperature evolution, predicted using TC6-TC7-TC8 posterior.

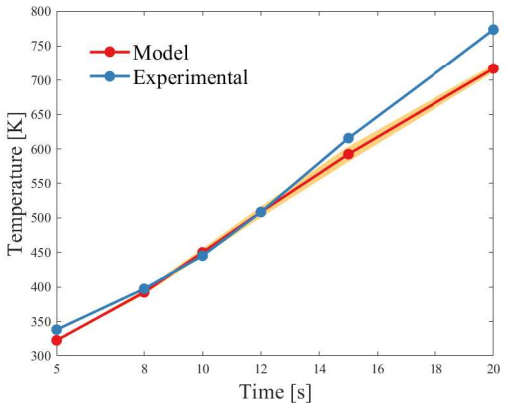

Figure 8.14: TC2 temperature evolution, predicted using TC6-TC7-TC8 posterior.

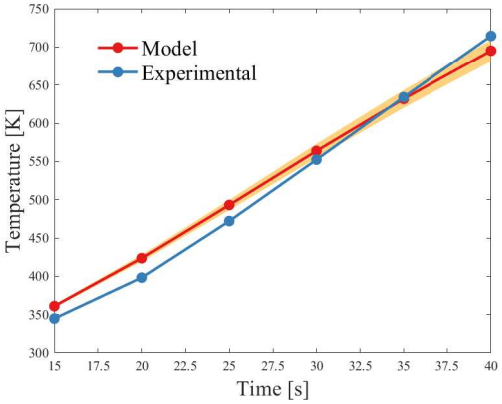

Figure 8.15: TC3 temperature evolution, predicted using TC6-TC7-TC8 posterior.

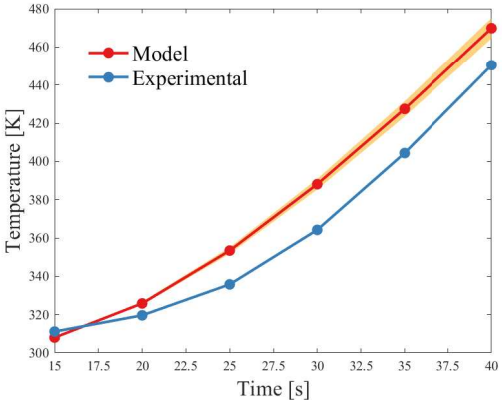

Figure 8.16: TC4 temperature evolution, predicted using TC6-TC7-TC8 posterior.

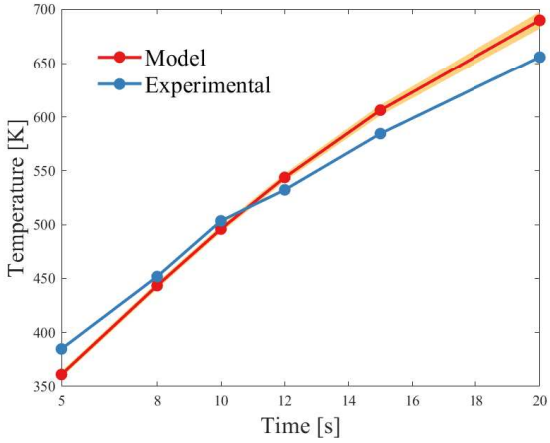

Figure 8.17: TC10 temperature evolution, predicted using TC6-TC7-TC8 posterior.

Time (s)	Mean - Exp (K)	STD (K)
3	259.86	22.20
4	237.32	24.61
5	210.87	26.70
6	180.59	28.56
8	146.62	30.27
10	81.89	33.30

Table 8.10: TC1, differences and STD.

Time (s)	Mean - Exp (K)	STD (K)
5	15.48	0.87
8	5.10	3.47
10	5.03	5.65
12	0.11	7.86
15	23.21	11.05
20	56.41	15.83

Table 8.11: TC2, differences and STD.

Time (s)	Mean - Exp (K)	STD (K)
15	16.20	2.12
20	25.40	4.30
25	21.02	6.72
30	11.40	9.22
35	2.29	11.66
40	19.29	14.00

Table 8.12: TC3, differences and STD.

Time (s)	Mean - Exp (K)	STD (K)
15	3.14	0.25
20	6.22	0.81
25	17.65	1.68
30	23.88	2.78
35	23.10	4.02
40	19.43	5.36

Table 8.13: TC4, differences and STD.

Time (s)	Mean - Exp (K)	STD (K)
5	23.69	2.93
8	8.42	3.15
10	7.58	3.20
12	11.54	3.61
15	21.81	4.94
20	33.86	7.81

Table 8.14: TC10, differences and STD.

Overall, these results, as the ones calculated using set 6, show that the posterior distributions obtained are reliable, given that they predict with good accuracy the other thermocouples' temperature evolution. For what concerns the temperature estimation in TC5 and TC9, both sets lead to very close results, so no major differences in using one or the other are registered.

8.4 Correlation between α and q_0

Inferring q_0 and α within the same simulation allows the calculation of a correlation matrix through *UQLab*. This 2x2 matrix expresses the correlation between the two inferred parameters, with values ranging from -1 to 1. A correlation equal to 1 expresses a perfect positive correlation, indicating that when one of the two increases, the other increases proportionally. At the same time, a correlation of -1 indicates a perfect negative correlation, that is when a parameter increases, the other decreases proportionally. Full independence between the two is implied by a 0, even though the inverse is not always true by necessity.

The correlation matrix thus provides valuable information on how the parameters interact. This can be particularly useful for identifying non-obvious relationships between parameters or for better understanding the structure of the model used within the inference framework. A correlation matrix calculated through *UQLab*, for two inferred parameters, has the following structure:

$$A = \begin{array}{|c|cc|} \hline & q0 & \alpha \\ \hline q0 & 1 & A_{1,2} \\ \alpha & A_{2,1} & 1 \\ \hline \end{array}. \tag{8.2}$$

Along the matrix's diagonal values will always be equal to 1, as every variable is perfectly correlated with itself, while nondiagonal values express the actual correlation between the two variables. Also, values in position $A_{2,1}$ and $A_{1,2}$ will always be identical since correlation is a measure of linear association that does not depend on the order of the variables. Said so, we will later refer to correlation factors and not matrices, only mentioning the values of $A_{2,1}$ (equal to $A_{1,2}$).

Correlation factors have been calculated for every inference analysis performed in Section 8.1 with 1000 iterations each, and reported in Table 8.15. Particular attention has to be given to the most informative sets, such as sets number 6, 8, 10, 11, and 12. For these sets, even analyses with more iterations confirmed a correlation factor of -0.93 or -0.94.

In fact, as can be seen from the Table and the next Figures, the sets give a strong (nearly perfect) negative correlation between α and q_0. This means that when α decreases, q_0 is expected to increase. This type of correlation may be due to the fact that a decrease in α is led by a decrease in $k_{\parallel}$, given that $k_{\perp}$ is constant. So, as thermal conductivity in the radial direction decreases, and so the material isolation increases, a higher value of q_0 is then required to obtain the same temperature in a specific TC location.

Also, as Table 8.15 shows, sets not including radial thermocouples, such as set numbers 1, 2, 3, or 4, which contains axials TCs, all show lower values of correlation between the inferred parameters, when compared to the other sets. This seems to show that the correlation intensifies when a radial thermocouple is introduced.

Perhaps, this may be also due to a non-investigated relation between the information gain and the correlation factor. In fact, all the sets that show low values of correlation, also show low KL values as reported in Table 8.1.

Set	Thermocouples	Correlation Factor
1	TC3	-0.80
2	TC4	-0.56
3	TC5	0.02
4	TC6	-0.12
5	TC7	-0.88
6	TC8	-0.94
7	TC5-TC7	-0.87
8	TC5-TC8	-0.94
9	TC6-TC7	-0.87
10	TC6-TC8	-0.94
11	TC5-TC7-TC8	-0.93
12	TC6-TC7-TC8	-0.94

Table 8.15: Correlation factors for every set simulated, σ^2 = 100 K.

Overall, a correlation between α and q_0 indicates that the value of one parameter depends on the value of the other. This leads to the conclusion that experiments performed by Milos and Chen [4] are probably not suitable to determine their value separately, as they instead give information about (α, q_0) couples. However, this correlation is most probably born by using our specific thermal model within the framework.
In fact, in reality, these two parameters are probably not related, as α is a characteristic of the material, while q_0 is the stagnation point cold wall heat flux, and they are assessed with different methods.

In the next Figures, the joint posterior PDF of the two variables α and q_0 can be visualized. The attention is given to the three most informative sets for which posterior distributions have been discussed in Section 8.2.

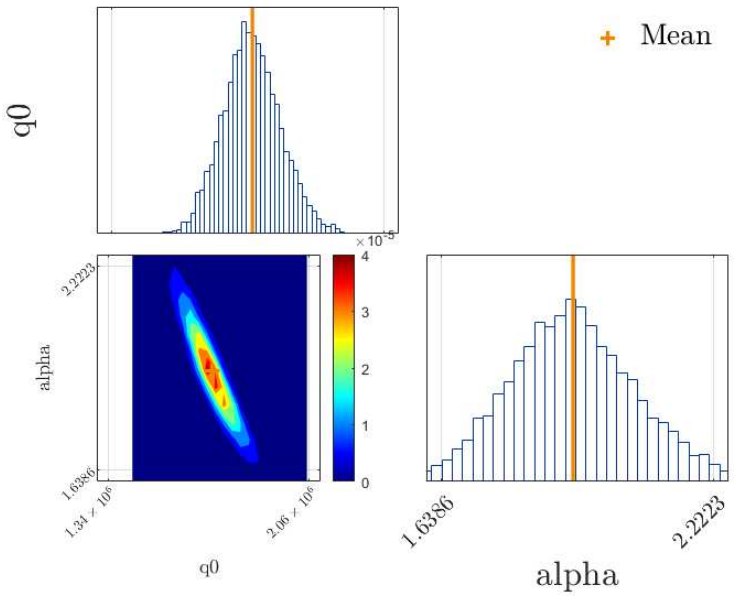

Figure 8.18: Variables distributions and joint PDF, obtained with TC8.

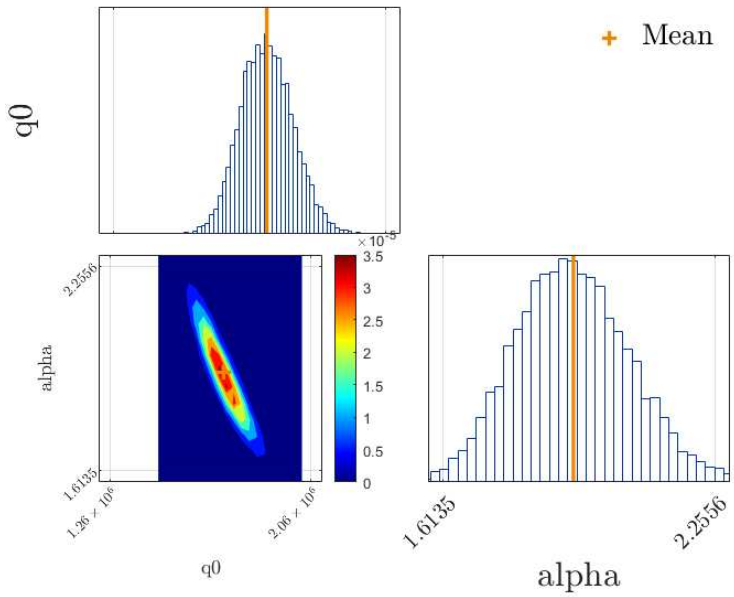

Figure 8.19: Variables distributions and joint PDF, obtained with TC5-TC8.

The joint PDF is visualized as a heatmap. The color intensity within the heatmap corresponds to the density of samples: regions with warmer colors (red to yellow) indicate higher densities of samples, whereas cooler colors (blue) represent lower densities. The elliptical shape of the high-density region on the heatmap indicates a strong negative correlation between q_0 and α.

The variable's marginal distributions are displayed on the top and right sides of the heatmap. These histograms represent the frequency distributions of the samples for the parameters, as derived from the Bayesian inference analysis. To transition from these discrete histograms to smooth curves, mostly used in the context of this thesis, a kernel density estimation (KSD) technique has been employed using the *ksdensity* command in *MATLAB*. This transition has been done as it helps visualize the density of samples in a more interpretable form.

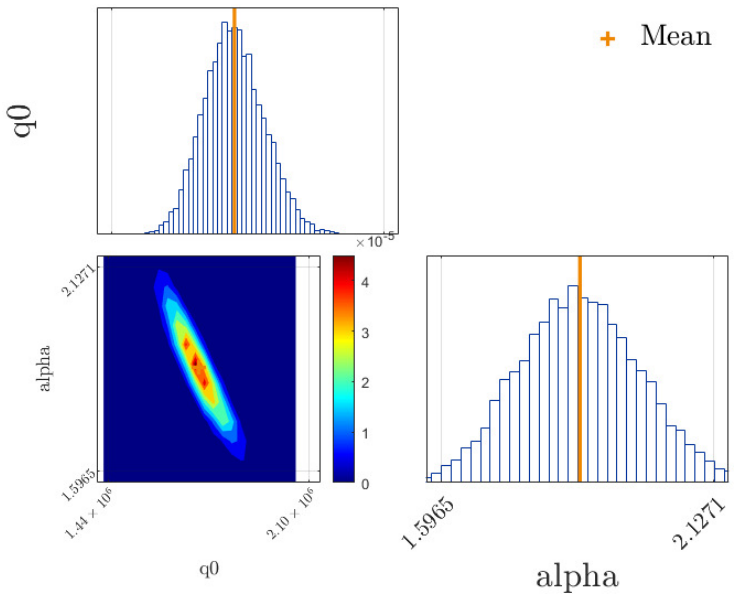

Figure 8.20: Variables distributions and joint PDF, obtained with TC6-TC7-TC8.

However, as correlation factors for sets number 3 and 4 show, some sets define a weaker correlation between the parameters, when compared to the rest of the sets simulated. This suggests the possibility of performing inference with a sequential analysis, in which one parameter is inferred using a specific set, and the other is inferred using a different one. In particular, as shown in Table 8.1, sets 3 and 4 are found to be more informative about q_0, as they do not include radial thermocouples.

Between these two, set 4 offers more information gain about q_0 than set 3, as it includes TC6, which is closer to the surface and hence feels more the heat flux conducted through the sample. Given that, a sequential analysis could be employed by inferring q_0 using set 4 and inferring α using a different set that offers a high KL on this parameter.

8.5 Sequential analysis

Following the considerations just discussed, a sequential analysis has been performed. In order to verify its reliability, results have been compared with the ones previously obtained by inferring both parameters within the same simulation and using the most informative sets. In particular, a sequential analysis has been employed by inferring q_0 using set 4, containing TC6, and inferring α using set 6, which includes TC8. Set 4 has been chosen between the two sets which define an almost null correlation, sets 3 and 4, as it offers the highest values of KL on both parameters, hence being more informative. In particular, its KL on q_0 is way higher than the one for α, as displayed in Table 8.1, so it has been used for estimating this parameter.

The first analysis has been run by setting α as constant and equal to 1.90 when inferring q_0 through set 4. This value has been chosen as it represents the arithmetic mean between the mean values of the posterior distributions of the most informative sets reported in Section 8.2. After inference has been performed on q_0, the mean value of its posterior distribution has been taken and used as a constant value for q_0 when inferring α through set 6.

Both analyses of inference have been run with 100 parallel chains and 300000 iterations, using the AIES sampler. When inferring q_0, its Seed has been chosen as equal to the Case 2 reported in Table 7.7, while, when inferring α, to the Case 1 reported in Table 7.3.

Set 4: Inference on q_0

The inference performed on q_0, as can be depicted in the next Table and Figure, led to very different results when compared to the ones obtained with the most informative sets. Both the mean and standard deviation values of the posterior are very different from the values previously calculated, such as the ones reported in Table 7.9, for example.

Parameter	Mean	STD
q_0	191.73	60.18

Table 8.16: Mean and STD values of the posterior, reported in W/cm^2.

This disagreement is most probably due to the very low value of Kullback Divergence on this parameter that this set gives. In fact, the KL value on q_0 for set 4 is 0.15, which is way lower when compared to the ones obtained for the most informative sets, for example, 2.11 for set 6, as reported in Table 8.1. This is made visually very clear when looking at the posterior PDF obtained, which is really wide, indicating a high inaccuracy.

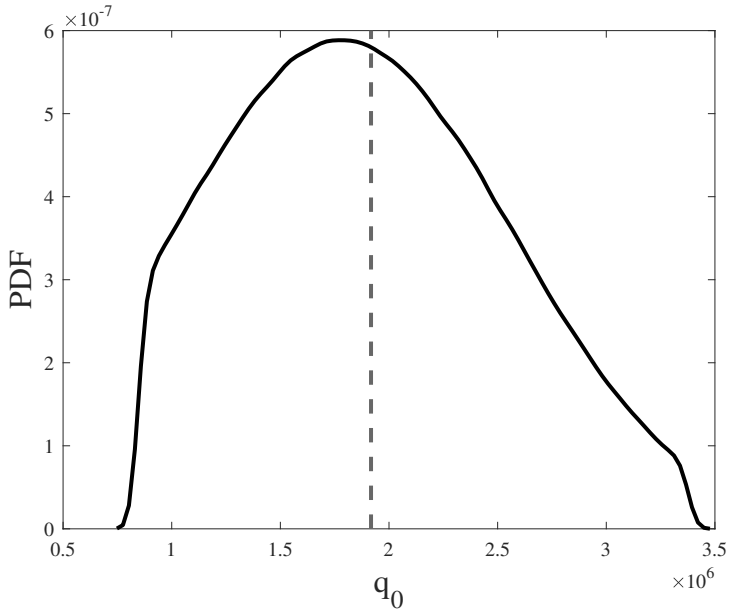

Figure 8.21: Posterior PDF of q_0, with dashed line indicating the mean value. Values are in W/cm^2.

Set 6: Inference on α

The other parameter, α, has then been inferred by taking q_0 constant and equal to 191.73 W/cm^2, the mean value of the posterior previously obtained.

Parameter	Mean	STD
α	1.6	0.035

Table 8.17: Mean and STD values of the posterior, reported in W/cm^2.

As can be visualized from the next Table and Figure, the inaccuracy regarding q_0 reflects on the estimation of α too. So, even though set 6 offers a really high KL on this parameter, the inaccurate estimation previously obtained through set 4 leads to obtaining a posterior distribution of α which is still very different from the ones obtained through the most informative sets. This is a common aspect of sequential analyses, where the first analysis always introduces a bias into the second one.

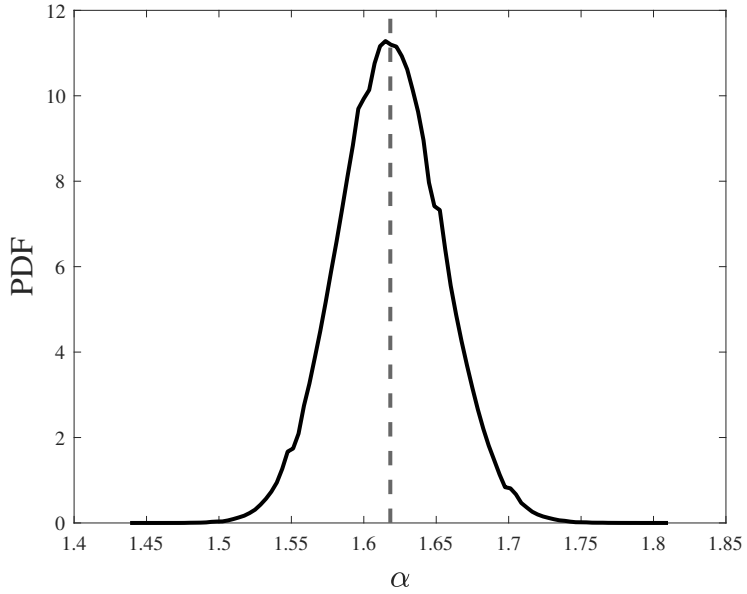

Figure 8.22: Posterior PDF of α, with dashed line indicating the mean value.

Chapter 9

Conclusions

The work performed in this thesis allowed for assessing the thermal conductivities ratio and the stagnation point cold wall heat flux, respectively α and q_0, through a Bayesian inference analysis.

Before running the inference, the studies performed assessed the independence of the inference process with respect to three different Seeds and to the Sampler chosen for the MCMC mechanism. The convergence of the 100 parallel chains has also been investigated, for both parameters, by using the Gelman-Rubin statistics. For what concerns the discrepancy model, two values have been found by inferring the variance of the discrepancy distribution. By performing these analyses, the inference framework later used has been set.

The inference analyses have then been run by using three of the most informative thermocouple sets, early defined by using the KL Divergence. The discrepancy used was 100 K. Analyses using a discrepancy of 36 K did not highlight major differences. Using a lower discrepancy variance simply allowed for obtaining higher KL values and lower STDs, as was expected, but without leading to changes on which thermocouple sets were more informative. The sets used, with a 100 K discrepancy variance, were set 6, containing TC8, set 8 containing TC5 and TC8 and finally set 12, containing TC6, TC7, and TC8. Through these, the nominal values of the two variables early declared as constant from the work of Milos and Chen [4] have been refined. In particular, the analyses allowed for obtaining posterior log-normal statistical distributions, finally accounting for uncertainties, as displayed in Section 8.2.

Of the posteriors obtained, the mean value of α has been found to be 1.92 using two of the most informative sets, set 6 and 8, while equal to 1.87 using set 12. For what concerns q_0, its value has been found to be 171.26 W/cm^2 by using sets 6 and 8, and 173.70 W/cm^2 using set 12. These, along with a visualization of the distributions, have been reported in Chapter 8.

Between the three sets used, set 12, demonstrated a slightly higher information gain for both variables, as reported in Table 8.1. Also, set 12 showed a lower value of standard deviation of the posterior distributions for both variables, when compared to the other two sets (i.e. a narrower and so more precise distribution). Following these considerations, it can be said that this set is probably slightly more reliable than the other two in estimating the parameters' distributions.

The results obtained (i.e. the posterior distributions) with these three most informative sets have later been tested, by using them to estimate the temperature evolution in other thermocouples. These forward propagation analyses showed that the posteriors calculated using these three sets can estimate with good accuracy the temperature over time in other TCs with the only great disagreement found to be in estimating the temperature in TC1. As these analyses showed that the posteriors previously calculated can estimate with good precision the experimental data in other thermocouples, they have ultimately demonstrated that the previously found posteriors are reliable.

The inference performed on both parameters within the same simulation also allowed for the investigation of the correlation factor between the two. In particular, this value has been calculated for all the sets tested, as reported in Table 8.15. The most informative sets (i.e. the most reliable) defined a correlation of -0.94. This strong negative correlation defines that when α decreases, q_0 is expected to increase almost proportionally. The correlation between the two parameters indicates that they are dependent, and so, their value cannot be determined separately. However, two sets, set 3 and 4, defined an almost null correlation between the two. This result has suggested the possibility of running an inference analysis sequentially, by inferring one parameter with one of those two sets, and the other parameter with a different one.

The sequential analysis has so been performed by inferring q_0 using set 4 and inferring α using set 6. The sequential analysis highlighted very different results from the ones obtained with the sets previously tested by inferring both parameters at the same time. This is most probably due to the fact that the KL Divergence for set 4 (even though being higher than the one in set 3) is still very low compared to the most informative sets. For instance, the KL, for q_0, is 0.15 for set 4, while being 2.11 for set 6.

This lack of information, besides providing a value of q_0 that is very different from the previous ones when inferring this variable, has also been reflected in the inference regarding the variable α, since the average value of the previously found q_0 distribution was then used in this simulation. At the same time, both set 3 and set 4, give even lower KL values about α. The sequential analysis employed has so been found to be not reliable since set 3 and set 4 are not informative about either one of the two parameters. Nevertheless, the possibility of performing a sequential analysis could still be really interesting in the context of this work, given that by inferring other parameters higher KL values could be found in the sets that give an almost zero correlation.

Appendix

1.

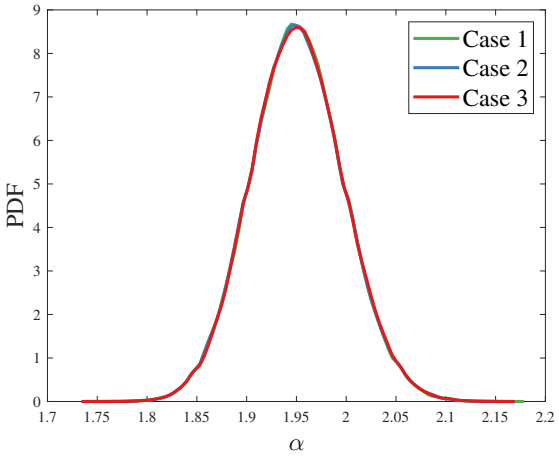

Figure 9.1: Posterior PDFs of α with different Seeds for TC8.

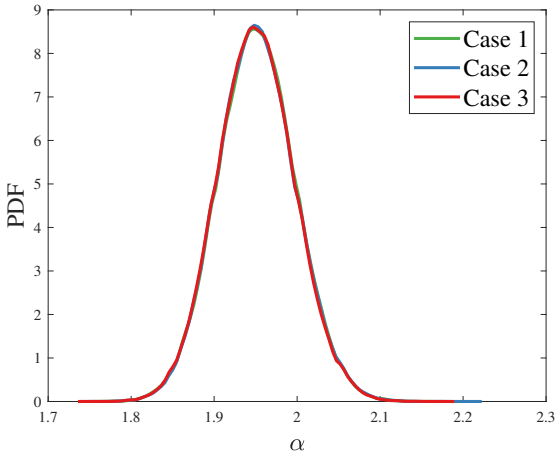

Figure 9.2: Posterior PDFs of α with different Seeds for TC6-TC8.

2.

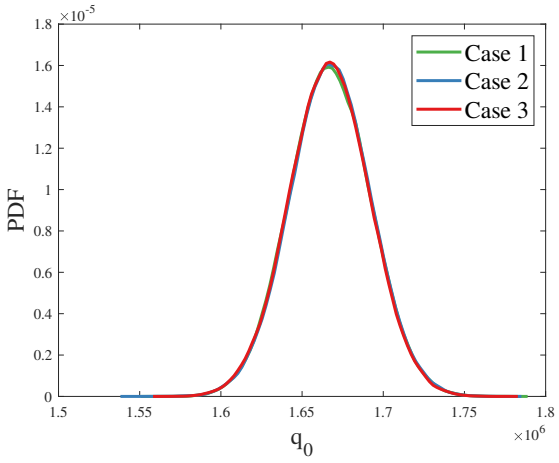

Figure 9.3: Posterior PDFs of q_0 with different Seeds for TC8.

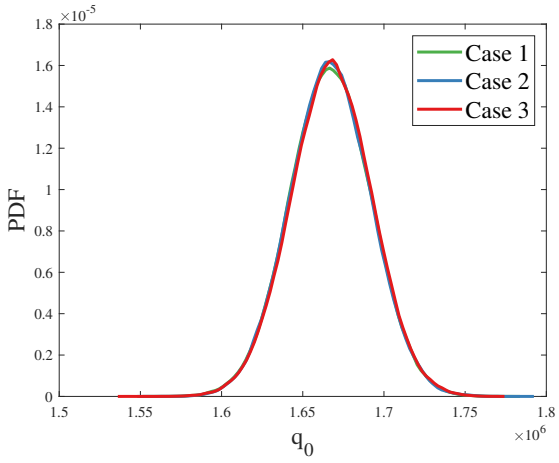

Figure 9.4: Posterior PDFs of q_0 with different Seeds for TC6-TC8.

3.

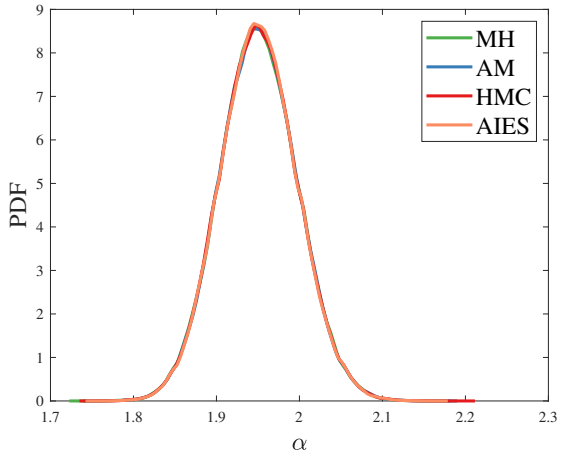

Figure 9.5: Posterior PDFs of α with different samplers for TC8.

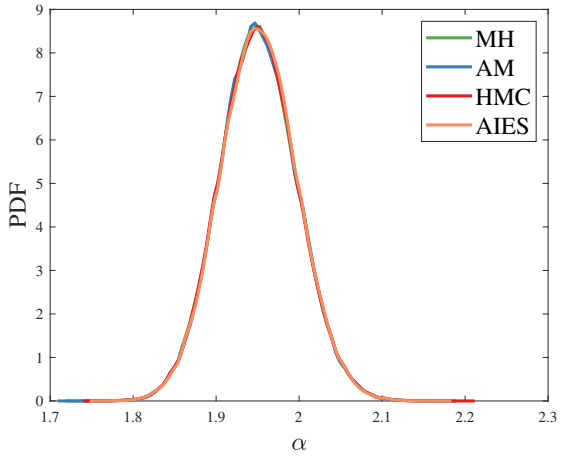

Figure 9.6: Posterior PDFs of α with different samplers for TC6-TC8.

4.

Set	**Thermocouples**	**α Mean**	**α STD**	**q_0 Mean**	**q_0 STD**	**α KL**	**q_0 KL**
1	TC3	1.86	0.30	172.38	3.72	0.12	2.85
2	TC4	2.01	0.35	142.57	6.80	0.01	2.17
3	TC5	2.06	0.36	217.79	73.74	0.01	0.06
4	TC6	2.00	0.34	178.00	44.90	0.01	0.36
5	TC7	1.94	0.31	157.28	37.21	0.04	0.82
6	TC8	1.91	0.08	171.50	4.85	0.96	2.58
7	TC5-TC7	1.94	0.31	157.26	27.46	0.04	0.81
8	TC5-TC8	1.91	0.09	171.50	4.87	0.96	2.58
9	TC6-TC7	1.89	0.30	161.59	26.67	0.08	0.84
10	TC6-TC8	1.91	0.08	171.53	4.86	0.96	2.58
11	TC5-TC7-TC8	1.86	0.07	173.92	4.68	1.08	2.62
12	TC6-TC7-TC8	1.86	0.07	174.09	4.67	1.10	2.62
13	TC5-TC6-TC7-TC8	1.86	0.07	174.08	4.67	1.09	2.62

Table 9.1: Inference for $\sigma^2 = 36$ K, q_0 in W/cm^2. Mean and STD referred to posterior PDFs. KL have been calculated with an increment of $\varepsilon = 10^{-10}$.

5.

Set	Thermocouples	$\varepsilon = 10^{-12}$	$\varepsilon = 10^{-10}$	$\varepsilon = 10^{-08}$
1	TC3	0.0206	0.0206	0.0206
2	TC4	0.0088	0.0088	0.0086
3	TC5	0.0073	0.0073	0.0072
4	TC6	0.0093	0.0093	0.0092
5	TC7	0.0167	0.0167	0.0165
6	TC8	0.5306	0.5306	0.5306
7	TC5-TC7	0.0298	0.0298	0.0298
8	TC5-TC8	0.5253	0.5253	0.5253
9	TC6-TC7	0.0372	0.0372	0.0372
10	TC6-TC8	0.5090	0.5090	0.5090
11	TC5-TC7-TC8	0.6553	0.6553	0.6553
12	TC6-TC7-TC8	0.6641	0.6641	0.6641
13	TC5-TC6-TC7-TC8	0.6478	0.6478	0.6475

Table 9.2: α KL calculation dependency on ε, σ^2 = 100 K.

6.

Set	Thermocouples	$\varepsilon = 10^{-12}$	$\varepsilon = 10^{-10}$	$\varepsilon = 10^{-08}$
1	TC3	0.1189	0.1189	0.1184
2	TC4	0.0098	0.0098	0.0096
3	TC5	0.0092	0.0092	0.0089
4	TC6	0.0076	0.0075	0.0075
5	TC7	0.0355	0.0355	0.0352
6	TC8	0.9640	0.9640	0.9640
7	TC5-TC7	0.0387	0.0387	0.0386
8	TC5-TC8	0.9610	0.9610	0.9610
9	TC6-TC7	0.0786	0.0786	0.0785
10	TC6-TC8	0.9608	0.9608	0.9603
11	TC5-TC7-TC8	1.0789	1.0788	1.0788
12	TC6-TC7-TC8	1.0982	1.0982	1.0980
13	TC5-TC6-TC7-TC8	1.0902	1.0902	1.0901

Table 9.3: α KL calculation dependency on ε, σ^2 = 36 K.

7.

Set	Thermocouples	$\varepsilon = 10^{-12}$	$\varepsilon = 10^{-10}$	$\varepsilon = 10^{-08}$
1	TC3	2.6170	2.6170	2.6168
2	TC4	1.9128	1.9128	1.9127
3	TC5	0.0586	0.0585	0.0585
4	TC6	0.1549	0.1549	0.1549
5	TC7	0.7102	0.7102	0.7103
6	TC8	2.1061	2.1061	2.1060
7	TC5-TC7	0.7033	0.7033	0.7031
8	TC5-TC8	2.1142	2.1140	2.1140
9	TC6-TC7	0.7467	0.7467	0.7467
10	TC6-TC8	2.1028	2.1027	2.1027
11	TC5-TC7-TC8	2.1697	2.1697	2.1695
12	TC6-TC7-TC8	2.1692	2.1692	2.1690
13	TC5-TC6-TC7-TC8	2.1685	2.1685	2.1683

Table 9.4: q_0 KL calculation dependency on ε, σ^2 = 100 K.

8.

Set	Thermocouples	$\varepsilon = 10^{-12}$	$\varepsilon = 10^{-10}$	$\varepsilon = 10^{-08}$
1	TC3	2.8546	2.8546	2.8545
2	TC4	2.1704	2.1703	2.1703
3	TC5	0.0586	0.0585	0.0585
4	TC6	0.3644	0.3644	0.3642
5	TC7	0.8171	0.8171	0.8170
6	TC8	2.5809	2.5809	2.5809
7	TC5-TC7	0.8066	0.8066	0.8065
8	TC5-TC8	2.5789	2.5789	2.5789
9	TC6-TC7	0.8378	0.8378	0.8375
10	TC6-TC8	2.5805	2.5805	2.5804
11	TC5-TC7-TC8	2.6152	2.6152	2.6150
12	TC6-TC7-TC8	2.6182	2.6182	2.6181
13	TC5-TC6-TC7-TC8	2.6178	2.6178	2.6174

Table 9.5: q_0 KL calculation dependency on ε, σ^2 = 36 K.

9.

Time (s)	Mean - Exp (K)	STD (K)
15	24.11	0.08
20	11.89	0.02
25	0.69	0.01
30	3.49	0.05
35	5.05	0.01
40	4.73	0.02

Table 9.6: TC5 estimation, differences and STD, obtained with TC8 posterior.

10.

Time (s)	Mean - Exp (K)	STD (K)
15	24.11	0.08
20	11.91	0.02
25	0.69	0.01
30	3.59	0.05
35	5.05	0.01
40	4.74	0.02

Table 9.7: TC5 estimation, differences and STD, obtained with TC6-TC7-TC8 posterior.

11.

Time (s)	Mean - Exp (K)	STD (K)
5	40.57	3.72
8	23.23	3.83
10	11.07	3.74
12	0.49	4.41
15	13.14	6.63
20	19.93	11.14

Table 9.8: TC9 estimation, differences and STD, obtained with TC8 posterior.

12.

Time (s)	Mean - Exp (K)	STD (K)
5	42.18	3.36
8	24.81	3.57
10	12.19	3.61
12	0.02	4.35
15	13.66	6.50
20	22.17	10.79

Table 9.9: TC9 estimation, differences and STD, obtained with TC6-TC7-TC8 posterior.

Bibliography

[1] Stefano Marelli and Bruno Sudret. UQLab: A Framework for Uncertainty Quantification in MATLAB. *2nd International Conference on Vulnerability and Risk Analysis and Management*, University of Liverpool, United Kingdom:pp. 2554–2563, (ICVRAM 2014).

[2] Stewart V. McDougall Fabrizio Bisetti Aditya Vinod. Early response of ablative materials to arcjet testing. *AIAA SCITECH 2023 Forum*, 2023.

[3] NASA Ames Research Center. https://reentry.arc.nasa.gov/. *Hypervelocity Reentries*.

[4] F.S. Milos Y-K. Chen. Ablation and thermal response property model validation for Phenolic Impregnated Carbon Ablator. *Journal of spacecraft and Rockets*, 2010.

[5] Panerai Magin Anst Mansour Blondeau Torres Herrador, F.Coheur. Competititve kinetic model for the pyrolisis of the Phenolic Impregnated Carbon Ablator. *Aerospace Science and Technology Vo. 100*, 2020.

[6] F.S. Milos and Y.-K. Chen. Two-Dimensional Ablation, Thermal Response, and Sizing Program for Pyrolyzing Ablators. *Journal of Spacecraft and Rockets*, Vol.46(6):1089–1099, 2009.

[7] F.S. Milos and Y.-K. Chen. Ablation and Thermal Response Program for Spacecraft Heatshield Analysis. *Journal of Spacecraft and Rockets*, Vol. 36 No.3, 1999.

[8] Driver D. M. Wright M. J. Hwang H. H. Edquist K. T. Beck, R. A. S. and S. A. Sepka. Development of the Mars Science Laboratory Heatshield Thermal Protection System. *Journal of Spacecraft and Rockets,*, Vol. 51(No. 4):pp. 1139–1150, 2014.

[9] K. A. Trick and T.E. Saliba. Mechanisms of the pyrolysis of phenolic resin in a carbon/phenolic composite. *Carbon*, Vol. 33(No. 11):pp. 1509–1515, 1995.

[10] Luke T. Samuel N. Surviving the heat: the application of Phenolic Impregnated Carbon Ablators. *University of Pittsburgh Swanson School of Engineering*.

[11] Ellerby D. T. Switzer M. R. Agrawal, P. and T. H. Squire. Multidimensional Testing of Thermal Protection Materials in the Arcjet Test Facility. *10th AIAA/ASME Joint Thermophysics and Heat Transfer Conference*, 2010.

[12] J. Lachaud and N. N. Mansour. Porous-Material Analysis Toolbox Based on OpenFOAM and Applications. *Journal of Thermophysics and Heat Transfer*, Vol. 28(No. 2):pp. 191–202.

[13] Scoggins J. Magin T. Meyer M. Lachaud, J. and N. Mansour. A generic local thermal equilibrium model for porous reactive materials submitted to high temperatures. *International Journal of Heat and Mass Transfer*, Vol. 108:pp. 1406–1417, 2017.

[14] van Eekelen T. Scoggins J. B. Magin T. E. Lachaud, J. and N. N. Mansour. Detailed chemical equilibrium model for porous ablative materials. *International Journal of Heat and Mass Transfer*, Vol. 90:pp. 1034–1045, 2015.

[15] Lachaud J. Panerai F. Tang C. Meurisse, J. B. E. and N. N. Mansour. Multidimensional material response simulations of a full-scale tiled ablative heatshield. *Aerospace Science and Technology*, Vol. 76:pp. 497–511, 2018.

[16] Johnson C.E. Rasky D. J. Hui F. Y. C. ta Hsu M. Chen T. S. Chen Y.-K. Paragas D. Tran, H.K. and L. Kobayashi. Phenolic Impregnated Carbon Ablators (PICA) as Thermal Protection Systems for Discovery Missions. *NASA Technical Memorandum*, 110440, 1997.

[17] Heinemann J. M. Goldstein H. E. Chen Y.-K. Terrazas-Salinas I. Balboni J. A. Olejniczak-J. Covington, M. A. and E. R. Martinez. Performance of a Low Density Ablative Heat Shield Material. *Journal of Spacecraft and Rockets*, Vol. 45, No. 2:pp 237–247, 2008.

[18] C.B. Moyer and M. R. Wool. Aerotherm Charring Material Thermal Response and Ablation program. *User's manual, Aerotherm Corp.*, 1970.

[19] Tyler Dae Devlin Daniel Xiang Daniel Kunin, Jingruo Guo. seeing-theory.brown.edu. *Seeing Theory*.

[20] Donald B. Rubin Andrew Gelman. Inference from Iterative Simulation Using Multiple Sequences. *Statist. Sci. 7 (4) 457 - 472*, November, 1992.

[21] Johnson C. Rasky D. Hui F. Hsu M.-T Tran, H. and Y. Chen. Phenolic Impregnated Carbon Ablators (PICA) for Discovery class missions. 1996.

[22] Ludovic Benistant Ben Huberman. https://towardsdatascience.com. *Towards Data Science*.

[23] Ananya19b. https://www.analyticsvidhya.com. *Analytics Vidhya*.