Quantification of model-form uncertainties affecting the calibration of a carbon nitridation model by means of Bayesian Model Averaging

Anabel del Val, Thierry Magin, Pietro Marco Congedo

Quantification of model-form uncertainties affecting the calibration of a carbon nitridation model by means of Bayesian Model Averaging

Anabel del Val[a,b,*], Thierry E. Magin[a], Pietro M. Congedo[b]

[a] *von Karman Institute for Fluid Dynamics, Chaussée de Waterloo 72, 1640 Rhode-St Genèse, Belgium*
[b] *Inria, Centre de Mathématiques Appliquées, École Polytechnique, IPP, Route de Saclay, 91128 Palaiseau Cedex, France*

Abstract

Severe epistemic uncertainties not only can affect the prescription of parameters within a given model but also the choice of models we make to interpret and infer from experimental data. In this work, we incorporate experimental, parametric and model-form uncertainties in the calibration of a carbon nitridation model. The model-form uncertainties considered stem from the different modeling choices that are taken as valid representations of a set of plasma wind tunnel experiments. To this end, we define a Bayesian model averaging strategy where the marginal posteriors of the nitridation reaction efficiencies are weighted by the marginalized likelihoods of the experimental data for each proposed model. First, Bayes factors are computed to possibly discard invalid models. The baseline model, a thermal equilibrium stagnation line flow with nitridation as only surface reaction, performs as well as all the alternative models proposed, which range from adding surface recombination reactions to considering thermal non-equilibrium in the gas and gas-surface interface. The presence of nitrogen recombination reactions is shown to broaden the support of the nitridation marginal posteriors considerably, allowing it to take on larger values. Lastly, a Bayesian model averaged Arrhenius law for the nitridation efficiencies is computed for a range of surface temperatures.

Keywords: Uncertainty Quantification, Bayesian Inference, Thermal Protection Systems, Ablation

1. Introduction

Space vehicles plunging into dense planetary atmospheres release a significant amount of kinetic energy in a very short window of time. This energy is transformed into thermal energy through the formation of a strong bow shock ahead of the vehicle. This harsh environment can generate extreme heat loads ($> 10,000$ J/cm^2) [1] which are to be partially absorbed by the vehicle's surface. Aerospace engineers have largely figured out how to efficiently protect the integrity of spacecraft as well as astronauts and cargo from the damage derived from such intense heating. Thermal Protections Systems (TPS) are designed and used for this purpose. In particular, ablative materials such as rigid carbon fibers or silicon composites, impregnated with an organic resin matrix, are broadly used to deal with surface heating for high speed atmospheric entries ($\sim$ 11km/s).

The characterization of ablation phenomena plays an important role in the development of theoretical models that can predict material degradation and recession of TPS, both important aspects in the design of ablative TPS [2]. As such, experiments and models are often used to understand the physics of ablation and improve our predictive capabilities. On the experimental side, the stochastic nature of the data must be accurately described and modeled to produce reliable experimental data on which to base our analyses. On the modeling side, many different sources of uncertainties can affect the predictions considerably. Objectively characterizing and quantifying such uncertainties is important to make comparisons to experimentally observed quantities useful, providing a more consistent/quantifiable way of defining new research directions.

Overall, the process of inferring ablation parameters from experimental data poses many questions that span beyond our experimental capabilities to the assumptions in our models. We still grapple with understanding our own

*Corresponding author: von Karman Institute for Fluid Dynamics, Chaussée de Waterloo 72, 1640 Rhode-St-Genèse, Belgium. *Email address:* ana.isabel.del.val.benitez@vki.ac.be (A. del Val)

experimental facilities to the point that we must admit ignorance about the relevant physical processes for different experimental conditions in some cases. While experiments are giving us answers to questions, sometimes we must admit that it is not clear what are the questions to which the answers pertain and this is significantly slowing down our progress. In this regard, severe lack of knowledge greatly affects our capabilities to build predictive models.

In particular, the carbon nitridation reaction $C_s + N \rightarrow CN + 0.34eV$ on a solid (s) carbon surface is still hard to predict accurately at temperatures above 1,000 K. The models found in the aerothermodynamics literature are derived empirically by fitting experimental data [3–7]. Even in the cases where uncertainty is taken into account by performing stochastic calibrations, the results only reflect experimental and parametric uncertainty within the given structure of a model [8, 9]. In the cases where different models were entertained, such as in the work of Upadhyay et al. [8], the overall resulting uncertainty in the final calibrated model was never discussed on the basis of model predictive performance and complexity, which prevented such model-form uncertainties from being characterized and included in the end results.

In this work, we address the fact that there exists severe lack of knowledge in the calibration of carbon nitridation reaction efficiencies from plasma wind tunnel data. Experimental and parametric uncertainty, together with a selection of model scenarios, are taken into account to produce a comprehensive picture of nitrogen ablation for selected plasma wind tunnel testing conditions. We use the experimental dataset produced by Helber et al. [10] for which they subjected a graphite material sample to a subsonic plasma flow. They recorded material recession rates and local CN mass densities in the boundary layer by means of spectroscopy measurements. The fact that the material tested is graphite in a pure nitrogen flow gives us an unprecedented window for testing hypotheses concerning the flow and wall states without having to add important uncertainties stemming from other physical processes such as pyrolysis and oxidation. We consider several modeling scenarios to explain the experimental data and predict the observed recession rates and CN densities. The baseline model is a thermal equilibrium stagnation line flow with only nitridation as surface mechanism. Calibrations performed for this model indicate that its assumptions hold for the experimental conditions considered. We use the baseline model as reference or null hypothesis against which three alternative models of additive physical richness are compared.

The first alternative model adds nitrogen recombination reactions at the gas-surface interface by incorporating a recombination efficiency parameter. As recombination reactions compete with nitridation reactions for the available atomic nitrogen, it is important to evaluate their impact in our calibrations. Moreover, we are also interested in learning about nitrogen recombination from these experiments. A second alternative model expresses the same physics but adopts a Surface Energy Balance (SEB) to resolve the surface temperature instead of imposing it as a given boundary condition. The SEB model reduces the complexity in terms of number of parameters to be inferred, given that the surface temperature is now a prediction of the model instead of another free parameter to be imposed. Furthermore, the new parametrization that the SEB introduces comes with more constrained parameter spaces such as the material emissivity. On the other hand, this model has to be tested against an additional observation, the surface temperature. The last alternative model also retains all the physics of the previous models but relaxes the conditions of thermal equilibrium at the gas and gas-surface interface. It remains agnostic about the thermal state of the gas by assuming two distinct thermal baths. At the gas-surface interface, it solves a SEB for each bath, introducing yet a new parametrization that also weighs in the comparison of models.

First, each model proposed is calibrated separately for each experimental condition. The hypothesis testing study is then made through the computation of Bayes factors which compare the resulting nitridation reaction efficiencies marginal posteriors at the order of magnitude level between the baseline and each alternative model. Two different aspects affect the computation of such metric: the goodness-of-fit to the observations that each model can predict as well as their complexities in terms of number of model parameters to be inferred and the extent of their prior spaces. The final nitridation reaction efficiency posterior distributions obtained under each successful model are weighted according to their marginalized likelihoods, and summed using Bayesian Model Averaging (BMA) to account for the epistemic uncertainties stemming from these different valid modeling assumptions. Further, an additional original contribution of this work is to combine all available measurements for the calibration of stochastic Arrhenius laws under each considered model, extending the computation of nitridation efficiencies to a broader span of surface temperatures for which there are no reliable experimental data. The marginal posteriors of the Arrhenius laws parameters are then weighted by their corresponding marginalized likelihoods to generate a model-averaged Arrhenius law that incorporates model-form uncertainties for a broad span of surface temperatures. The overall aim of this work is then twofold: to incorporate model-form uncertainties in the characterization of a carbon nitridation model and to

invalidate or, conversely, find new research directions that would eventually lead to dropping models, refining our understanding of the physics in the process.

The article is organized as follows. Sec. 2 addresses the experimental data, theoretical models and numerical methods used in this work. Sec. 3 presents the Bayesian model averaging methodology in detail. Sec. 4 shows and discusses the obtained results. Sec. 5 presents the conclusions and outlook.

2. Overview of experimental data and modeling scenarios

We introduce a brief note on the experimental data used for this study followed by the different modeling choices we can make to retrieve the carbon nitridation model we are after. The experiments and models can be found in more detail in other published works [10–12] so we focus only on recalling the important aspects.

2.1. On the experimental data

The experimental scenario is depicted in much detail in [10]. In this section, we briefly describe the measurements and their associated uncertainties.

The experiments were performed to rebuild nitridation reaction efficiencies $\gamma_{\mathrm{N}}^{\mathrm{CN}}$ for graphite at high surface temperatures. For this, Helber et al. [10] experimentally observed the response of graphite samples to a subsonic plasma flow produced by the von Karman Institute's (VKI) Plasmatron facility (see Fig. 1). This process creates a reacting boundary layer around the exposed sample where species resulting from the ablation reactions are injected. Overall, Helber et al. were interested in measuring the rate of recession of the graphite sample due to the ablation reaction $\mathrm{C_s + N \rightarrow CN}$, as well as the radiation signature from the CN molecule injected in the boundary layer. Additional measurements such as dynamic pressure P_{d} within the plasma jet, and static pressure P_{s} in the Plasmatron chamber are also included in the dataset.

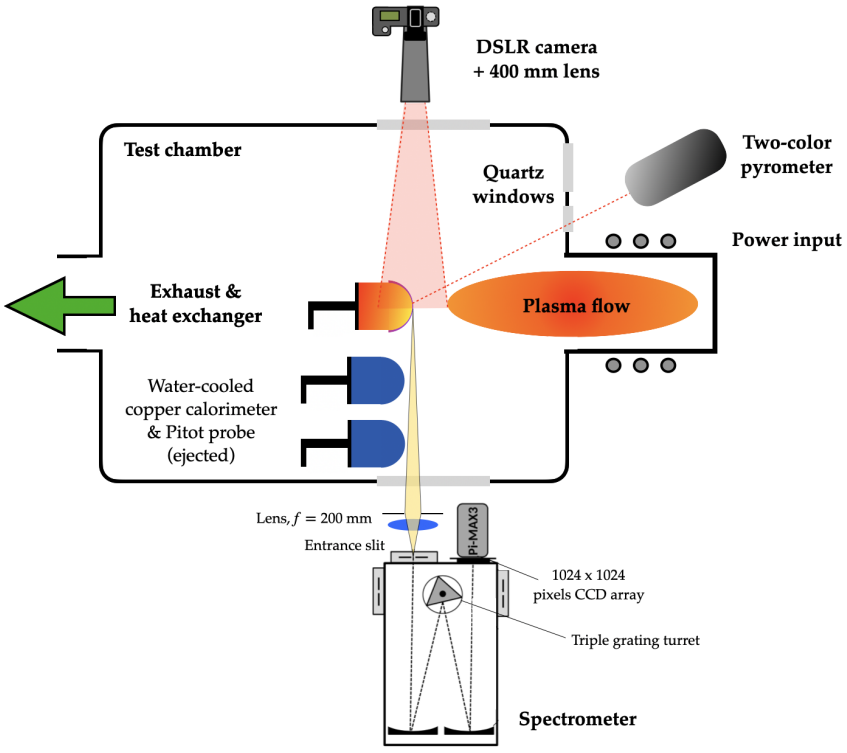

Figure 1: Schematic view of the ablation experimental set-up for the Plasmatron facility. Figure adapted from Helber [**?**]

All the different probes and test sample holders were exposed sequentially to the same plasma flow and the nitridation measurements were recorded until steady-state was maintained for 300 s. The following tables summarize the different measurements and uncertainties used in this work. In Table 1, the measured quantities are the dynamic pressure P_d, the wall temperature T_w, the chamber static pressure P_s, the recession rate $\dot{s}$ and the local CN densities ρ_{CN} estimated from spectra measurements locally resolved. For the particular case G6, CN density measurements are not taken into account given the inconsistencies identified [9].

Table 1: Experimental data.

Experiment ID	P_d [Pa]	T_w [K]	P_s [Pa]	$\dot{s}$ [μm/s]	ρ_{CN} [μg/m^3]
G4	231	2225	1500	1.41	385
G5	268	2410	1500	1.64	800
G6	312	2535	1500	2.60	-
G7	330	2575	1500	2.51	1050

The experimental uncertainties are assumed independently distributed Gaussian distributions $\mathcal{N}(\mu, \sigma)$ with σ taken considering the bounds given on the measurements as 2σ level of confidence [13–15].

Table 2: Experimental uncertainties.

Experiment ID	P_d [Pa]	T_w [K]	P_s [Pa]	$\dot{s}$ [μm/s]	ρ_{CN} [μg/m^3]
G4	$\mathcal{N}(231, 2.31)$	$\mathcal{N}(2225, 11.12)$	$\mathcal{N}(1500, 22.5)$	$\mathcal{N}(1.41, 0.26)$	$\mathcal{N}(385, 52)$
G5	$\mathcal{N}(268, 2.68)$	$\mathcal{N}(2410, 12.05)$	$\mathcal{N}(1500, 22.5)$	$\mathcal{N}(1.64, 0.27)$	$\mathcal{N}(800, 100)$
G6	$\mathcal{N}(312, 3.12)$	$\mathcal{N}(2535, 12.67)$	$\mathcal{N}(1500, 22.5)$	$\mathcal{N}(2.60, 0.27)$	-
G7	$\mathcal{N}(330, 3.3)$	$\mathcal{N}(2575, 12.87)$	$\mathcal{N}(1500, 22.5)$	$\mathcal{N}(2.51, 0.25)$	$\mathcal{N}(1050, 162)$

2.2. Baseline model

We use a one dimensional form of the full Navier-Stokes equations in order to perform efficient simulations of a high temperature, reacting flow with complex physico-chemical models and boundary conditions. Originally proposed in [16], this formulation is called "Dimensionally Reduced Navier-Stokes Equations (DRNSE)" and, in the present study, it is implemented in an in-house software developed by Munafò [11]. This code is able to reproduce the flow along the stagnation-line of spherical and cylindrical bodies [16].

In the solver, the spatial discretization is performed by means of a cell-centered finite volume method. The numerical advective fluxes are computed by means of the AUSM$^+$-up2 scheme [17], which splits the flux into a convective and a pressure terms. Second order accuracy in space is achieved by using an upwind reconstruction to obtain the cell interface variables. The diffusive fluxes and source terms are both evaluated in terms of primitive variables, namely, species densities ρ_i, velocity components u, v, and temperature T. The value of the primitive variables and their gradients at the volume interface are computed by a weighted average and a central finite difference approximation, respectively. The cell center value of the primitive variables to be used in the evaluation of the diffusive source terms are computed using a two-point central finite-difference approximation. A fully implicit scheme is used for the present study. The code is coupled with the Mutation^{++} library [12]. Species thermodynamic properties are obtained from Rigid-Rotor Harmonic-Oscillator computations, and relative mixture quantities are derived from pure species quantities. The transport properties are derived from kinetic theory, which provides relationships for macroscopic transport coefficients based on microscopic collision integrals. The chemical production rates for species, based on elementary chemical reactions, including third body, are calculated by taking the forward reaction rate coefficients specified by the user in an Arrhenius law form. The backward rate coefficient is determined by satisfying the equilibrium relation. The library has been designed, implemented and extensively tested to ensure high-fidelity together with low computational costs.

A coupled ablative boundary condition is also included in a generalized formulation [18] in Mutation^{++} based on surface mass and energy balances. Several other authors have applied this approach to study the ablative and non-ablative gas-surface interaction of thermal protection materials with the surrounding flowfield [19, 20]. The advective and the diffusive fluxes at the gas-surface interface are obtained by solving the surface species mass balance. This balance is obtained limiting the control volume of the mass conservation equation to the thin lamina representing the gas-surface interface [10]

$$\left[\rho_{\mathrm{w},i}(\boldsymbol{u}_{\mathrm{g}} - \boldsymbol{u}_{\mathrm{r}}) + \rho_{\mathrm{w},i}\boldsymbol{V}_{\mathrm{w},i}\right] \cdot \boldsymbol{n} = \dot{\omega}_{\mathrm{w},i}, \ \forall i \in \mathcal{S}, \tag{1}$$

where the source term $\dot{\omega}_{\mathrm{w},i}$ represents the sum for species i of the production rates of individual heterogeneous reactions with the surface $\dot{\omega}_{\mathrm{w},i} = \sum_{r=1}^{n_{\mathcal{R}_{\mathrm{w}}}} \dot{\omega}^{r}_{\mathrm{w},i}$, with $n_{\mathcal{R}_{\mathrm{w}}}$ being the number of considered reactions. For this model, we do not need to include a surface energy balance given that the temperature at the wall T_{w} is measured and imposed.

We need to introduce closure terms for the production of species $\dot{\omega}^{r}_{\mathrm{w},i}$. The approach for the production term considered in this work is a phenomenological one because there is only so much that we can learn from the present data. This type of models are quite popular in the aerothermodynamics literature for their straightforward implementation [21]. We are only concerned with ablation models for surface reactions involving nitrogen on the side of the flow and carbon on the material side. The model for the production rate is defined

$$\dot{\omega}^{r}_{\mathrm{w},i} = \gamma^{r}_{i} m_i \mathcal{F}^{\downarrow}_{i}, \ \forall i \in \mathcal{S}, \ \forall r \in n_{\mathcal{R}_{\mathrm{w}}}, \tag{2}$$

with γ^{r}_{i} representing the reaction efficiency, m_i the particle i mass, and $\mathcal{F}^{\downarrow}_{i}$ the impinging flux of species i.

The inlet boundary condition used for the simulations of such reacting flows comprises species densities ρ_i, temperature T, and velocity components u, v. The information needed to prescribe the boundary condition at the free stream pertains to a supersonic inlet even though we want to carry out subsonic simulations. This methodology has been investigated and discussed in [22], where it was verified that the solution of the Riemann problem at the boundary interface automatically screens the information needed from the inlet.

All in all, we intend to find the flow solutions that are compatible with our experimental observations under the considered model. This means that from the observations reported by Helber et al. [10], we want to infer the inlet conditions and the wall conditions in such a way that we are left with a population of possible flow solutions. After taking into account some physical considerations for the definition of the inlet quantities [9], for given Plasmatron operating conditions, we are left with the following expressions for the simulated counterparts of the observations of recession rate and local CN densities

$$\dot{s} = \dot{s}\left(P_{\delta}, T_{\delta}, P_{\mathrm{d}}, T_{\mathrm{w}}, \gamma^{\mathrm{CN}}_{\mathrm{N}}\right), \tag{3}$$

$$\rho_{\mathrm{CN}} = \rho_{\mathrm{CN}}\left(P_{\delta}, T_{\delta}, P_{\mathrm{d}}, T_{\mathrm{w}}, \gamma^{\mathrm{CN}}_{\mathrm{N}}\right), \tag{4}$$

where the subscript δ makes reference to the quantities defined at the boundary layer edge. These expressions constitute the dependencies of the forward model.

2.3. Alternative modeling scenarios

Several model assumptions stand at the basis of the baseline model. The flow is assumed in thermal equillibrium and the ablative wall conditions amount to only nitridation reactions. Furthermore, measuring the surface temperature also allows us to test ablation energy balance models by using the wall temperature as the quantity to be reproduced. In the available literature regarding experimental studies for nitridation, flow simulations are a needed component to compute atomic nitrogen concentrations at probe locations. The resulting nitridation reaction efficiencies rely heavily on such computations which are full of assumptions about the flow and material behaviors. Suzuki et al. [5, 23] used flow simulations that did not account for nitridation at the surface. Helber et al. [10] showed that the atomic nitrogen concentration profile along the stagnation line of the boundary layer is significantly different (~10%) when nitridation is turned on/off. Suzuki et al. [23] underestimated nitridation reaction efficiencies this way.

The thermal estate of the gas is not expected to impact the previous estimations of nitridation reaction efficiencies given that the recession rates and surface temperatures are accurately measured. Nevertheless, this set of experimental

data presents a great opportunity to study the evidence for thermal non-equilibrium and the adequacy of surface energy balance models. Together with the thermodynamics and energy wall conditions, nitrogen recombination reactions are added to our models. As recombination reactions compete with nitridation reactions for the available atomic nitrogen, it is important to evaluate their impact in our calibrations. Moreover, we are also interested in learning about nitrogen recombination from these experiments.

Overall, three different alternative model scenarios of additive physical richness are presented. They all build on the previous baseline model, $\mathcal{M}_0$:

- $\mathcal{M}_1$: This model accounts for nitrogen recombination at the surface by introducing an extra parameter $\gamma_{\mathrm{N}}^{\mathrm{N_2}}$ in the surface mass balance through the nitrogen recombination reaction $\mathrm{N(s) + N \rightarrow N_2}$. Included in the original model of Zhluktov and Abe [24], nitrogen recombination is a competing mechanism to nitridation, consuming available atomic nitrogen which, in turn, has the potential to impact the former reaction in a coupled fashion. This model also considers the flow to be in thermal equilibrium and imposes the surface temperature. The implicit model relationships between inputs and outputs for this model reads

$$\dot{s} = \dot{s}\left(P_\delta, T_\delta, P_\mathrm{d}, T_\mathrm{w}, \gamma_{\mathrm{N}}^{\mathrm{CN}}, \gamma_{\mathrm{N}}^{\mathrm{N_2}}\right), \tag{5}$$

$$\rho_{\mathrm{CN}} = \rho_{\mathrm{CN}}\left(P_\delta, T_\delta, P_\mathrm{d}, T_\mathrm{w}, \gamma_{\mathrm{N}}^{\mathrm{CN}}, \gamma_{\mathrm{N}}^{\mathrm{N_2}}\right). \tag{6}$$

- $\mathcal{M}_2$: In this model the thermal state is considered in equilibrium, surface recombination is taken into account, and the surface temperature is now a prediction of the model through a Surface Energy Balance (SEB)

$$\left[\rho_\mathrm{w}(\boldsymbol{u}_\mathrm{g} - \boldsymbol{u}_\mathrm{r})h_\mathrm{w} + \sum_{i\in\mathcal{S}} \rho_{\mathrm{w},i} h_{\mathrm{w},i} \boldsymbol{V}_{\mathrm{w},i} + \lambda_\mathrm{w}\nabla T_\mathrm{w} + \lambda_\mathrm{w}^{\mathcal{V}}\nabla T_\mathrm{w}^{\mathcal{V}}\right]\cdot\boldsymbol{n} = \dot{\Omega}_{\mathrm{w},i},\ \forall i \in \mathcal{S}. \tag{7}$$

  where the terms $\sum_{i\in\mathcal{S}} \rho_{\mathrm{w},i} h_{\mathrm{w},i} \boldsymbol{V}_{\mathrm{w},i}$, $\lambda_\mathrm{w}\nabla T_\mathrm{w}$ and $\lambda_\mathrm{w}^{\mathcal{V}}\nabla T_\mathrm{w}^{\mathcal{V}}$ are the diffusive and conductive heat fluxes for translational and vibrational thermal baths, respectively. Quantities $\rho_\mathrm{w}, \boldsymbol{u}_\mathrm{g}, \boldsymbol{u}_\mathrm{r}$ and h_w are the gas density at the surface, the velocity of the gas leaving the interface, the velocity at which the interface is moving, and the enthalpy at the surface, respectively. The term $\boldsymbol{V}_{\mathrm{w},i}$ refers to the diffusion velocity of species i at the gas-surface lamina, and $\lambda_\mathrm{w}, \lambda_\mathrm{w}^{\mathcal{V}}$ are the thermal conductivities for translational and internal groups of energies [9]. The direction normal to the surface pointing towards the gas is designated by $\boldsymbol{n}$ while the surface temperature for the vibrational thermal bath is $T_\mathrm{w}^{\mathcal{V}}$. As it is recalled, thermal equilibrium is considered in this model both in the gas phase and gas-surface interface for which the condition $T_\mathrm{w} = T_\mathrm{w}^{\mathcal{V}}$ is further imposed. The only source term taken into account in Eq. (7) as the effect of the surface in the surrounding gas is the radiative flux. This term follows the Stefan-Boltzmann law for a body in thermodynamic equilibrium at temperature T_w

$$\dot{\Omega}_\mathrm{w} = \sigma\epsilon\left(T_\mathrm{w}\right)^4, \tag{8}$$

  where σ is the Stefan-Boltzmann constant and ϵ the emissivity coefficient of the surface.

  In this work, the ablation modeling is restricted to carbon preforms as it constitutes the material of choice for wind tunnel testing for nitridation. For carbon preforms, the binding resine is non-existent, allowing us to drop the modeling of the physics associated to pyrolysis in the energy balances [20, 25]. Consequently, Eq. 7 is the result of simplifications that neglect the presence of pyrolysis gases and approximate the conductive heat flux following steady-state ablation conditions in the solid phase and integrating over the semi-infinite material [26, 27].

  The implicit model relationships between inputs and outputs for this model reads

$$\dot{s} = \dot{s}\left(P_\delta, T_\delta, P_\mathrm{d}, \gamma_{\mathrm{N}}^{\mathrm{CN}}, \gamma_{\mathrm{N}}^{\mathrm{N_2}}\right), \tag{9}$$

$$\rho_{\mathrm{CN}} = \rho_{\mathrm{CN}}\left(P_\delta, T_\delta, P_\mathrm{d}, \gamma_{\mathrm{N}}^{\mathrm{CN}}, \gamma_{\mathrm{N}}^{\mathrm{N_2}}\right), \tag{10}$$

$$T_{\mathrm{w}} = T_{\mathrm{w}}\left(P_\delta, T_\delta, P_{\mathrm{d}}, \gamma_{\mathrm{N}}^{\mathrm{CN}}, \gamma_{\mathrm{N}}^{\mathrm{N_2}}\right). \tag{11}$$

- $\mathcal{M}_3$: This model assumes thermal non-equilibrium by considering two different thermal baths for the translational-rotational and vibrational-electronic-electron temperatures. Furthermore, it is also agnostic about the wall surface energy condition by allowing it to be in non-equilibrium by additionally solving the surface energy balance

$$\left[\rho_{\mathrm{w}}(\boldsymbol{u}_{\mathrm{g}} - \boldsymbol{u}_{\mathrm{r}})h_{\mathrm{w}}^{\mathcal{V}} + \sum_{i\in\mathcal{S}} \rho_{\mathrm{w},i} h_{\mathrm{w},i}^{\mathcal{V}} \boldsymbol{V}_{\mathrm{w},i} + \lambda_{\mathrm{w}}^{\mathcal{V}} \nabla T_{\mathrm{w}}^{\mathcal{V}}\right] \cdot \boldsymbol{n} = \dot{\Omega}_{\mathrm{w}}^{\mathcal{V}\mathrm{T}} + \dot{\Omega}_{\mathrm{w}}^{\mathrm{C}\mathcal{V}}, \tag{12}$$

where the convective and conductive terms in Eq. (7) are now only associated to the internal energies $\mathcal{V}$. The first source term $\dot{\Omega}_{\mathrm{w}}^{\mathcal{V}\mathrm{T}}$ is the relaxation of internally excited particles by collision with the surface (heterogeneous quenching). This term was derived by adapting the volumetric Landau-Teller gas-gas collisions formulation [28] to surface gas-surface collisions. The expression reads

$$\dot{\Omega}_{\mathrm{w}}^{\mathcal{V}\mathrm{T}} = \sum_{i\in\mathcal{S}} \alpha_i \mathcal{F}_i^{\downarrow} m_i \left(h_{\mathrm{w},i}^{\mathcal{V}}\left(T^{\mathcal{V}}\right) - h_{\mathrm{w},i}^{\mathcal{V}}(T)\right), \tag{13}$$

where α_i is a parameter that considers that only the α_i fraction of the collisional events of species i leads to a relaxation. The term $\mathcal{F}_i^{\downarrow}$ represents the impinging flux of species i and m_i is the particle mass of species i. The second source term in Eq. (12) is a chemical source term which takes into account that not all the energy from chemical reactions is accommodated on the surface. It reads

$$\dot{\Omega}_{\mathrm{w}}^{\mathrm{C}\mathcal{V}} = \sum_{i\in\mathcal{S}} (1-\beta_i)\left(h_{\mathrm{w},i}^{\mathcal{V}} - h_{\mathrm{w},i}\right), \tag{14}$$

where the β_i coefficient is an adaptation for a two-temperature model of the energy accommodation term defined by Halpern and Rosner [29], and that a $1-\beta$ quota of the heat flux is transferred into the internal energy modes of the products. More details on the two-temperature surface energy balance here presented can be found in the work of Capriati et al. [30].

Finally, surface recombination reactions are also included. The implicit model relationships between inputs and outputs for this model reads

$$\dot{s} = \dot{s}\left(P_\delta, T_\delta, T_\delta^{\mathcal{V}}, P_{\mathrm{d}}, \gamma_{\mathrm{N}}^{\mathrm{CN}}, \gamma_{\mathrm{N}}^{\mathrm{N_2}}, \alpha, \beta, \epsilon\right), \tag{15}$$

$$\rho_{\mathrm{CN}} = \rho_{\mathrm{CN}}\left(P_\delta, T_\delta, T_\delta^{\mathcal{V}}, P_{\mathrm{d}}, \gamma_{\mathrm{N}}^{\mathrm{CN}}, \gamma_{\mathrm{N}}^{\mathrm{N_2}}, \alpha, \beta, \epsilon\right), \tag{16}$$

$$T_{\mathrm{w}} = T_{\mathrm{w}}\left(P_\delta, T_\delta, T_\delta^{\mathcal{V}}, P_{\mathrm{d}}, \gamma_{\mathrm{N}}^{\mathrm{CN}}, \gamma_{\mathrm{N}}^{\mathrm{N_2}}, \alpha, \beta, \epsilon\right). \tag{17}$$

The main characteristics of each modeling scenario are depicted in the following Table 3

Table 3: Modeling scenarios considered in this work.

Modeling assumptions	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
Thermal non-equilibrium	No	No	No	Yes
Surface energy balance	No	No	Yes	Yes
Surface recombination $\gamma_{\mathrm{N}}^{\mathrm{N_2}}$	No	Yes	Yes	Yes

3. Bayesian model averaging methodology

In this section, we present the methods used to identify whether or not the proposed alternative models can explain the experimental data comparatively well, and how we quantify the resulting uncertainty stemming from such different modeling choices.

3.1. Bayesian inference

Performing Bayesian analyses requires the specification and computation of the different ingredients of the Bayes' rule. We can formulate such rule in a general way for each model considered as

$$\mathcal{P}(\boldsymbol{q}^{(i)}|\boldsymbol{y}^{(i)}_{\text{obs}}, \mathcal{M}_i) = \frac{\mathcal{P}(\boldsymbol{y}^{(i)}_{\text{obs}}|\boldsymbol{q}^{(i)}, \mathcal{M}_i)\, \mathcal{P}(\boldsymbol{q}^{(i)}|\mathcal{M}_i)}{\int_\Omega \mathcal{P}(\boldsymbol{y}^{(i)}_{\text{obs}}|\boldsymbol{q}^{(i)}, \mathcal{M}_i)\, \mathcal{P}(\boldsymbol{q}^{(i)}|\mathcal{M}_i) d\boldsymbol{q}^{(i)}} \quad \forall i \in \{0, ..., 3\}, \tag{18}$$

where $\boldsymbol{q}^{(i)}$ is the vector of parameters to be inferred which results from the stagnation line modeling and physical considerations presented in Sec. 2 for each model, and $\boldsymbol{y}^{(i)}_{\text{obs}}$ is the vector of observations for model i. Furthermore, for compact notation, we split the observations and parameters to be inferred in two distinct subsets each. The subset $\boldsymbol{q}^{(i)'}$ whose parameters have direct measured counterparts, and $\boldsymbol{y}^{(i)'}_{\text{obs}}$ containing the corresponding measurements. The subset $\boldsymbol{q}^{(i)''}$ contains the parameters that appear implicitly in the model outputs dependencies, while $\boldsymbol{y}^{(i)''}_{\text{obs}}$ is the vector containing the measurements that are to be compared against the corresponding model outputs. We also define the set of model outputs $\boldsymbol{Y}^{(i)}$ which contains the model outputs to be compared against the corresponding $\boldsymbol{y}^{(i)''}_{\text{obs}}$ set, such as recession rates, CN densities and wall temperatures. This partition of the parameters, observations and model outputs sets helps setting up the likelihood function with a more compact notation. Table 4 shows an overview of the different parameters, observations and model outputs considered under each model.

Table 4: Vectors of parameters to be calibrated, observations and model outputs for the different modeling scenarios.

Vectors	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
Parameters $\boldsymbol{q}^{(i)}$				
$\boldsymbol{q}^{(i)'}$	$P_\delta, P_\text{d}, T_\text{w}$	$P_\delta, P_\text{d}, T_\text{w}$	P_δ, P_d	P_δ, P_d
$\boldsymbol{q}^{(i)''}$	$T_\delta, \gamma_\text{N}^\text{CN}$	$T_\delta, \gamma_\text{N}^\text{CN}, \gamma_\text{N}^{\text{N}_2}$	$T_\delta, \gamma_\text{N}^\text{CN}, \gamma_\text{N}^{\text{N}_2}$	$T_\delta, T_\delta^\mathcal{V}, \gamma_\text{N}^\text{CN}, \gamma_\text{N}^{\text{N}_2}, \alpha, \beta, \epsilon$
Observations $\boldsymbol{y}^{(i)}_{\text{obs}}$				
$\boldsymbol{y}^{(i)'}_{\text{obs}}$	$P_\text{d}^\text{meas}, P_\text{s}^\text{meas}, T_\text{w}^\text{meas}$	$P_\text{d}^\text{meas}, P_\text{s}^\text{meas}, T_\text{w}^\text{meas}$	$P_\text{d}^\text{meas}, P_\text{s}^\text{meas}$	$P_\text{d}^\text{meas}, P_\text{s}^\text{meas}$
$\boldsymbol{y}^{(i)''}_{\text{obs}}$	$\dot{s}^\text{meas}, \rho_\text{CN}^\text{meas}$	$\dot{s}^\text{meas}, \rho_\text{CN}^\text{meas}$	$\dot{s}^\text{meas}, \rho_\text{CN}^\text{meas}, T_\text{w}^\text{meas}$	$\dot{s}^\text{meas}, \rho_\text{CN}^\text{meas}, T_\text{w}^\text{meas}$
Model outputs $\boldsymbol{Y}^{(i)}$	$\dot{s}, \rho_\text{CN}$	$\dot{s}, \rho_\text{CN}$	$\dot{s}, \rho_\text{CN}, T_\text{w}$	$\dot{s}, \rho_\text{CN}, T_\text{w}$

The definition of the likelihood function encapsulates the forward model given that the parameters to be inferred have an implicit relation with the simulated observations. In Sec. 2, we referenced the models used to simulate the reacting boundary layer encountered in the experiments. Specifically, we can predict recession rates, CN densities and wall temperature (only for model $\mathcal{M}_2$ and $\mathcal{M}_3$) which we can then compare to the observations. Having our different sets of measurements, parameters and outputs defined, for each experimental condition and model we propose a general likelihood of the form

$$\mathcal{P}(\boldsymbol{y}_{\text{obs}}^{(i)}|\boldsymbol{q}^{(i)}) \propto \prod_{j=1}^{N^{(i)'}} \exp\left[-\frac{|y_{\text{obs},j}^{(i)'} - q_j^{(i)'}|^2}{2\sigma^2_{y_{\text{obs},j}^{(i)'}}}\right] \times \prod_{k=1}^{N^{(i)''}} \exp\left[-\frac{|y_{\text{obs},k}^{(i)''} - Y_k^{(i)}|^2}{2\sigma^2_{y_{\text{obs},k}^{(i)''}}}\right], \tag{19}$$

where $q_j^{(i)'}, y_{\text{obs},j}^{(i)'}$ represent the jth elements of the $\boldsymbol{q}^{(i)'}, \boldsymbol{y}_{\text{obs}}^{(i)'}$ sets, while $y_{\text{obs},k}^{(i)''}$ is the kth element of the $\boldsymbol{y}_{\text{obs}}^{(i)''}$ set. The terms $N^{(i)'}$ and $N^{(i)''}$ refer to the number of elements in the prime $\boldsymbol{q}^{(i)'}, \boldsymbol{y}_{\text{obs}}^{(i)'}$ sets, and the observations double prime $\boldsymbol{y}_{\text{obs}}^{(i)''}$ set, respectively. The term $Y_k^{(i)}$ is the kth element of the $\boldsymbol{Y}^{(i)}$ set. The standard deviations of the elements in both sets of measurements are expressed as $\sigma_{y_{\text{obs}}^{(i)'},j}$ and $\sigma_{y_{\text{obs}}^{(i)''},k}$. In practice, given that the forward model evaluations that compose the sets $\boldsymbol{Y}^{(i)}$ are expensive, we approximate them with Gaussian process (GP) surrogates whose L_2-norm thresholds are kept under 1% [9].

Besides the likelihood function, we need to define prior distributions for all the parameters in each $\boldsymbol{q}^{(i)}$. In the available literature for tests in the VKI Plasmatron and stagnation line plasma flows [13–15, 31], we can only establish some bounds a priori. In a maximum entropy set-up [32] with this a priori information, the resulting prior distributions are uniform over the prescribed range. Table 5 shows the chosen prior distributions.

Table 5: Parameters to be calibrated and their priors for the different modeling scenarios.

Inference parameter	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$	Prior
Dynamic pressure P_{d}	Yes	Yes	Yes	Yes	$\mathcal{U}[200, 360]$
Edge translational temperature T_δ	Yes	Yes	Yes	Yes	$\mathcal{U}[9000, 13000]$
Static pressure P_δ	Yes	Yes	Yes	Yes	$\mathcal{U}[1200, 1700]$
Surface temperature T_{w}	Yes	Yes	No	No	$\mathcal{U}[2000, 4000]$
Nitridation efficiency $\gamma_{\text{N}}^{\text{CN}}$	Yes	Yes	Yes	Yes	$\log \mathcal{U}[-4, 0]$
Recombination efficiency $\gamma_{\text{N}}^{\text{N}_2}$	No	Yes	Yes	Yes	$\log \mathcal{U}[-4, 0]$
Emissivity ϵ	No	No	Yes	Yes	$\mathcal{U}[0.5, 1]$
Energy acommodation coefficient β	No	No	No	Yes	$\mathcal{U}[-1, 1]$
Collisional efficiency α	No	No	No	Yes	$\mathcal{U}[0, 1]$

Depending on the algorithms used to solve Eq. (18), we generally do not need to compute the normalization constant in the denominator [33]. Its computation is not necessary to achieve the objective of sampling from the posterior distribution for a given model. When different models are to be compared in their predictive capabilities, the normalization constant becomes central and this is at the core of hypothesis testing studies using Bayesian statistics which we address in the following section.

3.2. Hypothesis testing

Performing a hypothesis testing study entails the evaluation of the posterior probabilities of each competing model (or hypothesis in our context). This comparison allows us to weigh the relative merit of each model. At first thought, we might consider the most suitable model to be the one that best fits the data. Difficulties quickly arise as we aim at comparing models of different complexity. More complex models involving the definition of more parameters can generally be more expressive and fit very well the data. In this case, the results of hypothesis testing studies have the philosophical principle of *Ockham's razor* built in them. This principle embodies the fact that the simplest models that can explain comparatively well the data are the preferred ones. The comparison is driven through the posterior ratio

$$\frac{\mathcal{P}(\mathcal{M}_i|\boldsymbol{y}_{\text{obs}})}{\mathcal{P}(\mathcal{M}_0|\boldsymbol{y}_{\text{obs}})}, \quad \forall i \in \{1, ..., 3\}, \tag{20}$$

for which $\mathcal{M}_0$ is used as reference in the denominator. This notation omits using superscripts (i) to denote each model's sets of parameters and observations for clarity. The comparison is performed at the order of magnitude level. If the posterior of the baseline model is several orders of magnitude below the competing model, we would prefer the

competing model and viceversa. If both posteriors are of similar orders of magnitude, the data do not provide enough evidence to preferentially support either model.

The posterior ratio can be decomposed according to the Bayes' rule

$$\frac{\mathcal{P}(\mathcal{M}_i|\boldsymbol{y}_{\text{obs}})}{\mathcal{P}(\mathcal{M}_0|\boldsymbol{y}_{\text{obs}})} = \frac{\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\mathcal{M}_i)}{\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\mathcal{M}_0)} \times \frac{\mathcal{P}(\mathcal{M}_i)}{\mathcal{P}(\mathcal{M}_0)}, \quad \forall i \in \{1, ..., 3\}, \tag{21}$$

where the posterior ratio is expressed as the product of the ratio of the marginalized likelihoods and priors of the different models. The marginalized likelihood of the data under the assumptions of each model is expressed as follows

$$\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\mathcal{M}_i) = \int_\Omega \mathcal{P}(\boldsymbol{y}_{\text{obs}}|\boldsymbol{q}, \mathcal{M}_i)\mathcal{P}(\boldsymbol{q})\mathrm{d}\boldsymbol{q}, \quad \forall i \in \{0, ..., 3\}, \tag{22}$$

where the term $\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\boldsymbol{q}, \mathcal{M}_i)$ is the likelihood in Eq. (18) for the different models under consideration. $\mathcal{P}(\boldsymbol{q})$ is the prior of the parameters $\boldsymbol{q}$. The marginalized likelihood is also the normalization constant of the posterior $\mathcal{P}(\boldsymbol{q}|\boldsymbol{y}_{\text{obs}}, \mathcal{M}_i)$ under each model. The prior ratio $\mathcal{P}(\mathcal{M}_i)/\mathcal{P}(\mathcal{M}_0)$ is taken as unity in this case. We consider all models to be equally likely to explain the experiments a priori. The results from the hypothesis testing study is entirely based on how well each model fits the data (values of the likelihood $\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\boldsymbol{q}, \mathcal{M}_i)$), and how complex each model is regarding the number and ignorance about each model parameter (dimensionality of the integral in Eq. (22) and overall support of the priors of $\boldsymbol{q}$).

Looking at Table 4, it is straightforward to see that the null hypothesis $\mathcal{M}_0$ is the simplest model in terms of model parameters, while $\mathcal{M}_3$ is the most complex one. For $\mathcal{M}_0$ to be definitely invalidated by the experimental evidence, its predictions must be considerably worse than the ones produced by any other model. We have to consider this assessment with care. We already know that the experimental data can be reasonably well explained with $\mathcal{M}_0$ [9], and the fact that a competing model does much better is reason enough to abandon the hypotheses behind $\mathcal{M}_0$, or at least modify them. If $\mathcal{M}_0$ was not good enough to begin with, we would not have any reason to consider it in the first place for our hypothesis testing study.

Typically, the evidences are compared through the Bayes factor

$$BF_{i0} = \frac{\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\mathcal{M}_i)}{\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\mathcal{M}_0)}, \quad \forall i \in \{1, ..., 3\}, \tag{23}$$

where BF refers to the Bayes factor which is generally expressed in logarithmic scale as

$$\log_{10} BF_{i0} = \log_{10}(\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\mathcal{M}_i)) - \log_{10}(\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\mathcal{M}_0)), \quad \forall i \in \{1, ..., 3\}. \tag{24}$$

For strong evidences in favour or against a certain hypothesis, indicative thresholds are found in the literature by Jeffreys [34] and Kass and Raftery [35]. In general, the Bayes factor should be above a threshold of $|\log_{10} BF| = 8$, which indicates that as long as the evidences for each model remain in within some orders of magnitude, we cannot conclude anything from the data. In this regard, it is straightforward to see that $BF_{10} = 7$ indicates that the data are 7 times more likely under $\mathcal{M}_1$ than under $\mathcal{M}_0$, whereas $BF_{10} = 0.2$ indicates that the observations are 5 times more likely under $\mathcal{M}_0$ than under $\mathcal{M}_1$.

3.3. Model averaging

Given the same prior odds for both models, the Bayes factor is a metric that allows us to quantify the overall best performing model. Apart from Bayes factors, other model selection criteria are commonly used in the literature. One criterion that closely approximates the Bayes factor is the Bayesian Information Criterion (BIC). Developed by Schwarz [36], the BIC serves as an asymptotic approximation to a transformation of the Bayesian posterior probability of a candidate model. The model corresponding to the minimum value of BIC is selected. Similarly, the Akaike Information Criterion (AIC) score rewards models that achieve a high goodness-of-fit score and penalizes them if they become overly complex [37]. The model with the lower AIC score is expected to strike a superior balance. A comprehensive overview of AIC and other popular model selection methods is given by Ding et al. [38].

In this work, we focus on computing Bayes factors as the metric used to discard different thermochemical models and ablation boundary conditions (Sec. 2). An issue that arises is the fact that Bayes factor thresholds are not very

sharp and there can be a grey area where the prior spaces for both models are comparable as well as their goodness-of-fit ratio. There can also be a trade-off between prior space, or complexity, and goodness-of-fit in such a way that the Bayes factor falls in the grey area. This would indicate that the more complex model manages to overcome the complexity penalization by fitting the data much better than the simpler model. In this case, discarding a given model is not possible on the basis of the data and prior information we possess.

Already Chamberlin [39] advocated the concept of "multiple working hypotheses", implying that there is not a single "right" model; instead, there are several well-supported models that are being entertained. If new experiments are well-designed to maximize their informational content, the results tend to support one or more models while providing less support for others. Repetition of this general approach leads to advances, combining experimental design with comparison of theoretical models. New or more elaborate models are consequently postulated, while models with little empirical support are gradually dropped from consideration.

Acknowledging this fact means incorporating model uncertainty stemming from the existence of more than one plausible model. Given these conditions, we may forego selection, retaining the uncertainty across the model space. In this work, we focus on Bayesian model averaging [40]. The model averaged posterior density for the model parameters $\gamma_{\mathrm{N}}^{\mathrm{CN}}$ shared between the baseline and each alternative model reads

$$\mathcal{P}(\gamma_{\mathrm{N}}^{\mathrm{CN}}|\boldsymbol{y}_{\mathrm{obs}}) = \sum_{i\in\{0,\dots,3\}} \mathcal{P}(\gamma_{\mathrm{N}}^{\mathrm{CN}}|\boldsymbol{y}_{\mathrm{obs}}, \mathcal{M}_i)\mathcal{P}(\mathcal{M}_i|\boldsymbol{y}_{\mathrm{obs}}). \tag{25}$$

The result is an overall marginal posterior which gathers the information on nitridation efficiencies provided by each valid model, incorporating this way model-form uncertainties in the calibration of $\gamma_{\mathrm{N}}^{\mathrm{CN}}$.

BMA is important for expressing model uncertainty in a consistent way, thus, avoiding over-confident inferences. Introduced by Roberts [41], who for the first time combined two expert opinions in the statistical literature, BMA has a challenging implementation due to the complex probability distributions it deals with and their integrations. Moreover, assigning prior distributions to each specified model can be a source of caveats when the choice is not obvious.

4. Results and discussion

In this section, we present and discuss the results of the analyses. First, the hypothesis testing results are laid out followed by the resulting model averaging outcomes. Summary statistics are provided for the final averaged models.

4.1. Hypothesis testing

The evidences $\mathcal{P}(\boldsymbol{y}_{\mathrm{obs}}|\mathcal{M}_i)$ are computed for each model and experimental case (see Table 6). As we recall, given the same prior odds to each competing model, the evidences dominate the computation of the Bayes factors for model selection. In the particular case of experiment G6, only recession rates and wall temperatures (when applicable) are used in the computations, producing evidences of several orders of magnitude lower than for the other experimental conditions. This is not an issue given that the comparisons are made among the models for each experimental case independently and on the basis of the same experimental data. The CN density measurement under G6 conditions was deemed untrustworthy in previous analyses [9], therefore, it is not used in the context of this study.

Table 6: Evidences or marginalized likelihoods $\mathcal{P}(\boldsymbol{y}_{\mathrm{obs}}|\mathcal{M}_i)$ of the different models for the given experimental conditions.

Experiment ID	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
G4	1.12×10^{11}	5.6×10^{10}	6.0×10^{7}	6.64×10^{7}
G5	7.31×10^{10}	6.76×10^{10}	2.85×10^{5}	7.73×10^{6}
G6	5.1×10^{4}	5.08×10^{4}	3.45×10^{1}	4.43×10^{1}
G7	3.7×10^{10}	3.76×10^{10}	5.0×10^{5}	3.7×10^{6}

Table 7 shows the resulting Bayes factors. There is no definitive evidence against the baseline model which considers thermal equilibrium and nitridation reactions as the only wall reactions. On the other hand, it means that the SEB model ($\mathcal{M}_2$) does not perform too poorly compared to imposing the measured surface temperature directly,

giving us some confidence in such models for ablative surfaces under the considered Plasmatron conditions. The baseline model $\mathcal{M}_0$ does pretty well compared against the other models for which the larger number of parameters to be inferred penalizes the resulting evidence. In the case of model $\mathcal{M}_2$, the offset of the temperature prediction is what lowers its evidence given that $\mathcal{M}_0$ and $\mathcal{M}_2$ have the same number of parameters $\boldsymbol{q}$. $\mathcal{M}_2$ also replaces T_w, which is a well-defined parameter within its experimental uncertainty for $\mathcal{M}_0$, with the emissivity ϵ which does not have a measured counterpart in these cases, also contributing to the lowering of $\mathcal{M}_2$'s evidence. Model $\mathcal{M}_2$ needs a perfect prediction of T_w to amount to the evidence of $\mathcal{M}_0$ which is not achieved, although the deviation is small to still consider the model good enough compared to the baseline.

Table 7: Bayes factors of the different models against the baseline model $\mathcal{M}_0$ for each experimental condition.

Experiment ID	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
G4	-0.3	-3.27	-3.22
G5	-0.034	-5.4	-4.0
G6	-0.0024	-3.17	-3.06
G7	0.01	-4.87	-4.0

From this study we cannot conclude that the baseline model is the truth and it indicates what is happening in the flow but rather that the experimental data are consistent with the baseline model as well as the other models. If the true flow was in non-equilibrium, other measurements would be needed in order to find the relevant evidence and truly invalidate the baseline model of thermal equilibrium. The same is applicable to the presence of recombination reactions. Concentrations of N_2 in the vicinity of the material surface that cannot be explained by the baseline model could definitely give evidence of recombination reactions and their relative importance. So far, the calibrations of the different models with this particular set of experimental data do not provide enough information about recombination reactions. Fig. 2 depicts the marginal posteriors for $\gamma_\mathrm{N}^{\mathrm{N}_2}$ under the different models that include such parameter. Almost flat posteriors with the same support as the priors are recovered except for the 1T SEB model $\mathcal{M}_2$ where the posterior support is reduced. In order for $\mathcal{M}_2$ to be able to explain the experimental observations, recombination reaction efficiencies cannot surpass a certain upper limit. Going over such limit would probably throw off the prediction of the surface temperature. Given that the other models have more parameters to control the surface temperature prediction, model $\mathcal{M}_2$ is able to place more constraints on $\gamma_\mathrm{N}^{\mathrm{N}_2}$. It can also be appreciated that as the wall temperature measured decreases, the posterior support for $\gamma_\mathrm{N}^{\mathrm{N}_2}$ shrinks toward lower values.

For more insight about model performance, we can look at the resulting residuals of the statistical inverse problem for the measured quantities. The inverse problem we are solving every time we perform an inference amounts to solving the equation

$$\boldsymbol{y}_\mathrm{obs} = \boldsymbol{Y}(\boldsymbol{q}) + \boldsymbol{E}, \tag{26}$$

where the stagnation line model evaluations are replaced by accurate GP surrogates $\boldsymbol{Y}$ for the predictions of recession rates $\dot{s}$, CN densities ρ_CN, and surface temperatures T_w. The error model is chosen as unbiased Gaussian distributions for which we prescribe the different standard deviations σ for each measurement. Once the model is calibrated and we have posteriors on the parameters $\boldsymbol{q}$, we propagate them through the model to generate the resulting distribution of $\boldsymbol{y}_\mathrm{obs} - \boldsymbol{Y}(\boldsymbol{q})$. This distribution is what we call the residual and it should match the error model distribution $\boldsymbol{E}$ initially prescribed for each measurement. In this case, we compute the residual as the distribution of $\boldsymbol{Y}(\boldsymbol{q}) - \boldsymbol{y}_\mathrm{obs}$ as it has a more intuitive interpretation when looking at the results. This distribution also has to amount to the error distribution given its symmetrical qualities.

Fig. 3 shows the residuals for case G5 for recession rate $\dot{s}^\mathrm{meas}$, CN density $\rho_\mathrm{CN}^\mathrm{meas}$ and surface temperature $T_\mathrm{w}^\mathrm{meas}$ under the different modeling scenarios. The rest of the results for the other experimental cases can be found in Appendix A. Overall, all models retrieve the noise of recession rate and CN density. Model $\mathcal{M}_2$ shows the largest deviations in the predictions which are still quite small. The 1T SEB model ($\mathcal{M}_2$) tends to over predict recession rates and under predict CN densities in the flowfield. The surface temperature is perfectly captured by the 2T SEB model $\mathcal{M}_3$ while the 1T SEB model $\mathcal{M}_2$ shows some deviation towards higher surface temperatures. The deviation is ~ 25 K which is rather small compared to the range of surface temperatures we deal with in this problem ($> 2,000$ K).

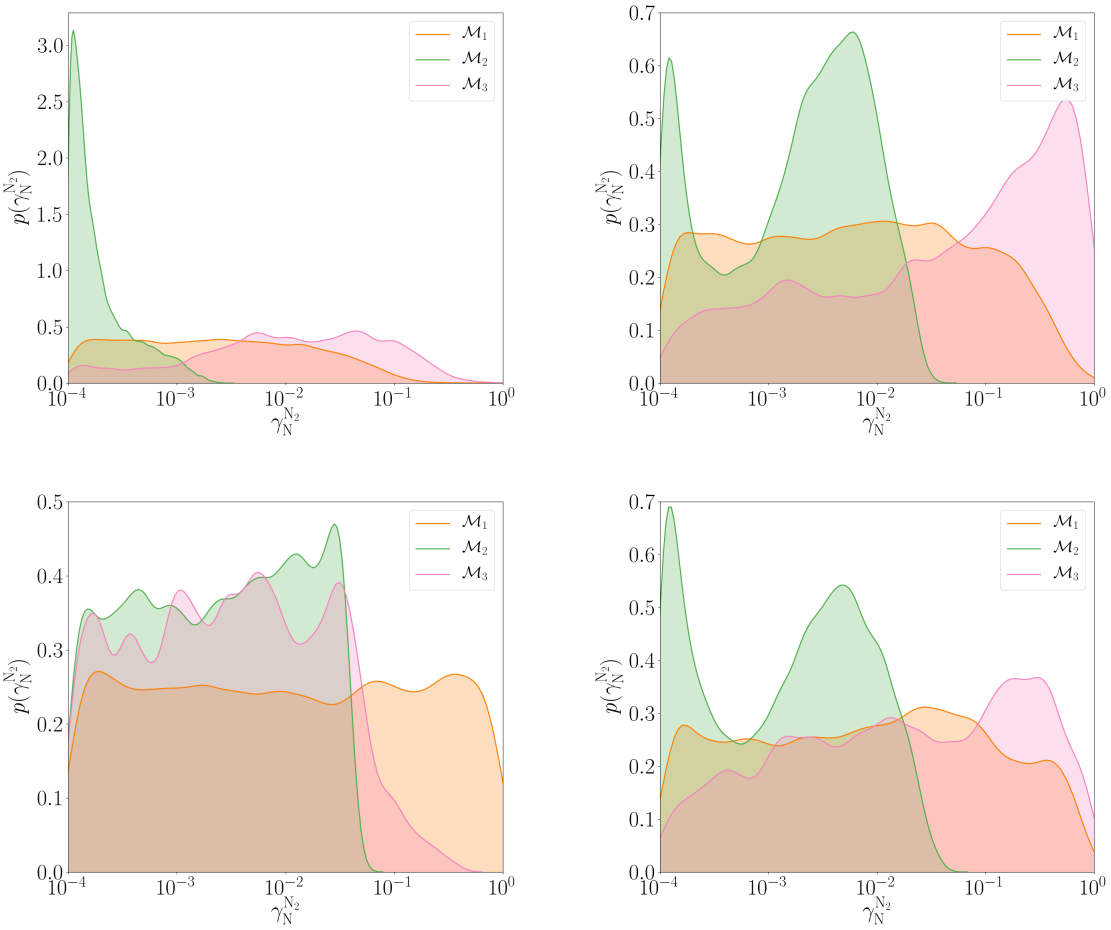

Figure 2: Posterior marginal distributions of $\gamma_N^{N_2}$ under the different modeling scenarios for G4 (upper left), G5 (upper right), G6 (lower left) and G7 (lower right).

It is interesting to reflect on how the 2T SEB model $\mathcal{M}_3$ manages to achieve a lower prediction of the recession rate and surface temperature, and yet perfectly recuperate the error of the CN density. The only possibility for this to happen is through available thermal non-equilibrium mechanisms which are not considered in the 1T SEB model $\mathcal{M}_2$. The 2T surface energy balance (Sec. 2) contains an additional term for the conductive contribution of the internal energy to the total energy balance. Through this mechanism, the 2T SEB model finds the combination of collisional efficiency α and energy accommodation factor β so that the internal energy contribution cools off the surface. Fig. 4 shows the resulting marginal posterior distributions as well as samples from the joint distribution of the α and β parameters. The collisional efficiency α is quite constrained by the experimental data to low values, while the energy accommodation coefficient β keeps all the prior support. The resulting distributions are shifted towards values that favour thermal non-equilibrium ($\alpha < 1, \beta < 1$).

Fig. 5 shows the propagated posterior solutions of the stagnation line model for both the 1T SEB and 2T SEB models with their corresponding confidence intervals. Promoting strong thermal non-equilibrium, as seen in the temperature profiles in Fig. 5, makes the surface excite the internal energy modes, consequently lowering the resulting translational surface temperature to match perfectly the experimental data. This feature of the 2T SEB model can easily decouple the surface temperature prediction from the rest of the temperature profile in the flowfield. The result is an overall increment of the flowfield temperatures with respect to the 1T SEB model.

The fact that this change in the temperature profiles can also alter the chemical composition in the flowfield to

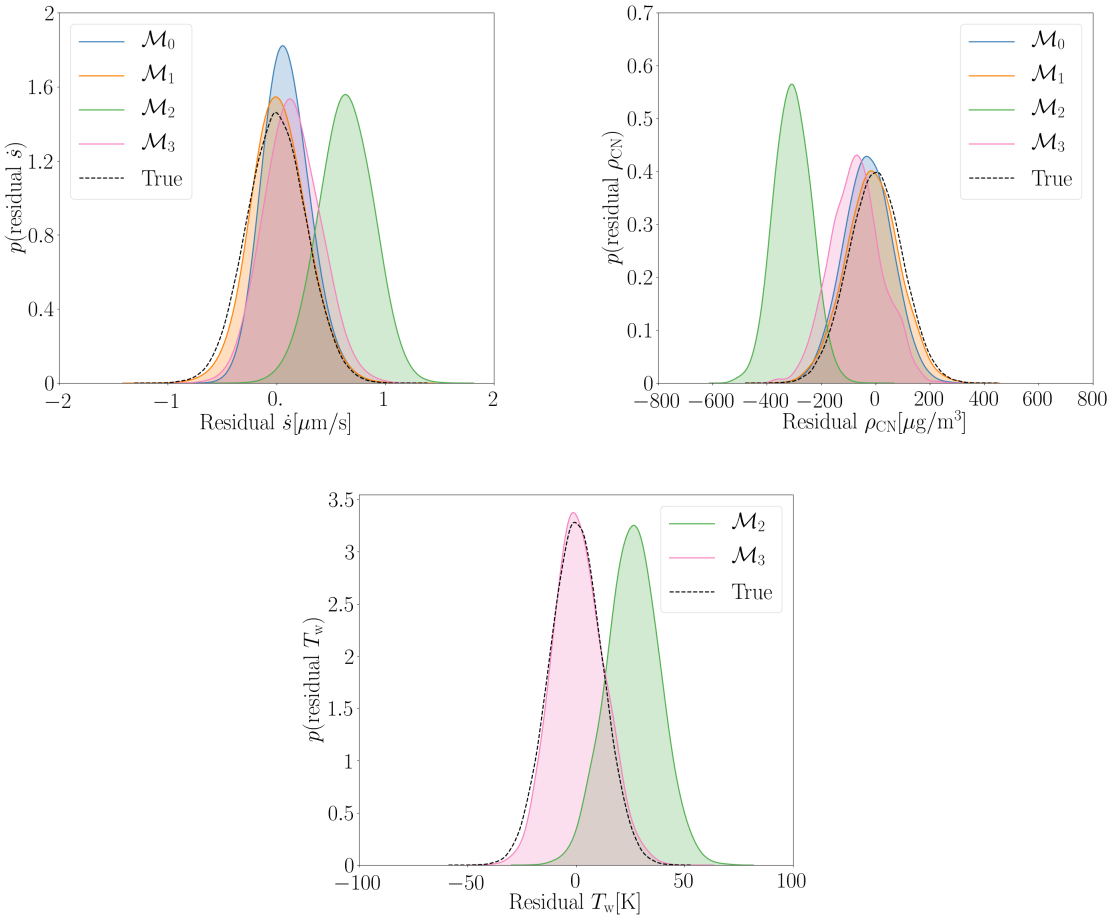

Figure 3: Residuals of recession rate $\dot{s}^{\text{meas}}$, CN density $\rho_{\text{CN}}^{\text{meas}}$ and surface temperature $T_{\text{w}}^{\text{meas}}$ for the different models under conditions G5.

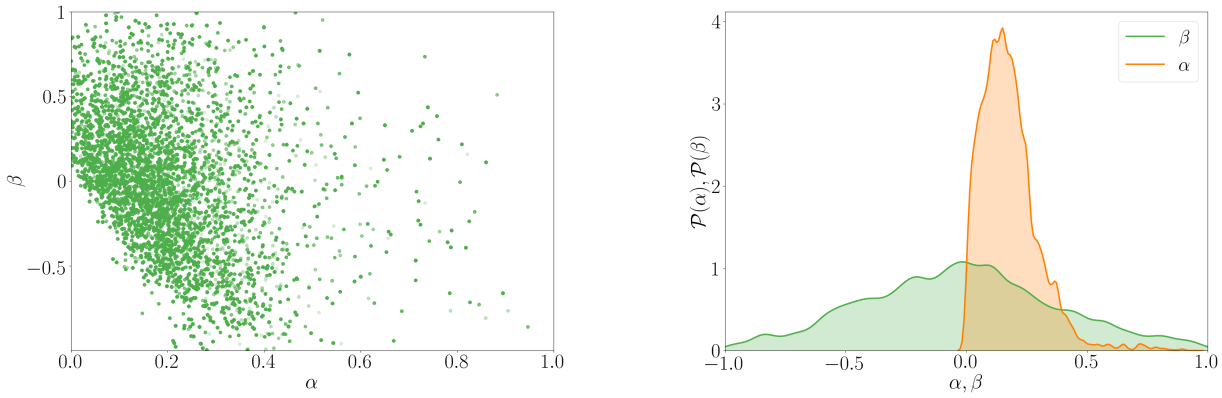

Figure 4: Left: 100,000 samples from the joint posterior distribution of α and β. Right: posterior marginal distributions of α and β under conditions G5.

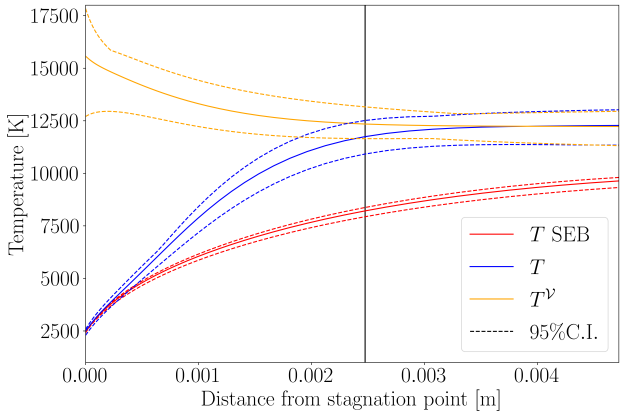

Figure 5: Propagated temperature profiles for the 1T SEB ($\mathcal{M}_2$) and 2T SEB ($\mathcal{M}_3$) models with their respective confidence intervals. The vertical black line represents the position of the spectrometer measurement.

match the observed ρ_{CN} (Fig. 6) brings forth the importance of another set of assumptions: the gas phase chemistry. In particular, the resulting overall depleition and creation of CN differs between the two models in such a way that a higher CN density is achieved at the spectrometer location for the same concentration of CN at the wall, dictated primarily by the recession rates predicted (schematic representation can be found in Fig. 7). Additional experiments for the proper calibration of the gas phase chemistry would be needed to have a realistic assessment on the evidence of thermal non-equilibrium at the wall.

Still, considering the same chemical mechanism and reaction rates, we could also propose new experimental avenues to find strong evidence, or lack thereof, of thermal non-equilibrium at the wall. In particular, the differences in the CN mole fraction profiles could be exploited further by considering the addition of a spectrometer measurement in a different location along the stagnation line. Such measurement could be performed downstream the current spectrometer measurement, although the accuracy and uncertainty estimation of the resulting CN concentration should be carefully addressed to be able to tell the two models apart. Upstream, the radiative signal would lose intensity and the observation would become more noisy. In any case, it represents a challenge to find stronger evidence in favour of thermal equilibrium or non-equilibrium at the wall.

4.2. Bayesian model averaging

All the modeling scenarios presented are so far consistent with the experimental data. We want to incorporate the uncertainty of these multiple modeling choices into the resulting calibrated nitridation reaction efficiencies $\gamma_{\mathrm{N}}^{\mathrm{CN}}$. We recall the model assumptions for which we remain uncertain. The baseline model $\mathcal{M}_0$ considers a flow in thermal equilibrium for which only one temperature is defined. As heterogeneous chemistry mechanism at the wall it only considers nitridation. Model $\mathcal{M}_1$ introduces the possibility of having nitrogen recombination reactions at the surface. Models $\mathcal{M}_2$ and $\mathcal{M}_3$ incorporate a surface energy balance model for the prediction of surface temperatures which are not imposed in these cases. Further, model $\mathcal{M}_3$ assumes the flow to be in a state of thermal non-equilibrium by assuming two distinct thermal baths.

Fig. 8 shows the resulting $\gamma_{\mathrm{N}}^{\mathrm{CN}}$ marginal posteriors under the different modeling scenarios. The resulting weighted average posterior is labelled as BMA. Overall, all the posteriors share supports and are consistent with each other. This is to be expected as $\gamma_{\mathrm{N}}^{\mathrm{CN}}$ dominates recession rate and surface temperature measurements which are quite accurate in all cases and are not substantially affected by the differential features of the models. We notice a broader support for the thermal equilibrium model with recombination $\gamma_{\mathrm{N}}^{\mathrm{N_2}}$ in it $\mathcal{M}_1$. As a competing mechanism, recombination reactions allow nitridation reactions to have higher efficiencies for the same observed recession rates. The atomic nitrogen available at the surface is more scarce for nitridation reactions when competing with recombination reactions. Higher nitridation efficiencies would still produce the same observed recession rates. The wider support is also extended to the averaged posterior given that model $\mathcal{M}_1$ has similar weight as all the other models, consequently widening the support for $\gamma_{\mathrm{N}}^{\mathrm{CN}}$. The models with the tightest supports are the baseline $\mathcal{M}_0$ and the 1T SEB $\mathcal{M}_2$ models. They are

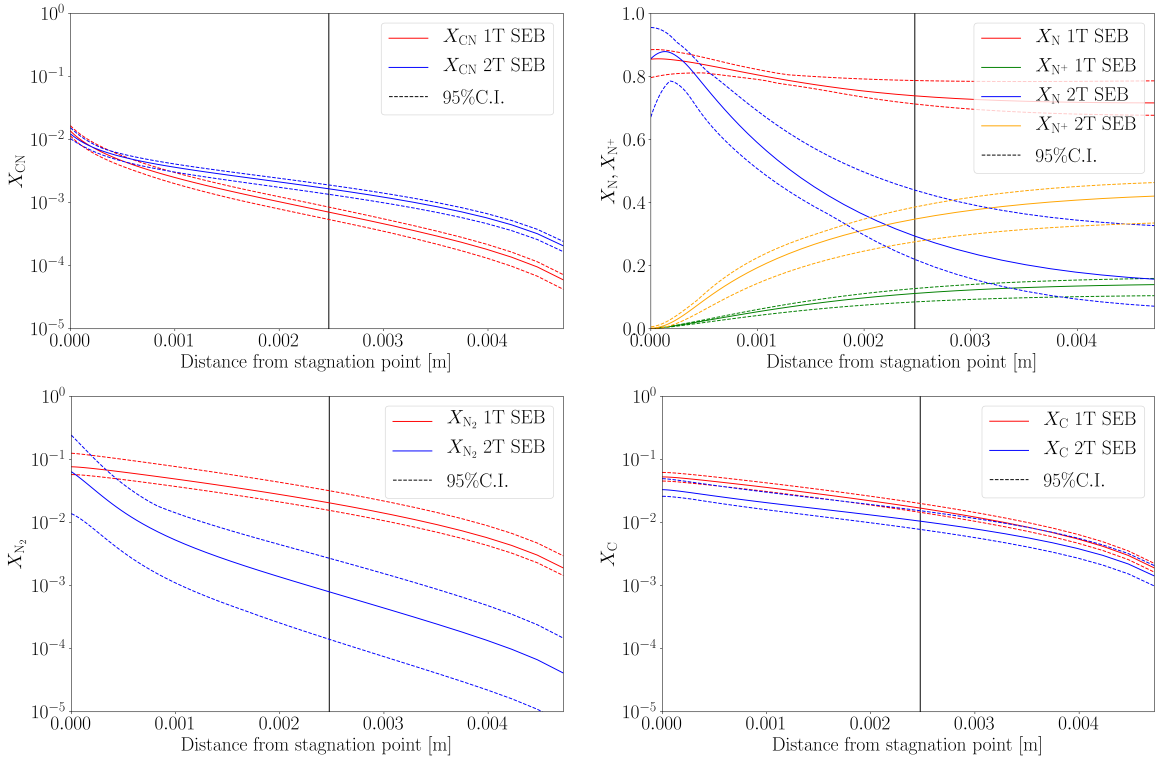

Figure 6: Propagated mole fraction profiles for the 1T SEB ($\mathcal{M}_2$) and 2T SEB ($\mathcal{M}_3$) models with their respective confidence intervals. The vertical black lines represent the position of the spectrometer measurement.

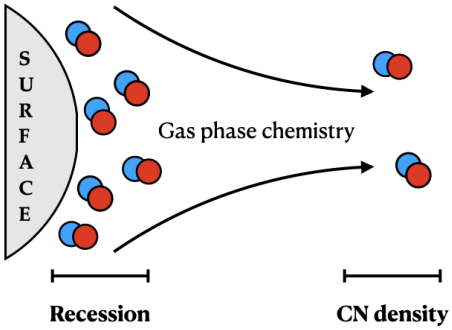

Figure 7: Schematic representation of the modulating effect of the gas phase chemistry in the calibration of the different stagnation line models. The CN concentration at the wall is dominated by the recession rate while the CN density at the spectrometer location is a combined result of the recession rate and the gas phase chemistry.

also the models with the least number of parameters to be inferred, producing more defined peaks and less uncertainty for γ_N^{CN}.

The most relevant summary statistics for the resulting BMA distributions are gathered in Table 8.

We also extend the model averaging methodology to the computation of an Arrhenius law that captures such model-form uncertainties for a broad span of surface temperatures. In this context, the calibration of an Arrhenius law is performed for each of the models proposed. It takes the form

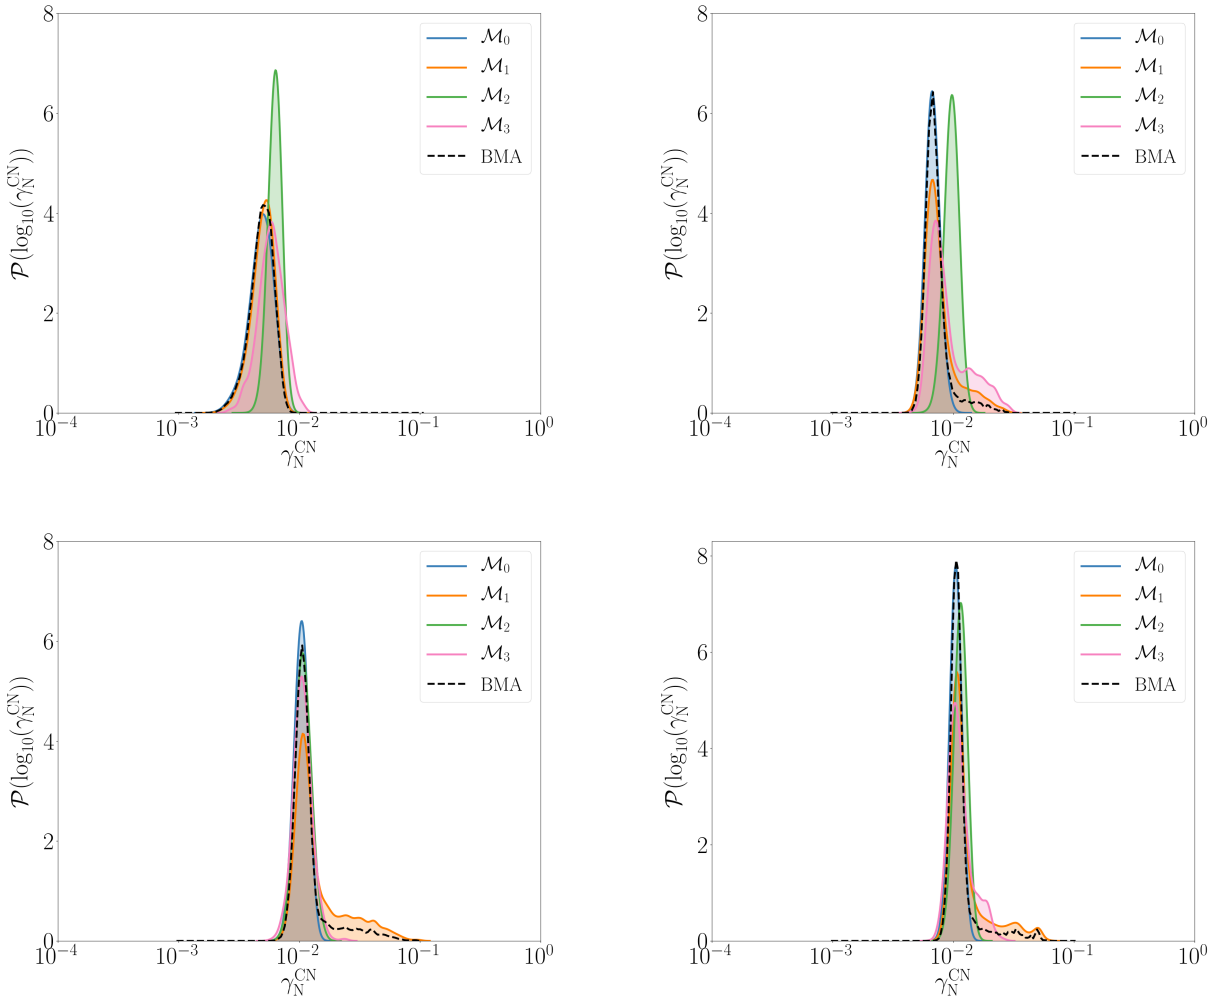

Figure 8: Posterior marginal distributions of γ_N^{CN} under the different modeling scenarios and the resulting averaged posteriors for G4 (upper left), G5 (upper right), G6 (lower left) and G7 (lower right).

Table 8: γ_N^{CN} posterior statistics of the resulting BMA distributions for the different experimental cases.

Experiment ID	Mean	95% C. I.	MAP
G4	0.005	[0.003, 0.007]	0.005
G5	0.007	[0.005, 0.013]	0.007
G6	0.012	[0.008, 0.026]	0.010
G7	0.012	[0.009, 0.024]	0.010

$$\gamma_N^{CN} = A \exp\left(\frac{-T_a}{T_w}\right), \tag{27}$$

with the pre-exponential factor A, and the activation energy (given in temperature units) T_a. For this analysis, both parameters are the same accross all experimental cases and considered model-dependent. The resulting marginal posteriors are then averaged a posteriori to compute the Bayesian model averaged Arrhenius law.

The inference problem needs to be re-formulated to account for the new parametrization introduced. For ease of notation, we unpack here a general likelihood function applicable to each model proposed in Sec. 2 by changing the $\boldsymbol{q}$, $\boldsymbol{y}_{\text{obs}}$ and $\boldsymbol{Y}$ sets to their specific ones depicted in Table 4. Along with the Arrhenius parameters, we need to add all the nuisance parameters to the calibrations. In this case, each nuisance parameter is represented as a different variable for each experimental case. The compact notation is as follows. For each case (i) in $\mathbf{G} = (\mathrm{G}_4, \mathrm{G}_5, \mathrm{G}_6, \mathrm{G}_7)$, we define the following likelihood

$$\mathcal{P}(\boldsymbol{y}_{\text{obs}}|\boldsymbol{q}) \doteq \prod_{i \in \mathbf{G}} \mathcal{P}(\boldsymbol{y}_{\text{obs}}^{(i)}|\boldsymbol{q}^{(i)}), \tag{28}$$

with $\mathcal{P}(\boldsymbol{y}_{\text{obs}}^{(i)}|\boldsymbol{q}^{(i)})$ defined for each model following Eq. (19). Fig. 9 depicts a schematic representation of the methodology. All experimental cases are used to jointly inform the Arrhenius law parameters for each model scenario. Model averaging is performed on such parameters to retrieve the final parameters which incorporate model-form uncertainty.

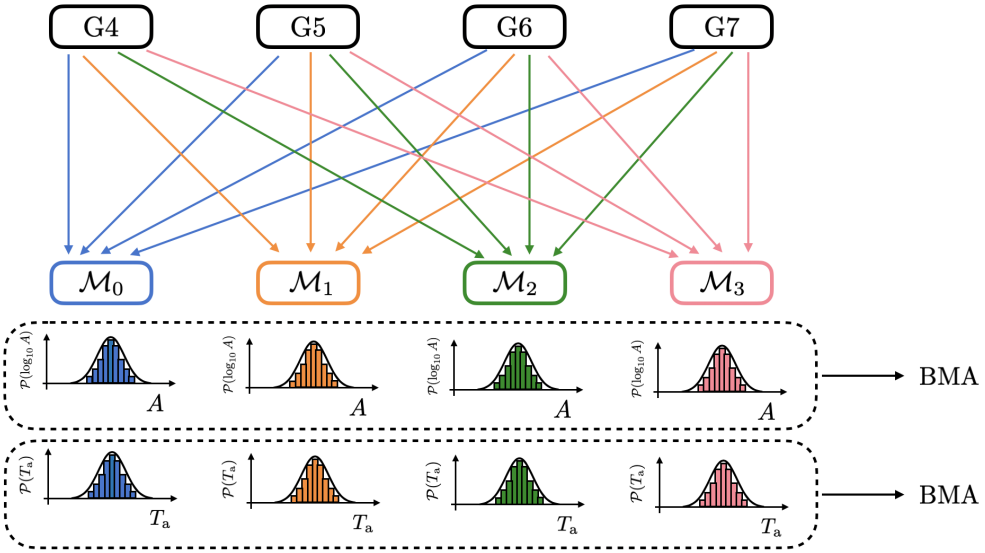

Figure 9: Schematic representation of the relation among the experimental data and model scenarios for the inference of the Arrhenius law parameters.

The prior distribution is taken as independent priors on all parameters, already specified in Sec. 3.1. For A we prescribe the same prior as for $\gamma_{\text{N}}^{\text{CN}}$ as it is an equivalent parameter, in this case, $A \sim \log \mathcal{U}[-4, 0]$. The prior on T_{a} is chosen wide enough with the same lower bound as the prior on T_{w}, given that we expect the activation energy to be at least as low as the a priori possible lowest T_{w}. It is not easy to prescribe a good upper bound as there is not so much literature on this particular reaction parameter. We choose the upper bound conservatively large based on the results of Suzuki et al. [23] and Helber et al. [10], having $T_{\text{a}} \sim \mathcal{U}[2000, 10000]$.

The results of the Arrhenius parameter marginal posteriors can be seen in Fig. 10. Overall there is good agreement in their support and the resulting BMA distributions follow closely the posterior obtained with the baseline model $\mathcal{M}_0$. The summary statistics are gathered in Table 9.

Table 9: Posterior statistics of the resulting BMA distributions for the Arrhenius parameters.

Arrhenius parameter	Mean	95% C. I.	MAP
A	0.253	[0.065, 0.502]	0.42
T_{a} [K]	8328.7	[4996.1, 9944.2]	9669

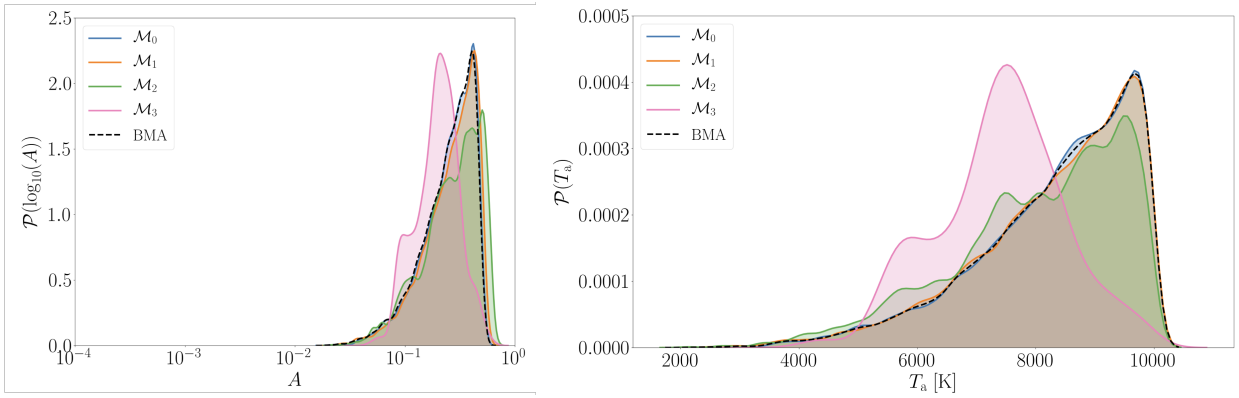

Figure 10: Calibrated Arrhenius parameters and their weighted average for the different models considered.

Propagating the pre-exponential factor and activation energy posteriors we get an estimation of the variation of the nitridation parameter γ_N^{CN} with the surface temperature T_w. Fig. 11 shows the different Arrhenius laws and the BMA with the 95% confidence intervals. The right plot of Fig. 11 shows the obtained calibrated laws plotted against the available experimental data. Even though the resulting laws are calibrated with the data from Helber et al. [10], we can project accurate confidence intervals to lower surface temperatures, where the experimental data available differ by several orders of magnitude. The data of Helber et al. are consistent with the experiments from Zhang et al. [6] for lower surface temperatures as well as Suzuki et al.'s [23] data. This apparent agreement enhances our confidence in the tools developed and the analyses performed.

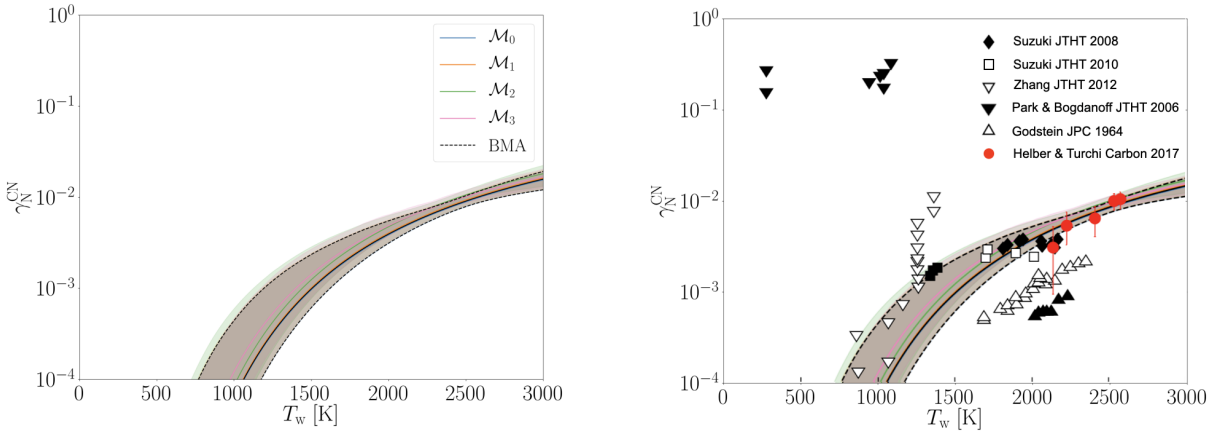

Figure 11: Calibrated Arrhennius laws and their weighted average for the different models considered (left). The Arrhenius laws are compared against the available experimental data (right).

A question that can be posed on these results is if defining a model averaged brings anything compared to just using the estimation recovered with the baseline model alone, given that the differences seem to be negligible. Indeed, we can already see from the case by case inferences that the SEB models have several orders of magnitude less evidence weight and this causes the BMA to basically be an average between $\mathcal{M}_0$ and $\mathcal{M}_1$ which very much agree in their Arrhenius parameters. One could argue that we could have saved ourselves some work by just extrapolating this reasoning from the case by case analyses to the Arrhenius parameters but the answer is not that simple. Arrhenius parameters introduce a new parametrization, one where there is coupling between the prediction of the wall temperature for the SEB models and the Arrhenius parameters themselves, leading to a non-linear constraint on the predicted γ_N^{CN} ($\log(\gamma_N^{CN}/A) + T_a/T_w = 0$, with $T_w = T_w(..., \gamma_N^{CN})$) for every pair of Arrhenius parameters. This issue, together with the fact that we infer using all data jointly, further influences the inferences on the parameters we are after. All in

all, it is hard to say a priori what this picture of the different models will look like. It is not until the exercise is done that we come to realise the similarities even under such constraints. Given the logarithmic nature of the scale of the nitridation parameter, averaging can be useful even if small deviations from model to model are found, considering that a small change in scale can mean a sizeable change in $\gamma_{\mathrm{N}}^{\mathrm{CN}}$ and, therefore, the recession rate to be predicted using this calibrated law. To illustrate this further, we can compute the recession rates resulting from using the parameters calibrated assuming the baseline model and the BMA model, respectively. The resulting predictions of recession rates in logarithmic scale for the same set of inputs except the Arrhenius parameters are

$$\log_{10} \dot{s}^{\mathrm{baseline}} = -8.16839, \tag{29}$$

$$\log_{10} \dot{s}^{\mathrm{BMA}} = -8.24106, \tag{30}$$

which when transformed to natural scale, we get a relative difference of $\sim 17\%$ in recession rates, that is one standard deviation of the experimental measurements. These differences can certainly impact the predictive capabilities of the model.

Another important question that arises regarding the results of Fig. 10 is the fact that the marginal distributions, especially those of T_{a}, seem to be constrained by their chosen priors. As said, the choice of priors for the Arrhenius parameters is challenging, given the poor literature on the subject. Even though the supports were chosen conservatively large, the posteriors seem to push their supports towards the higher end of the possible values for T_{a} (Fig. 10). It is interesting to look at the joint marginal posterior of both A and T_{a} given that they are strongly correlated. Fig. 12 shows a scatter plot of some samples taken from the joint marginal posterior. The current prior placed on T_{a} leads to a wall where the samples suddenly stop. Conversely, a wider prior relaxes the situation and now the limit of the posterior support is dictated by the prior placed on A. An Arrhenius law is no other thing than a macroscopic model with physically motivated parameters [42] that describe the frequency of collisions and the portion of those with enough energy to actually react based on the statistics of such encounters. As such, we find that the model becomes degenerate as it describes the same chemistry for large frequencies of collisions and large activation energies and viceversa. This issue is seen in Fig. 12 as the joint posterior goes on and on towards larger values and it only stops when specific constraints are placed, such as the ones stemming from the priors. The initially superficial question becomes more intricate, and it remarks the importance of looking at the microscopic nature of the flow to prescribe physically consistent priors that can tackle this kind of problems, therefore removing the possible bias that can be induced by selecting arbitrary priors.

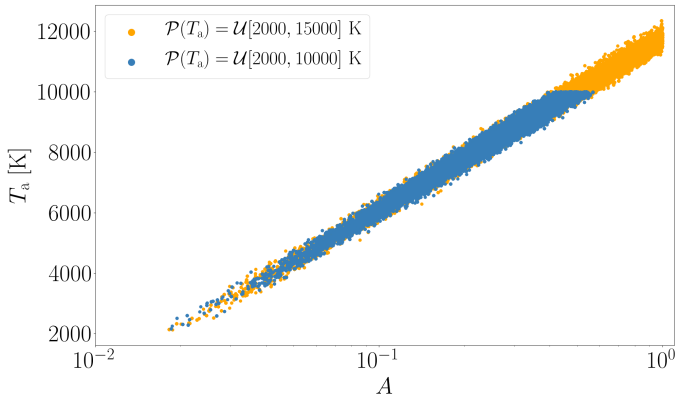

Figure 12: 10,000 samples from the joint marginal posterior of A and T_{a} for two different choices of priors for T_{a}.

5. Concluding remarks and outlook

This work presents a calibrated carbon nitridation model that takes into account experimental, parametric and model-form uncertainties. Experimental and parametric uncertainties are incorporated through dedicated Bayesian inference frameworks. Model-form uncertainties are quantified by performing model selection and averaging on the Bayesian posteriors. The final result is a calibrated Arrhenius law for nitridation reaction efficiencies spanning a range of surface temperatures whose parameters also include uncertainties stemming from different modeling assumptions about the physico-chemical nature of the plasma flow and gas-surface interface.

Different modeling scenarios are proposed in this work. We first define a baseline model against which to compare a set of alternative models in their predictive performances and complexities. The baseline model accounts for flow thermal equilibrium assumptions with nitridation as only surface mechanism; it also imposes the surface temperature, treating it as a free parameter. The model is calibrated and its predictive performance tested on the experimental dataset. The modeling assumptions hold for the conditions of the experiments giving us a solid base on which to build our subsequent hypothesis testing study. The alternative models relax several assumptions embedded in the baseline model. The first alternative model $\mathcal{M}_1$ adds nitrogen recombination reactions through an efficiency parameter that must be also calibrated. Model $\mathcal{M}_2$ builds on model $\mathcal{M}_1$ and includes a surface energy balance to predict the surface temperature instead of treating it as a free parameter to be calibrated. Finally, alternative model $\mathcal{M}_3$ considers nitrogen recombination reactions at the gas-surface interface and remains agnostic about the thermal state of the gas and gas-surface interface for which additional balance equations must be solved.

All these modeling choices are assessed in a hypothesis testing study. Overall no strong evidence is found against the baseline model and none of the considered models can be discarded. The calibrations of models with surface recombination can define an upper limit for the recombination efficiency parameter at best. The calibrated solutions of the thermal non-equilibrium model point to strong thermal non-equilibrium at the gas-surface interface, although the evidence is weak.

Model averaging is then performed to incorporate such model-form uncertainties in a Bayesian setting. The marginal posteriors of nitridation efficiencies parameters for each experimental condition are weighted according to their marginalized likelihoods assuming equal prior probabilities for all models. The results show how for cases G5-G7 the tail of the nitridation reaction efficiencies distributions towards higher values expands, enlarging the space of possible values for $\gamma_{\mathrm{N}}^{\mathrm{CN}}$. This growth of the $\gamma_{\mathrm{N}}^{\mathrm{CN}}$ uncertainty is mainly promoted by model $\mathcal{M}_1$ which includes recombination reactions while freely imposing the surface temperature. Furthermore, an averaged Arrhenius law is also obtained by weighting each Arrhenius model parameter by their corresponding marginalized likelihoods for each individual model. This work offers complete and accurate characterization of the uncertainties affecting the calibration of a carbon nitridation model, updating the current state-of-the-art databases.

The outcomes of this work direct our research efforts in this field towards two different aspects: generating new experimental data with which to be able to rule out some models, finding definitive evidence of thermal non-equilibrium as well as learning surface recombination; and incorporating priors for kinetic models, such as Arrhenius laws, that consistently reflect the microscopic nature of the flow without necessarily incurring into model degeneracies entirely controlled by the prior distributions chosen.

Acknowledgements

This work is fully funded by the European Commission H2020 programme, through the UTOPIAE Marie Curie Innovative Training Network, H2020-MSCA-ITN-2016, Grant Agreement number 722734. The authors would also like to thank Michele Capriati for his guidance in understanding and simulating the 2T SEB model. Prof. Serge Prudhomme is acknowledged for providing insightful comments which improved this work.

Appendix A. Complementary results

This appendix contains the results concerning the cases not included in Section 4.

Appendix A.1. Prediction residuals

Figs. A.13-A.15 show the residuals for cases G4, G6 and G7 for recession rate $\dot{s}$, CN density ρ_{CN} and surface temperature T_{w} under the different modeling scenarios. We recall that we only take into account the recession rate measurements for case G6 for which the CN density residual is not shown.

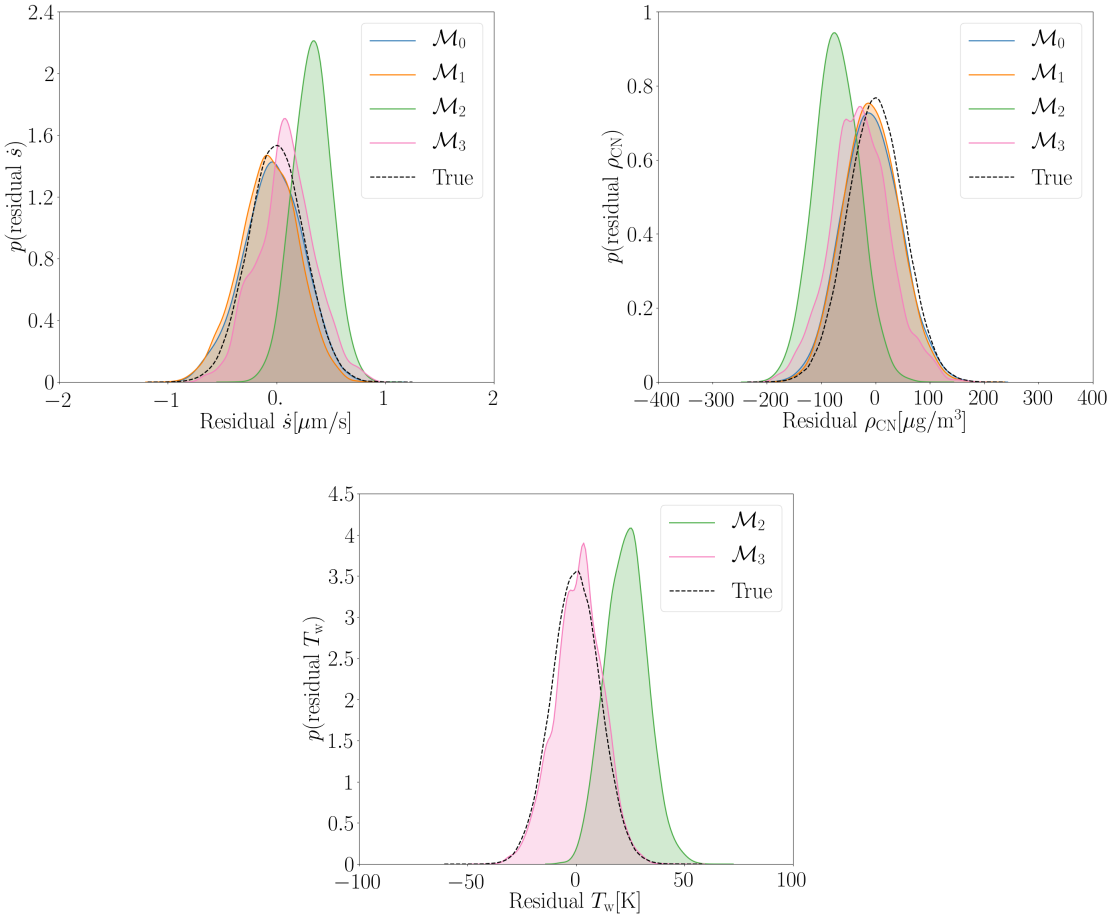

Figure A.13: Residuals of recession rate $\dot{s}$, CN density ρ_{CN} and surface temperature T_{w} for the different models under conditions G4.

Appendix A.2. Calibrated flowfield solutions

Figs. A.16-A.21 show the propagated posterior solutions of the stagnation line model for both the 1T SEB and 2T SEB models with their corresponding confidence intervals for cases G4, G6 and G7.

Appendix A.3. Posterior distributions of the surface parameters for the 2T model

Fig. A.22 shows samples of the 2D joint posterior distributions of the parameters α and β in the 2T surface balance model, as well as their marginal posteriors for cases G4, G6, and G7.

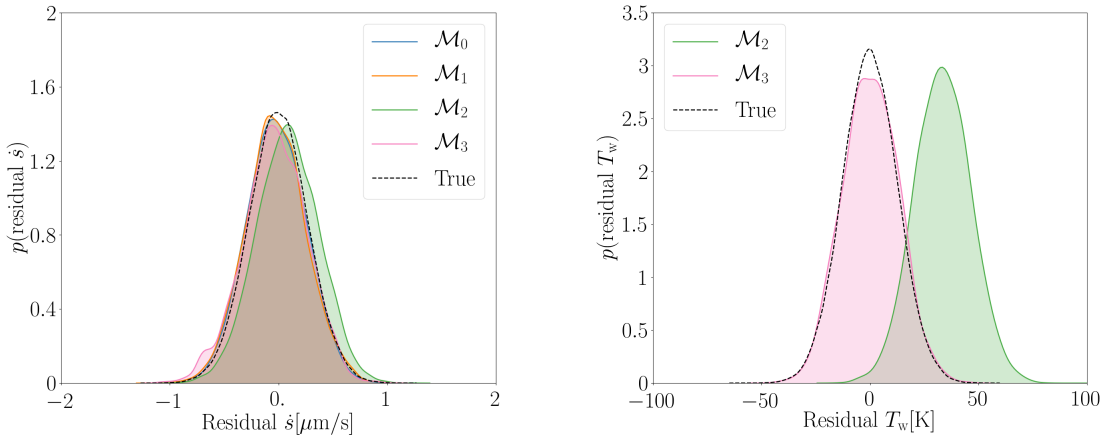

Figure A.14: Residuals of recession rate $\dot{s}$ and surface temperature T_w for the different models under conditions G6.

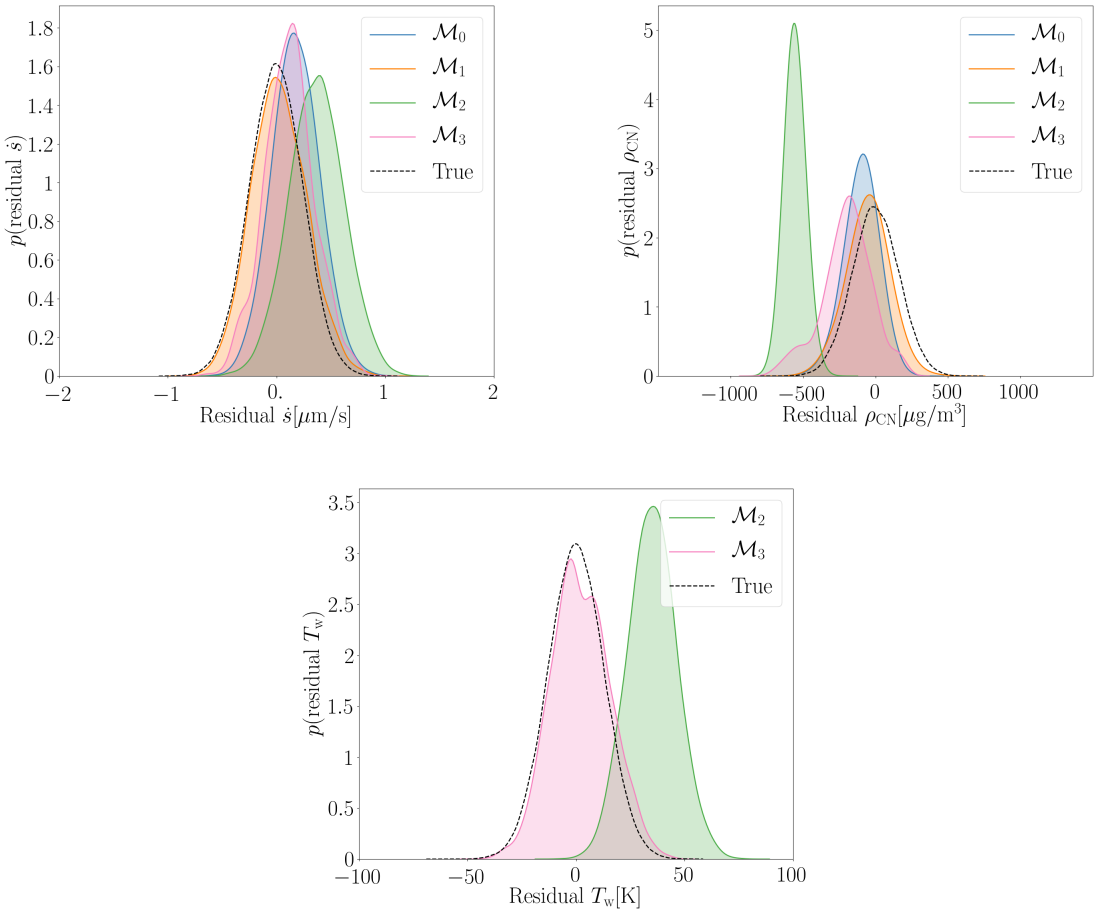

Figure A.15: Residuals of recession rate $\dot{s}$, CN density ρ_{CN} and surface temperature T_w for the different models under conditions G7.

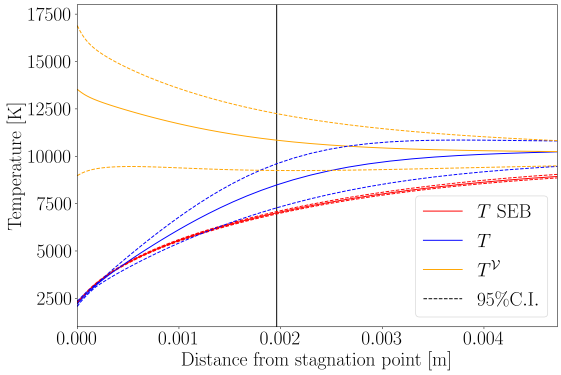

Figure A.16: Propagated temperature profiles for the 1T SEB ($\mathcal{M}_2$) and 2T SEB ($\mathcal{M}_3$) models with their respective confidence intervals for case G4. The vertical black line represents the position of the spectrometer measurement.

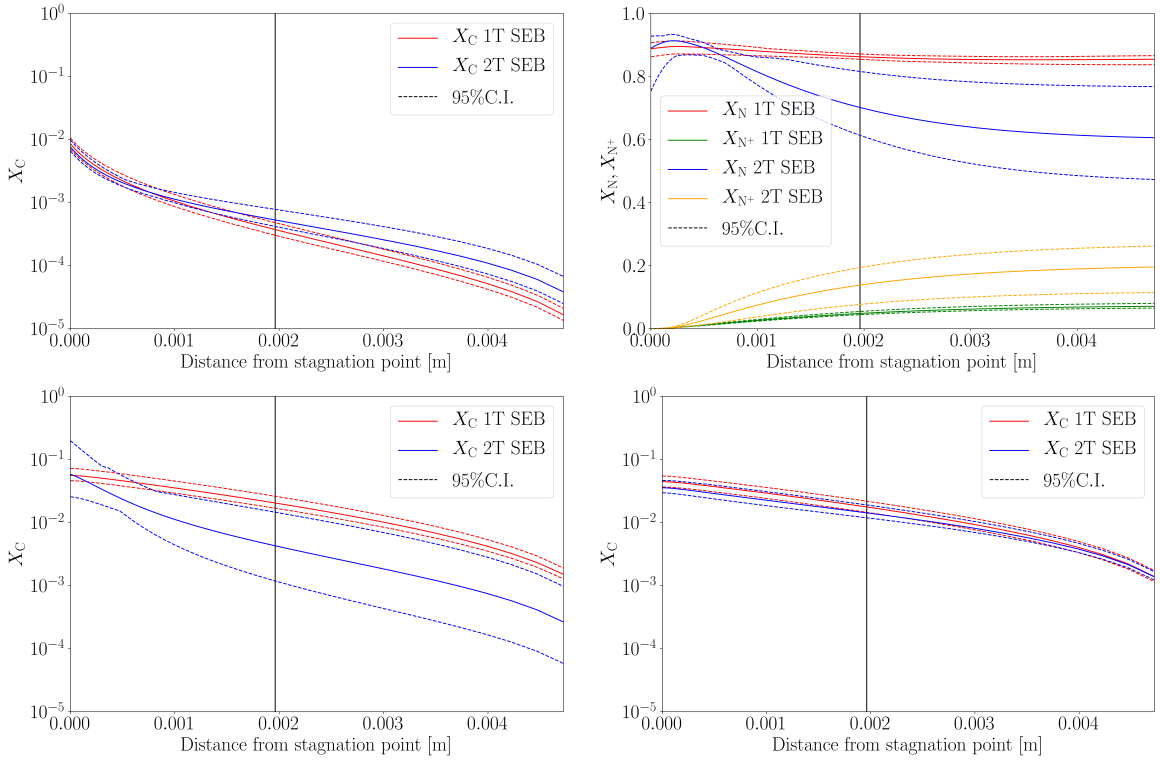

Figure A.17: Propagated mole fraction profiles for the 1T SEB ($\mathcal{M}_2$) and 2T SEB ($\mathcal{M}_3$) models with their respective confidence intervals for case G4. The vertical black lines represent the position of the spectrometer measurement.

References

[1] C. Davies, M. Arcadi, Planetary Mission Entry Vehicles Quick Reference Guide. Version 3.0., Tech. rep., NASA Technical Report 2006-3401 (2006).

[2] B. Helber, A. Turchi, J. B. Scoggins, A. Hubin, T. E. Magin, Experimental investigation of ablation and pyrolysis processes of carbon-phenolic ablators in atmospheric entry plasmas (2016), International Journal of Heat and Mass Transfer 100 (2016) 810–824. doi:https://doi.org/10.1016/j.ijheatmasstransfer.2016.04.072.

[3] H. W. Goldstein, The reaction of active nitrogen with graphite (1964), The Journal of Physical Chemistry 68 (1) (1964) 39–42. doi:10.1021/j100783a007.

[4] C. Park, D. Bogdanoff, Shock-tube measurements of nitridation coefficients of solid carbon (2006), Journal of Thermophysics and Heat Transfer 20 (3).

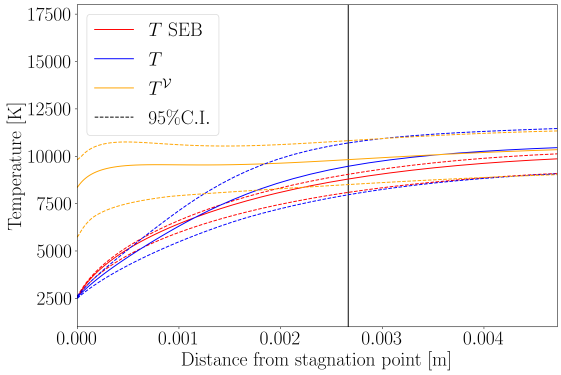

Figure A.18: Propagated temperature profiles for the 1T SEB ($\mathcal{M}_2$) and 2T SEB ($\mathcal{M}_3$) models with their respective confidence intervals for case G6. The vertical black line represents the position of the spectrometer measurement.

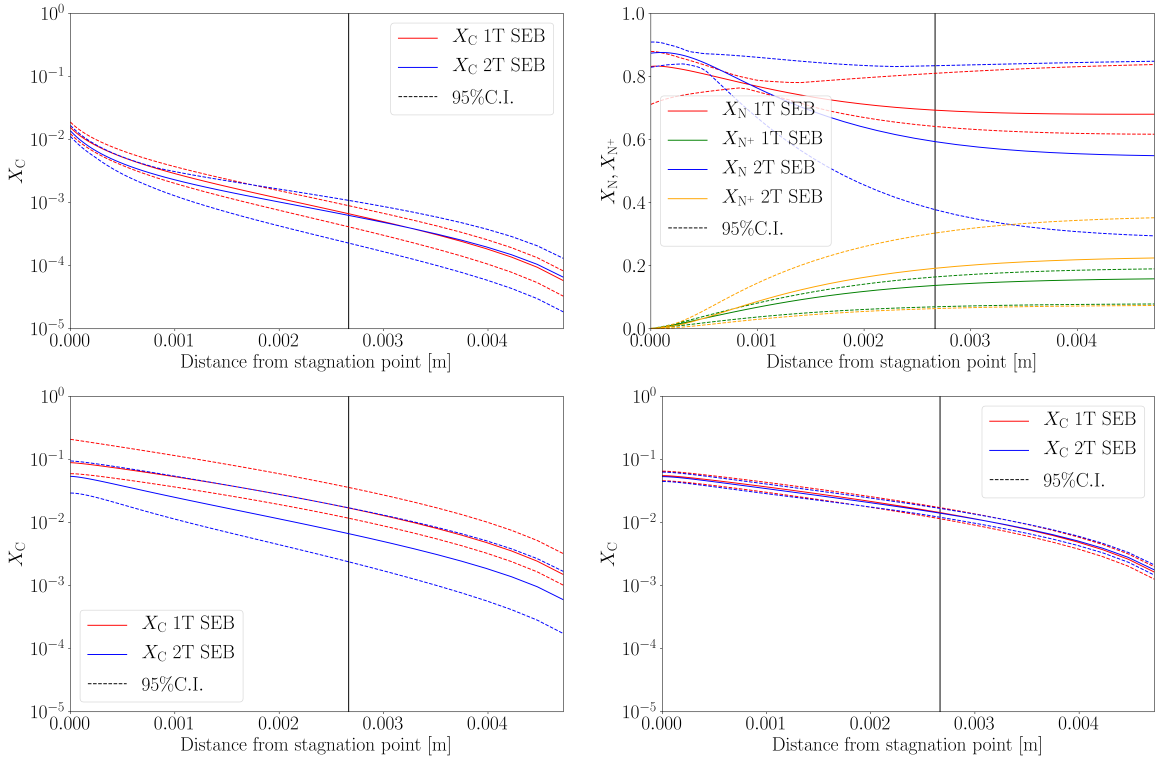

Figure A.19: Propagated mole fraction profiles for the 1T SEB ($\mathcal{M}_2$) and 2T SEB ($\mathcal{M}_3$) models with their respective confidence intervals for case G6. The vertical black lines represent the position of the spectrometer measurement.

[5] T. Suzuki, K. Fujita, T. Sakai, Graphite nitridation in lower surface temperature regime (2010), Journal of Thermophysics and Heat Transfer 24(1).

[6] L. Zhang, D. A. Pejaković, J. Marschall, M. Dougherty, D. G. Fletcher, Laboratory investigation of the active nitridation of graphite by atomic nitrogen (2012), Journal of Thermophysics and Heat Transfer 26 (1) (2012) 10–21. doi:10.2514/1.T3612.

[7] C. Park, D. W. Bogdanoff, Shock-tube measurement of nitridation coefficient of solid carbon (2006), Journal of Thermophysics and Heat Transfer 20 (3) (2006) 487–492. doi:10.2514/1.15743.

[8] Uncertainty quantification of a graphite nitridation experiment using a bayesian approach (2011), Experimental Thermal and Fluid Science 35 (8) (2011) 1588–1599. doi:https://doi.org/10.1016/j.expthermflusci.2011.07.010.

[9] A. del Val, Bayesian calibration and assessment of gas-surface interaction models and experiments in atmospheric entry plasmas, Ph.D. thesis, IPP/VKI (2021).

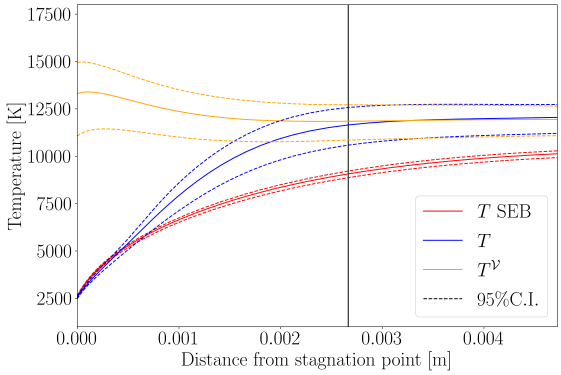

Figure A.20: Propagated temperature profiles for the 1T SEB ($\mathcal{M}_2$) and 2T SEB ($\mathcal{M}_3$) models with their respective confidence intervals for case G7. The vertical black line represents the position of the spectrometer measurement.

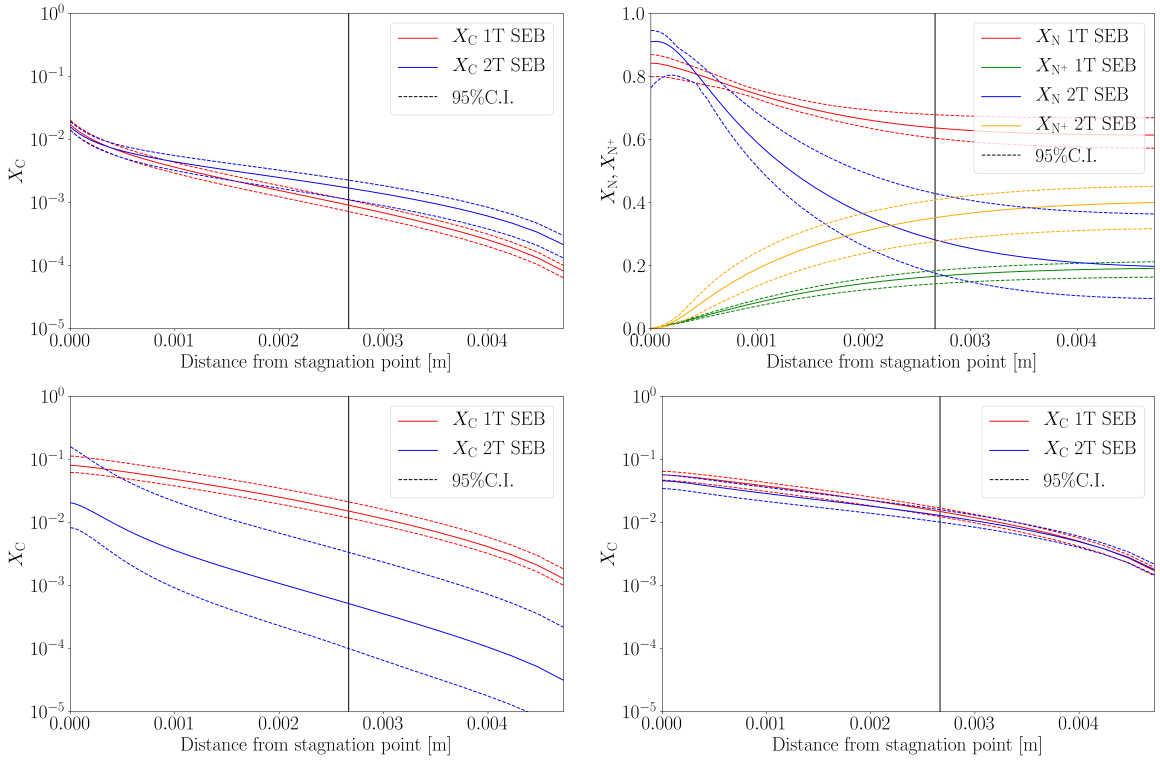

Figure A.21: Propagated mole fraction profiles for the 1T SEB ($\mathcal{M}_2$) and 2T SEB ($\mathcal{M}_3$) models with their respective confidence intervals for case G7. The vertical black lines represent the position of the spectrometer measurement.

[10] B. Helber, A. Turchi, T. E. Magin, Determination of active nitridation reaction efficiency of graphite in inductively coupled plasma flows (2017), Carbon 125 (2017) 582–594.

[11] A. Munafò, Multi-scale models and computational methods for aerothermodynamics, Ph.D. thesis, VKI/ECP (2014).

[12] J. B. Scoggins, V. Leroy, G. Bellas-Chatzigeorgis, B. Dias, T. E. Magin, Mutation++: Multicomponent thermodynamic and transport properties for ionized gases in c++ (2020), SoftwareX 12 (2020) 100575. doi:https://doi.org/10.1016/j.softx.2020.100575.

[13] B. Helber, Material response characterization of low-density ablators in atmospheric entry plasmas, Ph.D. thesis, VUB/VKI (2016).

[14] F. Panerai, Aerothermochemistry characterization of thermal protection systems, Ph.D. thesis, Università degli Studi di Perugia,von Karman Institute for Fluid Dynamics (2012).

[15] A. Viladegut, Assessment of Gas-Surface Interaction Modelling for Lifting Body Re-Entry Flight Design, Ph.D. thesis, VKI/UPC (2017).

[16] A. Klomfass, S. Müller, A quasi-onedimensional approach for hypersonic stagnation-point flows (1996), Technical report, Internal Report.

Figure A.22: Left: 100,000 samples from the joint posterior distributions of α and β. Right: posterior marginal distributions of α and β under conditions (from top to bottom) G4, G6 and G7.

[17] K. Kitamura, E. Shima, Towards shock-stable and accurate hypersonic heating computations: A new pressure flux for ausm-family schemes (2013), Journal of Computational Physics 245 (2013) 62–83. doi:https://doi.org/10.1016/j.jcp.2013.02.046.

[18] G. Bellas-Chatzigeorgis, P. F. Barbante, T. E. Magin, Energy accommodation coefficient calculation methodology using state-to-state catalysis applied to hypersonic flows (2020), AIAA Journal 58 (1) (2020) 278–290. doi:10.2514/1.J058543.

[19] P. Chen, N. Zabaras, I. Bilionis, Uncertainty propagation using infinite mixture of gaussian processes and variational bayesian inference (2015), Journal of Computational Physics 284 (2015) 291–333.

[20] A. Turchi, D. Bianchi, F. Nasuti, R. Marocco, Ablative material behavior in oxygen/methane thruster environment, in: 43rd AIAA Thermophysics Conference, 2012.
[21] J. Lachaud, J. Scoggins, T. Magin, M. Meyer, N. Mansour, A generic local thermal equilibrium model for porous reactive materials submitted to high temperatures (2017), International Journal of Heat and Mass Transfer 108 (2017) 1406–1417. doi:https://doi.org/10.1016/j.ijheatmasstransfer.2016.11.067.
URL https://www.sciencedirect.com/science/article/pii/S0017931016312480
[22] B. Dias, Thermal ablation and radiation modeling of meteor phenomena, Ph.D. thesis, Université Catolique de Louvain/VKI (2020).
[23] T. Suzuki, K. Fujita, T. Sakai, Graphite nitridation in lower surface temperature regime (2010), Journal of Thermophysics and Heat Transfer 24 (1) (2010) 212–215. doi:10.2514/1.43265.
[24] S. V. Zhluktov, T. Abe, Viscous shock-layer simulation of airflow past ablating blunt body with carbon surface (1999), Journal of Thermophysics and Heat Transfer 13(1) (1999) 50–59.
[25] A. Turchi, A gas-surface interaction model for the numerical study of rocket nozzle flow sover pyrolyzing ablative materials, Ph.D. thesis, Sapienza Università di Roma (2013).
[26] D. Bianchi, F. Nasuti, E. Martelli, Coupled analysis of flow and surface ablation in carbon-carbon rocket nozzles (2009), Journal of Spacecraft and Rockets 46(3).
[27] D. Bianchi, F. Nasuti, M. Onofri, E. Martelli, Thermochemical erosion analysis for graphite/carbon-carbon rocket nozzles (2011), Journal of Propulsion and Power 27(1).
[28] L. Landau, E. Teller, Theory of sound dispersion, Physikalische Zeitschrift der Sowjetunion, 1936.
[29] B. Halpern, D. Rosner, Chemical energy accommodation at catalyst surfaces. flow reactor studies of the association of nitrogen atoms on metals at high temperatures (1978), Journal of the Chemical Society, Faraday Transactions 74 (1978) 1883–1912.
[30] M. Capriati, G. Bellas-Chatzigeorgis, A. Turchi, B. Helber, T. E. Magin, Thermal non-equilibrium modeling for ablative gas-surface interaction, in: HiSST: 2nd International Conference on High-Speed Vehicle Science and Technology, 2020.
[31] G. Bellas, Development of advanced gas-surface interaction models for chemically reacting flows for re-entry conditions, Ph.D. thesis, VKI/Politecnico di Milano (2018).
[32] D. S. Sivia, J. Skilling, Data Analysis - A Bayesian Tutorial, 2nd Edition, Oxford Science Publications, Oxford University Press, 2006.
[33] H. Haario, E. Saksman, J. Tamminen, An adaptive metropolis algorithm (1998), Bernoulli 7 (1998) 223–242.
[34] H. Jeffreys, Theory of Probability, 3rd Edition, Oxford, Oxford, England, 1961.
[35] R. E. Kass, A. E. Raftery, Bayes factors (1995), Journal of the American Statistical Association 90 (430) (1995) 773–795. doi:10.1080/01621459.1995.10476572.
[36] G. Schwarz, Estimating the dimension of a model (1978), The Annals of Statistics 6 (1978) 461–464.
[37] H. Akaike, A new look at the statistical model identification (1974), IEEE Transactions on Automatic Control 19 (6) (1974) 716–723. doi:10.1109/TAC.1974.1100705.
[38] J. Ding, V. Tarokh, Y. Yang, Model selection techniques: An overview (2018), IEEE Signal Processing Magazine 35 (6) (2018) 16–34. doi:10.1109/MSP.2018.2867638.
[39] T. C. Chamberlin, The method of multiple working hypotheses (1965), Science 148 (3671) (1965) 754–759. doi:10.1126/science.148.3671.754.
[40] J. A. Hoeting, D. Madigan, A. E. Raftery, C. T. Volinsky, Bayesian model averaging: A tutorial (1999), Statistical Science 14 (4) (1999) 382–401.
[41] H. V. Roberts, Probabilistic prediction (1965), Journal of the American Statistical Association 60 (309) (1965) 50–62. doi:10.1080/01621459.1965.10480774.
[42] W. G. Vicenti, C. H. Kruger, Introduction to Physical Gas Dynamics, Krieger Pub Co., 1975.