



HAL
open science

Robust calibration of an air-carbon ablation model employing Plasmatron and molecular beam data

Vittorio Piro, Michele Capriati, Federico Bariselli, Pietro Marco Congedo, Thierry E. Magin

► To cite this version:

Vittorio Piro, Michele Capriati, Federico Bariselli, Pietro Marco Congedo, Thierry E. Magin. Robust calibration of an air-carbon ablation model employing Plasmatron and molecular beam data. von Karman Institute for Fluid Dynamics. 2025. <hal-05214270v1>

HAL Id: hal-05214270

<https://inria.hal.science/hal-05214270v1>

Submitted on 19 Aug 2025 (v1), last revised 2 Dec 2025 (v2)

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

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



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Robust calibration of an air-carbon ablation model employing Plasmatron and molecular beam data

Vittorio Piro

von Karman Institute for Fluid Dynamics, Aeronautics and Aerospace Department

(Dated: June 25, 2025)

Understanding and accurately predicting gas-surface interaction phenomena during a vehicle's re-entry phase is critical for the design of thermal protection materials. In this work, we aim to improve the robustness of an air-carbon ablation model by calibrating its reaction rates using both low-pressure molecular beam data and high-pressure Plasmatron data, as well as numerical models, following a Bayesian approach. Firstly, the subset of nitrogen chemical reactions is inferred, resulting in a finite-rate nitrogen model successfully calibrated and validated at both low- and high-pressure. A surrogate-based Bayesian framework is employed, comparing artificial neural network and Kriging approaches to mimic the nitrogen model predictions at high-pressure and accelerate posterior sampling. Then, for the first time, atomic oxygen reaction rates are also included in the calibration, along with low-pressure oxygen observations. Since the experimental data used for calibration do not include high-pressure conditions, the oxygen reaction rates for accuracy at high pressure could not be fully characterised. Nevertheless, the resulting models demonstrated superior performance compared to the state-of-the-art model under low-pressure conditions.

I. INTRODUCTION

During the re-entry phase of a hypersonic vehicle, the spacecraft reaches extremely high speeds (7 – 11 km/s) [1], resulting in significant kinetic energy to dissipate. Additionally, the bow shock formed in front of the blunt body generates a high-enthalpy environment leading to extremely high-temperatures on the surface of the spacecraft (up to 2000 K) [2]. To ensure the mission's safety, a Thermal Protection Material (TPM) is essential to absorb and dissipate this extreme heat. Among the multi-physics phenomena occurring in the shock layer of a spacecraft during the re-entry phase, the deviation from thermo-chemical equilibrium and the Gas-Surface Interaction (GSI) phenomena make the design of the TPM a very challenging task, requiring accurate experiments and simulations [3]. In the context of GSI phenomena, we refer to strong temperature and velocity gradients within the Boundary Layer (BL) developing in front of the re-entry vehicle that lead a chemically reacting gas to diffuse towards the surface and interact with it through mass and heat exchange. An illustration of the aerothermochemical phenomena occurring along the stagnation line of a hyperosonic re-entry vehicle is proposed in Fig. 1.

To design the TPM, flight testing, experimental on-ground facilities, and physico-chemical models are employed to characterise the state of the flow under these extreme conditions. The former provides direct measurements under real operational conditions, capturing properly also non-linear and unsteady effects. Compared to flight testing, on-ground facilities are less costly but have inherent limitations in fully replicating real-flight condi-

tions. This makes direct extrapolation from ground tests to flight challenging. Additionally, ground testing alone is not sufficient to analyse a significant number of flight configurations. Therefore, models and simulations are essential for making design decisions.

By definition, models are mathematical representations of physics and often incorporate simplifications and/or assumptions that limit their predictive accuracy. Uncertainty Quantification (UQ) techniques are, thus, employed to address these uncertainties by characterising, propagating, and perhaps lowering them in order to produce reliable predictions [4].




FIG. 1: Overview of the aerothermochemical phenomena along the stagnation line of a blunt body during re-entry phase: bow shock, shock layer where thermo-chemical Non-Equilibrium (NEQ) and Local Thermal Equilibrium (LTE) may coexist, and reactive boundary layer characterised by GSI. Adapted from [5].

The hypersonic community has just recently started to apply UQ techniques (e.g., reconstruction of free-stream conditions or improvement of the robustness of GSI chemical models). Cortesi et al. [6, 7] addressed the difficult inverse problem of rebuilding free-stream conditions (velocity and density) for atmospheric re-entry vehicles using pressure and heat flux measurements at the stagnation point. They employed a Bayesian approach [8] which, in general, offers a robust framework to characterise the uncertainty of specific Quantities of Interest (QoIs), which cannot be measured experimentally, by considering experimental observations along with their uncertainty. To overcome the computational cost of Computational Fluid Dynamics (CFD) simulations during Bayesian inference, they employed surrogate models [9, 10]. In general, these models are trained on a significantly reduced number of simulations compared to those necessary to conduct UQ study. However, they effectively reflect the fundamental correlation between input and output.

Upadhyay et al. [11] in 2011 were the first ones to focus on the Bayesian method applied to GSI studies. They quantified the uncertainty of the nitridation reaction efficiency starting from a graphite nitridation experiment, confirming its relevance for modelling thermo-chemical ablation of TPM used in Earth re-entry vehicles. More insights on GSI phenomena were provided by Del Val et al. [12, 13]. They proposed a novel Bayesian inference method to calibrate the catalytic recombination efficiencies affecting reusable TPM by combining different experimental data from the von Karman Institute (VKI) Plasmatron facility [14]. Lower uncertainty on the catalytic recombination parameters resulted from the calibration.

Contrary to [12, 13], the present work focuses specifically on ablative TPMs, typically made of carbon fibres reinforced with a matrix of resin, which converts thermal energy into material decomposition and removal. In this case, GSI phenomena result in both catalytic and ablative chemical reactions and their study is crucial for predicting the heat flux and the mass loss experienced by the hypersonic vehicle [15]. In order to model these surface chemical reactions, Finite-Rate Chemistry (FRC) models [16] can be employed. In this case, each elementary process occurring at the wall is characterised by a chemical reaction. One such model, used in this work, is the Air-Carbon Ablation (ACA) model from Prata et al. [17]. They interpreted the data from both Molecular Beam (MB) experiments [18] and plasma wind tunnels [2, 19] to present a new deterministic model that comprises 20 reaction mechanisms describing reactions between impinging O, N, and O₂ species with a carbon surface (e.g., adsorption, desorption, oxidation, nitrida-

tion, and recombination) and accounting for both dependencies of the reactions on temperature and pressure.

Originally, the parameters of the ACA model were determined deterministically, without accounting for both model and experimental uncertainties. This led to a model that is not fully consistent with experimental values. Thus, Capriati et al. [15] decided to perform a stochastic calibration of the model parameters by informing a subset of FRC model reactions, including only nitrogen atoms. They combined experimental observations from both MB experiments [18] and the VKI Plasmatron facility [2]. The latter allows for characterising the macroscopic properties of GSI phenomena (e.g., recession, surface temperature, and stagnation point heat flux) occurring on carbon-based TPMs. On the other hand, MB experiments allow the study of GSI phenomena from a microscopic point of view by bombarding a sample placed in a high-vacuum chamber with a beam of atoms/molecules. Although GSI phenomena are completely isolated in MB experiments, the pressure is very far from the real value. Therefore, Capriati et al. decided to employ also the more representative data from the VKI high-enthalpy plasma wind tunnel. This time, both experimental and model uncertainties were included during the analysis. The goal was to improve the robustness of the predictions of the model by computing new reaction rates.

Bandera [20] continued Capriati’s work, seeking to validate the updated stochastic model on new experimental data obtained in the VKI Plasmatron facility at pressures higher than those used for calibration [21]. The validation process consisted of propagating both model and experimental uncertainties for the new testing campaign. Bandera showed that Prata’s version is still more accurate than Capriati’s one in describing GSI at higher pressure. Thus, the validation failed, suggesting the need to perform a new calibration including these new experimental data.

The contribution of this work to the hypersonic community is to provide a robust FRC model capable of accurately describing the complex GSI phenomena between the ablative TPM and the surrounding gas during atmospheric re-entry. To this end, we perform a robust calibration of the ACA model, following a Bayesian approach, continuing the work of Capriati [15], to improve its predictivity under different pressure conditions than those used so far. We employ advanced UQ techniques to account for both experimental and model uncertainties. Since UQ analysis requires numerous calculations, we also determine the surrogate model to be used for achieving an adequate level of efficiency when CFD solvers are


too expensive to use.

Until now, the stochastic calibration was performed only on nitrogen as a test gas, but the ACA model also includes oxygen reactions. In fact, this project aims to extend, for the first time, the calibration to reactions involving both atomic nitrogen and atomic oxygen, a necessary step to elevate the model to true air-relevant conditions. Experimental data from a recent campaign at VKI [22] will serve as a benchmark to rigorously assess the predictive performance of this extended air-carbon ablation model.


The present work is structured as follows. Section II introduces the GSI phenomena, along with the FRC model employed in this study. Section III presents the ground-testing experiments used in the calibration process. Section IV details the overall methodology, with a focus on the numerical models and the Bayesian calibration approach. Sections V A and V B are dedicated to the stochastic analysis. Specifically, Section V A describes the construction and validation of surrogate models, while Section V B presents the Bayesian inference of the nitrogen reaction model. Section V C addresses the extended calibration of the air-carbon GSI model, combining both atomic nitrogen and oxygen measurements. Finally, Section VI summarises the main conclusions of the study and outlines possible future developments.

II. FINITE-RATE GAS-SURFACE INTERACTION

The applications treated in this work are characterised by high-temperatures (up to 10 000 K after the bow shock) and relatively low-pressures (200-20 000 Pa). Those high-temperatures cause reactions that alter the chemical composition of the gas, promoting dissociation of the molecules, leading to a multi-component, chemically reacting gas description. As said before, the resulting atoms diffuse towards the surface of the re-entry vehicle, reacting with it through ablative and catalytic reactions. With ablative reactions, we mean that an atom from the surface (b) reacts with one from the gas phase, e.g., oxidation of a carbon surface:



Catalytic reactions, on the other hand, occur when the surface (s) acts as a catalyst by promoting the recombination of two atoms in the gas-phase, for instance:



In both cases, not only the chemical composition of the BL is affected, but also the surface heat flux increases

because the reaction is exothermic. In this regard, the mass and energy balances on the reacting surfaces must be used to derive the material surface response.

Assuming a steady-state condition on the surface and considering an infinitesimal volume containing both the gas and the surface, the balance for a conservative quantity \mathbf{F} decreases to:

$$[\mathbf{F}_g - \mathbf{F}_s] \cdot \mathbf{n} = \dot{\Omega}_{\text{surf}}, \quad (3)$$

where \mathbf{F}_g and \mathbf{F}_s , respectively, are the fluxes in/from the gas- and the solid-interface, \mathbf{n} is the normal to the surface pointing the gas-phase, and $\dot{\Omega}_{\text{surf}}$ is the surface source term. As will be outlined in Section III, the surface temperature of the carbon samples is known for each experimental condition, therefore one needs to solve only a Surface Mass Balance (SMB) for each species, sketched in Fig. 2. The surface chemistry generates a species gradient at the surface, driving diffusion.




FIG. 2: Sketch of the SMB.

In case of SMB, \mathbf{F}_g consists of a convective part, $\rho_i \mathbf{u}$, and a diffusive part, \mathbf{j}_i . By assuming a frame of reference attached to the bulk of the solid, the velocity is expressed as the difference between the blowing velocity,

$$\mathbf{u}_g = \frac{\dot{m}_{\text{blow}}}{\rho}, \quad (4)$$

and the recession velocity,

$$\mathbf{u}_r = \frac{\dot{m}_{\text{abl}}}{\rho_s}, \quad (5)$$

where ρ and ρ_s are, respectively, the density of the gas and the material. Specifically, \dot{m}_{blow} is the total mass exiting the interface towards the gas phase (mostly due to chemical ablation and pyrolysis), while \dot{m}_{abl} is the mass flux only due to chemical ablation. In this work, any other phenomena, such as spallation, mechanical removal, and pyrolysis, are neglected. Moreover, due to the strong disparity in densities, we assume that $\mathbf{u}_r \ll \mathbf{u}_g$ [3]. Thus, \dot{m}_{blow} is mostly due to chemical ablation. Concerning the source term, $\dot{\Omega}_{\text{surf}}$ appear only due to chemical reactions which, for this work, can be only of catalytic and/or ablative nature. Therefore, the SMB for each species i reads:

$$\dot{w}_i = (\rho_i \mathbf{u}_g + \mathbf{j}_i) \cdot \mathbf{n} \quad \forall i \in [1, n_s]. \quad (6)$$

The closure of the SMB is obtained once the chemical production rate \dot{w}_i for each species is known. This can be done by employing phenomenological approaches, often referred to as "gamma" model [23], or FRC models [16]. In the first case, we necessitate the probability that a certain macroscopic reaction takes place. In this work, instead, we adopt the ACA FRC approach [17]. Thus, the surface chemical reaction terms read as follows:

$$\dot{w}_i = M_i \sum_r \nu_{i,r} k_r \prod_j \left(\frac{\rho_j}{M_j} \right)^{\nu_{j,r}}, \quad (7)$$

where M_i is the molar mass for the species i , $\nu_{i,r}$ is the stoichiometric reaction coefficient for the species i in the reaction r , and k_r is the reaction rate coefficient. Each elementary reaction occurring at the surface is described individually, with the rate k_r depending on the specific reaction mechanism.

Air Carbon Ablation model Prata et al. [17] derived 20 reaction mechanisms describing the ablative behaviour of a carbon surface when impinged by air. The rates were tuned starting from both low-pressure [18] and high-pressure [2, 19] experiments. In this work, we are interested in atomic oxygen and atomic nitrogen reactions, neglecting molecular oxygen contribution. They are reported, respectively, in Tab. I and Tab. II.

No Reaction	Rate Coefficient
1 $O + (s) \rightarrow O(s)$	$\frac{F_O}{B} \cdot X_1$
2 $O(s) \rightarrow O + (s)$	$\frac{2\pi m_O k_B^2 T_w^2}{A_v B h^3} \exp\left(-\frac{E_{a,2}}{T_w}\right)$
3 $O + O(s) + C(b) \rightarrow CO + O + (s)$	$\frac{F_O}{B} \cdot X_3 \exp\left(-\frac{E_{a,3}}{T_w}\right)$
4 $O + O(s) + C(b) \rightarrow CO_2 + (s)$	$\frac{F_O}{B} \exp\left(-\frac{E_{a,4}}{T_w}\right)$
5 $O + (s) \rightarrow O^*(s)$	$\frac{F_O}{B} \cdot X_5$
6 $O^*(s) \rightarrow O + (s)$	$\frac{2\pi m_O k_B^2 T_w^2}{A_v B h^3} \exp\left(-\frac{E_{a,6}}{T_w}\right)$
7 $O + O^*(s) + C(b) \rightarrow CO + O + (s)$	$\frac{F_O}{B} \cdot X_7 \exp\left(-\frac{E_{a,7}}{T_w}\right)$
8 $O^*(s) + O^*(s) \rightarrow O_2 + 2(s)$	$\sqrt{\frac{A_v}{B}} \cdot F_{O,2D} \cdot X_8 \exp\left(-\frac{E_{a,8}}{T_w}\right)$
9 $O(s) + O(s) \rightarrow O_2 + 2(s)$	$\sqrt{\frac{A_v}{B}} \cdot F_{O,2D} \cdot X_9 \exp\left(-\frac{E_{a,9}}{T_w}\right)$

TABLE I: Subset of the FRC reactions involving oxygen in the ACA model.

Considering a surface characterised by a total active site density B , this quantity can be expressed as the sum of the surface density of free reactive sites, denoted by s , and the surface density of sites already occupied by a gas-phase species G , denoted by $G(s)$. Here, $G = O$ or N . The term F_G denotes one quarter of the mean thermal speed of the atomic species, while $F_{G,2D}$ represents the mean thermal speed of the mobile adsorbed species on the surface; both are functions of the gas temperature.

In these expressions, m_G is the mass of atomic species, k_B and h are the Boltzmann and the Planck constants, respectively, and T_w is the surface temperature. Quantity A_v denotes Avogadro's number.

The symbols $E_{a,i}$ and X_i refer to the activation energies and pre-exponential factors, respectively, which are to be inferred through a Bayesian approach. The deterministic values proposed by Prata et al. are not reported in this section, but a detailed comparison with our stochastic predictions will be provided in Section V B and V C.

Looking at Tab I, Prata et al. proposed that there are two types of adsorbed atomic oxygen: one relatively weak, $O(s)$, and one relatively strong, $O^*(s)$. The latter has a higher desorption energy, which allows the model to capture the trend of continuous CO formation at high-temperatures, when $O(s)$ would already have desorbed [17]. Reactions (1) and (5) are the corresponding adsorption mechanisms. The model originally assigns different selectivities for adsorption in these states, which add up to unity. The desorption back into gas-phase of these two species is described by reactions (2) and (6): $O(s)$ desorbs more easily due to the lower desorption energy (equivalent to a C-O single bond). Reaction (3) describes the gas-phase dependent CO formation through $O(s)$, while reaction (7) due to $O^*(s)$. Both are important for model accuracy in describing the experiments at higher pressures. CO_2 formation is modelled in reaction (4) as a gas-phase dependent mechanism involving only $O(s)$, and for simplicity, no CO_2 mechanism dependent on $O^*(s)$ is included, as CO_2 production diminishes at higher temperatures where $O^*(s)$ persists [17]. Finally, both reactions (8) and (9) represent the gas-independent recombination mechanisms through $O^*(s)$ and $O(s)$. Prata et al. proposed an activation energy for both the reactions to best match the slope of increasing probability with temperature observed in the experimental data.

No Reaction	Rate Coefficient
10 $N + (s) \rightarrow N(s)$	$\frac{F_N}{B} \exp\left(-\frac{E_{a,10}}{T_w}\right)$
11 $N(s) \rightarrow N + (s)$	$\frac{2\pi m_N k_B^2 T_w^2}{A_v B h^3} \exp\left(-\frac{73971}{T_w}\right)$
12 $N + N(s) + C(b) \rightarrow CN + N + (s)$	$\frac{F_N}{B} \cdot X_{12} \exp\left(-\frac{E_{a,12}}{T_w}\right)$
13 $N + N(s) \rightarrow N_2 + (s)$	$\frac{F_N}{B} \cdot X_{13} \exp\left(-\frac{E_{a,13}}{T_w}\right)$
14 $N(s) + N(s) \rightarrow N_2 + 2(s)$	$\sqrt{\frac{A_v}{B}} \cdot F_{N,2D} \cdot X_{14} \exp\left(-\frac{E_{a,14}}{T_w}\right)$
15 $N(s) + C(b) \rightarrow CN + (s)$	$X_{15} \exp\left(-\frac{E_{a,15}}{T_w}\right)$

TABLE II: Subset of the FRC reactions involving nitrogen in the ACA model.

From Tab. II, reactions (10) and (11) of the finite-rate

nitridation model consist of adsorption and desorption processes. The N-atom adsorption is modelled as an activated process to capture the sharp drop in N_2 formation at low-temperatures in the MB experiments [17]. The desorption of N(s) into N is based on transition state theory: Prata et al. opted for a desorption energy corresponding to a double bond between C and N for describing the non-decreasing efficiency observed at high-temperature in the MB observations. For nitridation, they proposed one gas-phase dependent CN formation mechanism, reaction (12), for model accuracy at high-pressures. At the same time, one gas-phase independent mechanism, reaction (15), is also present for the trend at low-pressures. To conclude, two mechanisms of recombination are included. Reaction (13) represents a gas-dependent recombination mechanism, necessary to describe recombination processes over a wide range of pressures. Reaction (14) refers to gas-independent recombination mechanism observed in MB experiments [17].

During this work, we first focus on the finite-rate nitridation model to propose a new calibration based on the work of Capriati [15], which includes new high-pressure experimental points [21]. Then, we consider both models together to perform a joint calibration of the ACA model. This latter calibration is essential to ensure accurate predictions of GSI during atmospheric re-entry. Atomic nitrogen and oxygen coexist in the post-shock layer and compete for the same active sites on the surface, affecting, for instance, recombination rates and surface coverage. This calibration captures these coupled effects and improves the reliability of surface heat flux and ablation rate predictions.

III. EXPERIMENTS IN GROUND TESTING FACILITIES

As outlined in Section I, ground-based experiments are essential for the accurate design of the TPM of a re-entry vehicle. These experimental campaigns also provide the necessary data to perform a Bayesian inversion. In this section, we present the experimental campaigns used by Prata et al. [17] to deterministically calibrate the ACA model, as well as those used in our study to stochastically infer the same model parameters. The analysis conducted in this work draws on data from both MB and VKI Plasmatron experiments, which serve complementary roles. Specifically, MB observations offer valuable insights into the microscopic mechanisms of GSI, although the pressure conditions are not representative of actual re-entry scenarios. Conversely, plasma wind tunnel experiments enable the study of GSI phenomena on a macroscopic scale under pressure conditions that closely

mimic those experienced during atmospheric re-entry.

A. Molecular beam-surface scattering experiments

Murray et al. [18] performed molecular beam-surface scattering experiments by bombarding a carbon sample placed in a high-vacuum chamber ($p < 0.1$ Pa) with low-energy continuous beams of oxygen or nitrogen atoms. For the oxygen beam, an 85 mbar mixture of 5% O_2 in He was discharged with a high-pressure radio frequency source which allows the control of energy and direction. The nitrogen beam was produced using a 95 mbar mixture of 2.5% N_2 in He. The two beams were then expanded through a 0.48 mm diameter water-cooled quartz nozzle. The resulting beams have an average velocity of 2000 m/s and are composed, in mole fractions, of 48% O and 52% O_2 for the oxygen beam, and 18% N and 82% N_2 for the nitrogen beam [18]. Because of continuous beams, the determination of the product flux is not dependent on the time at which the products were formed on the surface. Moreover, due to the very low-pressure, they were able to study the interaction between the gas and the surface with the lowest number of collisions between the gas molecules.

The beams were directed at a heated sample surface, and the products scattering from the surface were detected with a rotatable mass spectrometer. This instrument identifies the number of scattering molecules as a function of time and their type according their mass and final direction. For the oxygen beam, four products were observed: O, CO, O_2 , and CO_2 . On the other hand, CN and N_2 for the nitrogen beam. The QoIs are the reaction efficiencies of nitridation, oxidation, and recombination, defined as:

$$\gamma = \frac{\text{flux of specific product}}{\text{flux of N/O atoms onto surface}}, \quad (8)$$

and quantified as a function of the surface temperature. The temperature of the sample was changed gradually from one experiment to another (from 800 to 1873 K). Before each test, the sample was first annealed at 1873 K for one hour to ensure reproducibility of the data. No results are presented in this section, as a comprehensive comparison with our model predictions will be provided in Section VB and VC.

B. VKI Plasmatron experiments

The high-pressure data chosen by Prata et al. [17] and Capriati et al. [15] to infer deterministically and statistically, respectively, the parameters of the ACA model

come from a subsonic experimental campaign carried out in the VKI Plasmatron facility [2]. Then, when Bandera [20] tested Capriati's model on experimental observations at higher pressure [21], always from the VKI Plasmatron facility, the model resulted in invalidation. Therefore, in this work, we decided to include also those experimental points in our calibration. In addition, a new experimental campaign was recently conducted in the VKI Plasmatron to provide the experimental observations needed to calibrate and validate the ACA model [22]. These new observations will be used to validate our new ACA model in future work, for which both atomic nitrogen and oxygen reactions have been calibrated.

The VKI Plasmatron is the most powerful Induction-Coupled-Plasma (ICP) wind tunnel in the world, equipped with an ICP torch powered by a high-frequency, high-power, and high-voltage generator. A high-purity plasma is obtained by heating gas through a coil and is blown in the form of a jet inside a test chamber at sub-atmospheric pressure. When the Local Heat Transfer Simulation [24] constraints are satisfied, the Plasmatron can replicate the reacting BL experienced by an object travelling at hypersonic speed by maintaining a high-enthalpy subsonic flow for an extended period, enabling the characterisation of the TPM. The control variables of the facility are the power supplied P , along with the chamber pressure p_s and the mass flow rate \dot{m} .

1. Experimental campaign 2017 - Nitrogen, $p_s = 15$ hPa

Helber et al. [2] performed a subsonic campaign to study the temperature dependence of the nitridation reaction efficiency on carbon surfaces by hitting a 25 mm radius hemispherical graphite sample with a nitrogen plasma flow. The pressure of the chamber containing the sample was set at $p_s = 15$ hPa. A digital camera was used to measure the recession rate at the stagnation point, and the surface temperature T_w was obtained with a pyrometer. A spectrometer was employed to retrieve both the gas temperature and the CN concentration.

CASE	p_d	P	T_w	\dot{m}_b	ρ_e	T_e	u_e
	[Pa]	[kW]	[K]	[g m ⁻² s ⁻¹]	[g m ⁻³]	[K]	[m/s]
G4	231	280	2225	2.49 ± 0.91	0.236	10005	554
G5	268	330	2410	2.89 ± 0.97	0.225	10280	562
G6	312	370	2535	4.41 ± 0.80	0.184	11040	846
G7	330	390	2575	4.56 ± 0.70	0.189	10970	859

TABLE III: Plasmatron test conditions (nitrogen, 15 hPa): dynamic pressure p_d , generator power P , surface temperature T_w , mass blowing rate \dot{m}_b , BL edge conditions (density ρ_e , temperature T_e , and velocity u_e).

They analysed seven different conditions, gradually increasing the power supplied to the Plasmatron, resulting in an increase in both T_w and recession. For each condition, the mass blowing rate \dot{m}_b was measured. For the first three conditions, no recession or mass loss was observed. In addition, the BL conditions (temperature T_e and velocity u_e) were reconstructed using a rebuilding procedure taking as input the dynamic pressure p_d and the heat flux q_w , along with p_s and PW (together with its efficiency) [2]. In Tab. III, the four test conditions are summarized.

2. Experimental campaign 2022 - Nitrogen, $p_s = 100, 200$ hPa

A second experimental campaign, at higher pressure, was carried out at VKI in 2022, in the frame of the US AFOSR research grant "PLASMUT" [21]. Accounting for these experimental points in the calibration allow us to study the reaction pressure dependence of the ACA model and to overcome the invalidation problem encountered by Bandera [20].

Two different types of raw graphite have been tested. The first is a superfine grain type, while the second is a high-strength, wear-resistant type. In the course of this work, we refer to the first as G1 and the second as G8. Contrary to experiments at $p_s = 15$ hPa, this time the experimental data lack heat flux at the stagnation point q_w . Thus, it was not possible to apply the rebuilding procedure to retrieve the quantities at the BL edge. Tab. IV presents a summary of the test results for the two graphite grains.

CASE	p_s	p_d	P	T_w	\dot{m}_b
	[hPa]	[Pa]	[kW]	[K]	[g m ⁻² s ⁻¹]
G1	100	56.7	370	2604	1.403 ± 0.414
G1	100	66.9	400	2725	1.226 ± 0.361
G1	200	36.2	370	2660	0.920 ± 0.369
G1	200	42.5	400	2704	1.592 ± 0.399
G8	100	58.8	370	2575	1.806 ± 0.508
G8	100	50.3	400	2688	2.345 ± 0.515
G8	200	33.4	370	2601	1.292 ± 0.486
G8	200	39.3	400	2674	1.561 ± 0.434

TABLE IV: Plasmatron test conditions (nitrogen, 100, 200 hPa).

G8 samples have a slightly lower density and a higher conductivity (1760 kgm⁻³ and 130 WK⁻¹m⁻¹) than the G1 samples (1810 kgm⁻³ and 85 WK⁻¹m⁻¹), which may explain the slightly higher surface temperature for G1 samples for each test condition. In the next sections, we

refer to these experimental points as GX_Y_Z, where X indicates the type of graphite grain, Y the chamber pressure, and Z the generator power.

3. Experimental campaign 2024 - Air, $p_s = 50, 200$ hPa

A new experimental campaign has been carried out at VKI in 2024, in the frame of the US AFOSR research grant "DRAGROUND FLIGHT" [22]. Unlike previous campaigns, this time test samples with a hemispherical shape of 50 mm were used. This geometry has been extensively tested and optimised over the last few decades for material degradation and ablation testing at the VKI. Graphite samples with a density of 1800 kgm^{-3} were used as test materials for this campaign. It is possible to derive the mass blowing rate from the density, test duration, and recession. The Plasmatron test conditions with the corresponding results are listed in Tab. V.

CASE	p_s [hPa]	p_d [Pa]	P [kW]	T_w [K]	τ [s]	Recession [mm]
2	200	7.16	150	1837	8451	0.0055
3	50	41.9	174	1879	850	0.0043
4	50	41.6	174	1761	1018	0.0051
5	50	131.7	350	2604	785	0.0086
6	200	31.8	270	2528	783	0.0083

TABLE V: Plasmatron test conditions (air, 50, 200 hPa).

IV. METHODOLOGY

From Section I, we understand that experiments and numerical simulations are necessary to study GSI phenomena. In this work, we do not consider them separately but combine them to characterise the reaction rates of the ACA FRC model. In fact, this model requires the specification of numerous reaction mechanisms and associated parameters, namely pre-exponential factor X_i , activation energies $E_{a,i}$, and total active site density B , many of which are poorly understood or lack experimental measurements. This high uncertainty suggests that a stochastic inference approach should be adopted to calibrate the model from the available experimental data. In this section, we explain the methodology for integrating the experiments described above and the corresponding numerical models. Such methodology is called Bayesian inversion [8].

As illustrated in Fig. 3, the two primary sources of information are the high-pressure measurements of mass blowing rate, \dot{m}_b , obtained from the VKI Plasmatron facility, and the low-pressure reaction efficiency data, γ ,

derived from MB experiments. Both datasets are associated with corresponding uncertainties. The VKI Plasmatron campaigns are numerically explained with CFD simulations coupled with the GSI model [5, 25]. On the other hand, MB experiments are described with a 0D reactor [17], i.e., a system of ordinary differential equations based on the law of mass action to be closed with certain assumptions. The inversion process requires repeated evaluations of the GSI model. To restore adequate levels of efficiency, the CFD solver is replaced by a surrogate model [9, 10].

The Bayesian inversion framework aims to update our prior beliefs about the model parameters using the observed data. The prior distributions are, in fact, the last input to the procedure and reflect the initial uncertainty about the parameters to be inferred. This is typically based on literature or expert assessment. The output of this procedure is called posterior distribution.

Finally, the resulting posterior distribution is used for validation. This step is essential to verify whether the inferred reaction parameters allow the model to reproduce independent experimental observations under conditions not used during calibration. It is important to acknowledge that the available experimental data may not fully constrain all model parameters. In particular, some reactions might remain poorly informed due to a lack of informations. As a result, the posterior distributions for certain parameters may still exhibit significant uncertainty, potentially limiting the model's ability to be predictive.

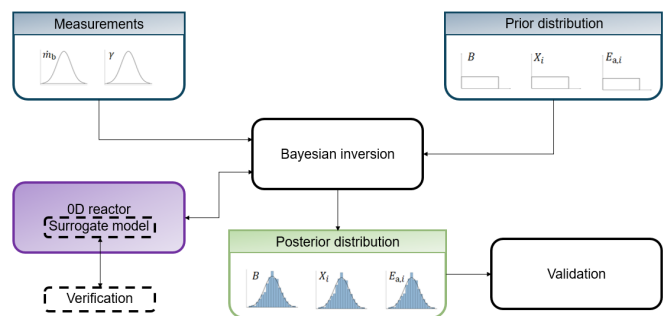


FIG. 3: Bayesian inference procedure.

The following provides a detailed description of the 0D reactor model and CFD solver employed in this study, along with the Bayesian inference framework and the surrogate modelling approach.

A. Numerical modelling

We present now the 0D model proposed by Prata et al. [17] based on the ACA FRC model and employed to describe the molecular beam experiments and the stagnation line code [25] used to simulate the reacting BL in the Plasmatron experiments.

1. 0D reactor model

Given the gas temperature and pressure, T and p , used to compute the gas composition of each species at the surface, along with the surface temperature T_w , it is possible to retrieve the model reaction efficiencies.

Defining the concentration of the gas-phase species as $[G]$ (or $[G(s)]$ if absorbed) and applying the law of mass action, the set of chemical production rates according to Tab. I reads:

$$\begin{aligned} \frac{d[O(s)]}{dt} = & k_1[O][s] - k_2[O(s)] - k_3[O][O(s)] \\ & - k_4[O][O(s)] - 2k_9[O(s)]^2 \end{aligned} \quad (9)$$

$$\begin{aligned} \frac{d[O^*(s)]}{dt} = & k_5[O][s] - k_6[O^*(s)] \\ & - k_7[O][O^*(s)] - 2k_8[O^*(s)]^2 \end{aligned} \quad (10)$$

$$\frac{d[CO]}{dt} = k_3[O][O(s)] + k_7[O][O^*(s)] \quad (11)$$

$$\frac{d[CO_2]}{dt} = k_4[O][O(s)] \quad (12)$$

$$\frac{d[O_2]}{dt} = k_8[O^*(s)]^2 + k_9[O(s)]^2 \quad (13)$$

It must be enlarged with the set from Tab. II:

$$\begin{aligned} \frac{d[N(s)]}{dt} = & k_{10}[N][s] - k_{11}[N(s)] - k_{12}[N][N(s)] \\ & - k_{13}[N][N(s)] - 2k_{14}[N(s)]^2 \\ & - k_{15}[N(s)] \end{aligned} \quad (14)$$

$$\frac{d[CN]}{dt} = k_{12}[N][N(s)] + k_{15}[N(s)] \quad (15)$$

$$\frac{d[N_2]}{dt} = k_{13}[N][N(s)] + k_{14}[N(s)]^2 \quad (16)$$

Assuming conservation of the total active site density,

$$B = [O(s)] + [O^*(s)] + [N(s)] + [s], \quad (17)$$

a steady-state solution of the system can be obtained by setting the rate of change in surface coverage equal to zero:

$$\frac{d[N(s)]}{dt} = \frac{d[O^*(s)]}{dt} = \frac{d[O(s)]}{dt} = 0. \quad (18)$$

As examples, the probability of O_2 recombination and CN nitridation through O-atom and N-atom fluxes, respectively, are:

$$\gamma_{O_2} = \frac{2(d[O_2]/dt)}{N_O}, \quad \gamma_{CN} = \frac{(d[CN]/dt)}{N_N}, \quad (19)$$

where $N_G = n_G \sqrt{k_B T / (2\pi m_G)}$ is the number of flux of particles impinging the surface, and n_G the number density of the gas species.

Fig. 4 and 5 serve as a verification of the correct implementation of the 0D reactor model. The results show that our model predictions perfectly overlap with those of Prata's model under the conditions of $T = 1000$ K and $p = 0.024$ Pa. This specific combination of temperature and pressure was chosen as it allows for the numerical reproduction of the same fluxes observed in the experimental campaign reported in [18].




FIG. 4: ACA model predictions of low-pressure oxygen reaction efficiencies.




FIG. 5: ACA model predictions of low-pressure nitrogen reaction efficiencies.

2. Stagnation line model

The VKI in-house software STAGLINE [25] is used to reproduce numerically the subsonic VKI Plasmatron experimental campaigns [2, 21, 22], specifically the reacting BL developing in front of the sample. According to the application treated in this manuscript, STAGLINE solves the multi-component Navier-Stokes (NS) equations required to describe a mix of n_s chemically reacting species in thermochemical non-equilibrium:

$$\frac{\partial \rho_i}{\partial t} + \nabla \cdot (\rho_i \mathbf{u} + \mathbf{j}_i) = \dot{\omega}_i, \quad \forall i \in [1, n_s], \quad (20)$$

$$\frac{\partial(\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + p \bar{\mathbf{I}} + \bar{\boldsymbol{\tau}}) = \mathbf{0}, \quad (21)$$

$$\frac{\partial(\rho E)}{\partial t} + \nabla \cdot (\rho \mathbf{u} H + \bar{\boldsymbol{\tau}} \cdot \mathbf{u} + \mathbf{q}) = 0. \quad (22)$$

Eq. 20 represents the mass conservation, and the corresponding terms were already defined in Section II. FRC is employed to compute $\dot{\omega}_i$ for species i . Concerning Eq. 21, p is the thermodynamic pressure of the mixture and $\bar{\boldsymbol{\tau}}$ refers to the viscous stress tensor. Looking at Eq. 22, E stands for the total energy and H is the total enthalpy. Neglecting the radiative component, \mathbf{q} is the total heat flux composed of both a conductive and a diffusive contribution. In order to close the NS equations, the computation of the thermodynamic properties of each species is required. For a thermally perfect gas, this can be achieved using the NASA polynomials [26]. The transport properties, instead, are computed from the kinetic theory of gases using a Chapman-Enskog perturbative solution method to the Boltzmann equation [27].

STAGLINE consists of a quasi-1D material response CFD solver for the reproduction of the flow field along the

stagnation line of a spherical body under the assumptions of steady, laminar, and axisymmetric flow. Thus, Munafò reduced the dimension of the problem by adopting the Dimensionally Reduced Navier-Stokes Equations formulation [28]. The spatial discretisation of these simplified conservation equations is based on the cell-centered Finite-Volume method. While, an implicit Backward-Euler method is employed for the time integration. The Roe scheme is used to compute the numerical fluxes [29].

STAGLINE is coupled with the MUTATION++ library [30] for the computation of chemical rates, thermodynamic and transport properties, and GSI balances.

The solver requires as input the temperature and velocity at the BL edge (T_e and u_e), together with the surface pressure and temperature (p_s and T_w), and provides the stagnation point mass blowing rate (\dot{m}_b). Since the surface temperature is known from the experiments, we assume thermal equilibrium at the wall at a single temperature. In this way, the ablative response of the material can be modelled by solving only the SMB for each species, already presented in Eq. 6. The computation of the chemical source term $\dot{\omega}_i$ appearing in the equation is done through MUTATION++ coupled with the ACA model.

B. Uncertainty Quantification

This section discusses the formulas for the UQ techniques used in this work. The Bayesian approach employed to solve the stochastic inverse problem is provided. Numerous numerical solutions need to be computed for the inverse problem. Surrogate models offer an efficient means of performing this task when the original models are too expensive. Their mathematical formulation is presented at the end of this section, with a specific emphasis on the Artificial Neural Network (ANN) [10] and the Gaussian Process (GP) Kriging methods [9, 31].

The UQ analyses undertaken in this manuscript utilise the UQLab software [32] (Uncertainty Quantification in MATLAB). The modules utilised include those for Kriging training [31] and Bayesian inversion [33]. The ANNs are trained using the PyTorch package in Python [34]. The optimisation of the corresponding hyperparameters is done through Optuna [35].

1. Stochastic inverse problem

The problem is called inverse because the goal is to send information about observations backwards to understand the model inputs. Bayesian frameworks, in-

deed, enable the mapping of uncertainty in the output of a model back to the uncertainties in the input of the model itself. This is different from the usual forward method, which consists of the determination of the uncertainty in the model’s response based on the stochastic nature of the model’s input.

To solve the Bayesian inverse problem, three major components are required: a computational forward model M , a set of input parameters \mathbf{x} to be inferred, and N_{obs} auxiliary experimental data, \mathbf{y}_{obs} . Bayes’ theorem states that the posterior distribution of \mathbf{x} is:

$$\pi(\mathbf{x}|\mathbf{y}_{\text{obs}}) = \frac{\pi(\mathbf{x})\pi(\mathbf{y}_{\text{obs}}|\mathbf{x})}{\int \pi(\mathbf{x})\pi(\mathbf{y}_{\text{obs}}|\mathbf{x}) d\mathbf{x}}, \quad (23)$$

where $\pi(\mathbf{x})$ represents the prior distribution of \mathbf{x} , and $\pi(\mathbf{y}_{\text{obs}}|\mathbf{x})$ denotes the likelihood of an observable quantity \mathbf{y}_{obs} given \mathbf{x} , and the integral at the denominator represents the marginal likelihood or evidence, ensuring the distribution integrates to 1. The first term expresses prior knowledge of \mathbf{x} before performing the inversion. It can result from previous calibration, tests, or expert assessment. When there are only a few observations (a case where the prior distribution has a dominant influence), it is preferable to specify a uniformly distributed non-informative prior to make sure that the likelihood has the greatest influence on the posterior. Log-uniform priors are a better option when there is a wide range of uncertainty about the parameters that spans multiple orders of magnitude [12]. When one strongly believes in a certain distribution, they can use informative priors, such as Gaussian experimental uncertainty, to provide specific information. The second term quantifies the probability of obtaining the outcome \mathbf{y} based on the input \mathbf{x} using a computational model $M(\mathbf{x})$.

As all models are inherently simplified representations of the real world [33], it is necessary to incorporate a discrepancy term,

$$\epsilon = \mathbf{y}_{\text{obs}} - M(\mathbf{x}), \quad (24)$$

where ϵ comprises both the model inaccuracy and the measurement noise. For the sake of simplicity, it is considered an additive Gaussian discrepancy with a zero mean value. Assuming, moreover, that the model error is very small compared to the experimental uncertainty, σ_i , on the measurement $y_{\text{obs},i}$, that the experimental uncertainty is known and normally distributed around the mean ($N(0, \sigma_i^2)$), and independence between N_{obs} , the likelihood reads:

$$\pi(\mathbf{y}_{\text{obs}}|\mathbf{x}) = \prod_{i=1}^{N_{\text{obs}}} N(y_{\text{obs},i}|M_i(\mathbf{x}), \sigma_i^2). \quad (25)$$

Generally, the evidence in Eq. 23 does not have a closed-form solution. Thus, the posterior distribution is obtained employing sampling methods, such as the Markov Chain Monte Carlo (MCMC) method [36]. The objective of an MCMC algorithm is to generate a series of points approximating the desired posterior distribution. In this work, since the parameters to infer show strong correlation between each other, the target posterior distribution is estimated using the MCMC method along with the Affine Invariant Ensemble algorithm [33]. It is based on proposing and subsequently accepting or rejecting candidate points. The algorithm is designed to address the challenge of exploring the state space effectively while ensuring convergence toward the desired distribution.

2. Gaussian Process Kriging

The entire process outlined above necessitates millions of calls of both the 0D reactor and STAGLINE. Even though the latter is a mid-fidelity CFD solver, it is still inefficient to be employed for this specific objective. Commonly, surrogate models are utilised, and one option that is widely employed is GP Kriging [31].

A Kriging model is constructed for each observation y of interest, with the latter being provided by a model $M(\mathbf{x})$ as a function of input quantity \mathbf{x} . In the field of metamodeling, given a set of n_t training points denoted as $\mathbf{X} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n_t)})$, along with their corresponding model realisations, $Y = (y^{(1)}, \dots, y^{(n_t)})$, the surrogate model is trained using these pairs. The goal is to use this model to predict the response for a new point, $\mathbf{x}^{(*)}$, at a lower cost. The equation below describes a Kriging metamodel:

$$M(\mathbf{x}) \approx M^K(\mathbf{x}) = \beta^T \mathbf{f}(\mathbf{x}) + \sigma^2 Z(\mathbf{x}, \omega). \quad (26)$$

The first term in Eq. 26, $\beta^T \mathbf{f}(\mathbf{x})$, represents the mean value (or trend) of the process. It captures the overall behaviour or systematic variation in the data, allowing the surrogate model to account for the general trend present in the observed responses. This term is characterised by the regression coefficients, β , to be found and the basis functions, $\mathbf{f}(\mathbf{x})$. The latter could be a constant, polynomials, trigonometric functions, splines, or other types of functions that are appropriate for the problem at hand. The second term, $\sigma^2 Z(\mathbf{x}, \omega)$, accounts for the fluctuations of the model’s trend from the real response. The symbol $Z(\mathbf{x}, \omega)$ stands for a Gaussian process with a zero-mean and a unit variance that is stationary and has a constant variance of σ^2 [31]. ω is used to denote the underlying probability space, which is characterised by the correlation function (or correlation family), $R = R(\mathbf{x}, \mathbf{x}', \theta)$. It

describes the correlation between two sample points in the output space which depends on \mathbf{x} , \mathbf{x}' , and the hyperparameters θ . In other words, R indicates the similarity between function values at two distinct sample points in the input space. Parameters σ^2 and θ also need to be determined.

To compute the mean and variance of the Kriging prediction, \hat{M}^K , for a new point $\mathbf{x}^{(*)}$, given noise-free responses and the observed data \mathbf{X} and \mathbf{Y} , follow these steps:

$$\mu_{\hat{M}^K}(\mathbf{x}^{(*)}) = \mathbf{f}(\mathbf{x}^{(*)})^T \hat{\beta} + \mathbf{r}(\mathbf{x}^{(*)})^T \mathbf{R}^{-1}(\mathbf{Y} - \mathbf{F}\hat{\beta}), \quad (27)$$

$$\sigma_{\hat{M}^K}(\mathbf{x}^{(*)}) = \sigma^2(1 - \mathbf{r}^T(\mathbf{x}^{(*)})\mathbf{R}^{-1}\mathbf{r}(\mathbf{x}^{(*)}) + \mathbf{u}^T(\mathbf{x}^{(*)})(\mathbf{F}^T\mathbf{R}^{-1}\mathbf{F})^{-1}\mathbf{u}(\mathbf{x}^{(*)})), \quad (28)$$

where \mathbf{F} represents the matrix of trends calculated at the training points and $\mathbf{r}(\mathbf{x}^{(*)})$ denotes the vector of cross-correlations between the prediction point $\mathbf{x}^{(*)}$ and each individual observation. The matrix \mathbf{R} is the correlation matrix of size $n_t \times n_t$, representing the correlations between all the training points. As seen in Eq. 27, the computation of the Kriging predictor necessitates the inversion of a matrix with dimensions $n_t \times n_t$. This represents one of the bottlenecks of the Kriging model. Moreover, it approximates well non-linear functions as long as the dimension of the input space is rather low. For this reason, we decided to compare the performances of the GP Kriging with ANNs on the same training points, given their ability to handle high-dimensional space well.

In this work, we employ an ordinary Kriging. The trend is characterised by a constant yet unknown value denoted as $\beta^T \mathbf{f}(\mathbf{x}) = \beta_0 f_0(\mathbf{x}) = \beta_0$, where it is conventionally defined that $f_0(\mathbf{x}) = 1$. Examining Eq. 27, it can be shown that for ordinary Kriging, the functions simplify to $\mathbf{f} = 1$ and $\mathbf{F} = 1$. The Matern52 family of kernels is used in this study, leading to a GP with relatively smooth behaviour. Since the input dimension is higher than one, a multi-dimensional ellipsoidal anisotropic correlation function is employed. In this case, a single length scale for each dimension, θ_k , is defined as the response's behaviour cannot be assumed to be uniform across all dimensions. To conclude, the estimation of the most suitable hyperparameters β , σ^2 , and θ_k is done by solving an optimisation problem. Maximum-Likelihood is adopted as estimation method. For more details refer to [31].

3. Artificial Neural Network

The MultiLayer Perceptron (MLP) is a commonly used neural network. MLP is composed of multiple layers, including an input layer, one or more hidden layers, and an

output layer, where each layer contains a set of perception elements known as neurones. A graphical representation of the MLP topology is shown in Fig. 6.

A detailed description of ANN structures is provided in [10]. Generally, the input layer contains n input variables $\mathbf{X} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)})$ and the output layer the corresponding m realisations $\mathbf{Y} = (\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(m)})$. Given a generic layer l , the input layer is denoted by $l = 1$ and the output layer by $l = L$, where L is the total number of layers in the network. Each layer is composed of multiple neurons, and the network is structured such that every neuron in layer l is connected to every neuron in layer $l + 1$. This type of architecture is referred to as a fully connected network. These connections are characterised by a set of weights and biases: for each neuron i in layer l , the bias b_i^l adjusts the neuron's output, providing additional flexibility during training to improve the overall fit of the model. The weights $w_{i,j}^l$, on the other hand, quantify how much the output of neuron i in layer l contributes to the activation of neuron j in layer $l + 1$. The number of neurones in the l -th layer is defined as s^l .




FIG. 6: Example of MultiLayer Perceptron topology [37].

A non-linearity is applied to the output of each neurone, typically in hidden layers, making the MLP capable of learning complex patterns in the data. This is done by non-linear activation function, defined as σ^l , like ReLU, Tanh, or ELU. Defining the output of neuron i -th of the layer l -th as y_i^l , the predictions can be computed as follows:

$$y_j^{l+1}(x) = \sigma^{l+1} \left(\sum_{i=1}^{s^l} w_{j,i}^l y_i^l(x) + b_j^{l+1} \right) \quad \forall i, \quad (29)$$

leading to the matrix form:

$$\mathbf{y}^{l+1}(\mathbf{x}) = \sigma^{l+1}(\mathbf{W}^{l+1}\mathbf{y}^l(\mathbf{x}) + \mathbf{b}^{l+1}). \quad (30)$$

In this work, linear activation functions are chosen for both the input and output layers of the network. In this

way, at the input layer we prevent the data from being distorted before being processed, and at the output layer we allow for continuous physical output values that are easier to interpret. It is clear that the remaining hyperparameters to be determined are the number of hidden layers, the number of neurones per layer, and the type of non-linear activation function for the hidden layers. As in GP, also in ANN an optimisation is necessary for this objective. In the analysis, we also include the learning rate, which controls how much the model’s weights are updated during training. It determines the step size taken in the direction of the gradient when minimising the loss function.

To determine the optimal hyperparameters, a general objective function is defined within the Optuna framework [35]. This objective function involves training the ANN and computing the Normalized Root Mean Square Error (NRMSE) on a separate validation set. The set of hyperparameters that yields the lowest NRMSE is retained. For the training process itself, the Mean Squared Error is used as the cost function, and the model parameters are optimized using the Adam algorithm [38]. The resulting ANN architecture consists of 3 hidden layers, each with 32 neurons and ELU as activation function.

To conclude, an ensemble of ANNs is employed in the course of the work to improve prediction accuracy and mitigate overfitting. Fixed the architecture, 50 distinct ANNs were trained using the procedure described above, each on a different bootstrap sample of the training data. The final prediction of the ensemble is obtained by averaging the predictions of all individual ANNs. The performances of the resulting ANN ensemble will be compared with the GP Kriging in the next section.

V. RESULTS

A. Surrogate models construction

The primary objective of this project is to calibrate the finite-rate nitridation model by expanding the experimental dataset originally used by Capriati [15], incorporating in particular the recent high-pressure nitrogen campaign conducted at subsonic conditions [21]. To this end, only one new surrogate model is required, tailored specifically to the pressure range of the new experimental campaign, namely $p_s = 100 - 200$ hPa. The training data for $p_s = 15$ hPa were already available from the earlier work by Capriati.

Both the GP Kriging and MLP surrogate models are trained on a comprehensive dataset of flow simulations provided by STAGLINE. The objective is to significantly

reduce computational cost and enable efficient evaluation of the mass blowing rate during the MCMC sampling. A total of 4000 simulations are used to train these surrogates, with training points sampled using the Latin Hypercube Sampling (LHS) technique [39] to ensure good coverage of the input space.

As detailed in Section IV A 2, the set of input variables required for each simulation includes T_e , p_s , u_e , T_w , and the GSI model. By applying Bernoulli’s equation with the viscous correction proposed by Homann [40], the velocity u_e can be linked to the dynamic pressure p_d . Finally, the reaction parameters to be inferred are those identified in Table II. Therefore,

$$\mathbf{x} = [T_e, p_s, p_d, T_w, \log_{10} X_{12}, \log_{10} X_{13}, \log_{10} X_{14}, \log_{10} X_{15}, E_{a,10}, E_{a,12}, E_{a,13}, E_{a,14}, E_{a,15}, \log_{10} B] \quad (31)$$

is the set of parameters on which the probabilistic input space is built. The latter is chosen to be sufficiently broad to include all the experimental operating conditions in Tab. IV and to encompass our previous beliefs about the model parameters in order to match the experimental observations. In Tab. VI the interval for each input is specified.

Parameter	Min	Max	Parameter	Min	Max
T_e [K]	8000	12000	$\log_{10} X_{15}$	5	12
p_s [Pa]	9000	21000	$E_{a,10}$ [K]	0	3000
p_d [Pa]	30	75	$E_{a,12}$ [K]	0	10000
T_w [K]	2350	2750	$E_{a,13}$ [K]	0	15000
$\log_{10} X_{12}$	-4	2	$E_{a,14}$ [K]	0	35000
$\log_{10} X_{13}$	-4	0	$E_{a,15}$ [K]	5	12
$\log_{10} X_{14}$	-4	0	$\log_{10} B$	-7	-3

TABLE VI: Input space used to train the surrogate models. Non-informative uniform prior distribution.

Non-informative uniform priors are chosen for all the quantities. The pre-exponential factors, X_i , and the total active site density, B , are not well known, and they have been expressed as \log_{10} because our ignorance of them spans different order of magnitude. Note that, before training the model, the inputs are normalised and the $\log_{10}(\dot{m}_b)$ is considered to make the training more robust.

Each time a surrogate model is trained, the final step is to assess its predictive quality. This is done by comparing the outputs of the surrogate model with the reference values computed by STAGLINE on an independent validation set. For both the Kriging and ANN models, an additional 150 validation points are generated using the LHS technique. The resulting normalised error, defined

as:

$$NRMSE = \sqrt{\frac{1}{N_v} \sum_{i=0}^{N_v} \frac{(Y_{v,i} - \hat{Y}_i)^2}{Y_{v,i}^2}} \cdot 100, \quad (32)$$

is computed. Quantity N_v denotes the number of validation points, and $Y_{v,i}$ and \hat{Y}_i represent the responses of the CFD solver and the surrogate model, respectively, evaluated at the same input point.

Figure 7 shows the comparison of the validation error in terms of NRMSE according to the number of training points used in the training. The results indicate that the ANN achieves superior accuracy, with an NRMSE of 1.2%, compared to 5% for the Kriging model. As mentioned in the previous section, Kriging struggles to accurately reproduce the behavior of the full numerical solver when operating in a high-dimensional input space, in this case 14 input parameters.

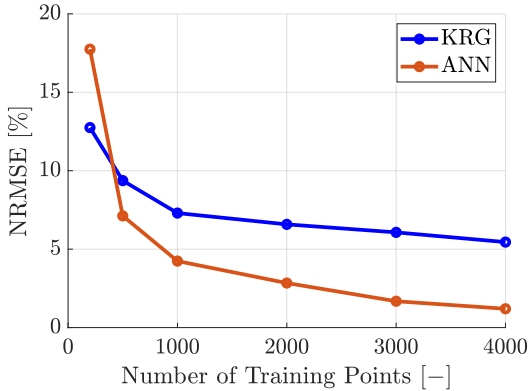


FIG. 7: Comparison of the NRMSE between the Kriging and ANN, based on a 14-dimensional input space.




FIG. 8: QQplot of the mass blowing rate: CFD response VS surrogate model response. Kriging values in blue. ANN in orange. CFD response on the validation points in black.

The QQplot shown in Fig. 8 further supports the conclusions from the NRMSE analysis. In this plot, the surrogate model predictions are plotted against the corresponding CFD solver responses on the validation set. While the Kriging model exhibits noticeable deviations from the reference values, the ANN predictions align closely with the 45-degree line, indicating a high level of agreement with the CFD outputs.

B. Finite-Rate Nitrogen Model Calibration

To increase the robustness of finite-rate nitrogen model predictions, the Bayesian framework is employed. According to Section IV B 1, to characterise the atomic nitrogen reactions presented in Tab. II, three key components are required: a set of prior distributions to be inferred, supporting experimental data accompanied by their associated uncertainties, and a forward numerical model. Our experimental observations consist of:

$$\mathbf{y}_{MB} = [\gamma_{CN}^{\text{exp}}, \gamma_{N_2}^{\text{exp}}] \quad \text{and} \quad \mathbf{y}_{PL} = [\dot{m}_b^{\text{exp}}], \quad (33)$$

where \mathbf{y}_{MB} represents the data set from MB experiments carried out by Murray et al. [18] and \mathbf{y}_{PL} refers to the observations at high-pressure from Helber et al. [2] and [21]. As previously mentioned, the MB experiments are modelled using the 0D reactor proposed by Prata et al. [17], while the VKI Plasmatron experiments are described using surrogate models trained on CFD simulations. Specifically, the Plasmatron campaign at $p_s = 15$ hPa is modelled using the Kriging surrogate developed by Capriati et al. [15]. For the higher-pressure campaign at $p_s = 100$ – 200 hPa, both the Kriging and ANN models presented in Section V A are employed.

It is important to note that although the last two surrogate models are trained on a 14-dimensional input space, the goal of the calibration is to infer only the reaction model parameters. The values of p_s , p_d , and T_w are taken directly from the experimental campaign and are assumed to be deterministic, without associated uncertainties. However, since T_e was not measured during the campaign reported in [21], it is treated as an additional unknown parameter in the inference process. The corresponding posterior distributions are not shown in the manuscript, as the results of the analysis did not yield significant insights as demonstrated by Capriati [4].

We performed two different Bayesian inferences, one using the Kriging surrogate model and the other using ANN. The objective was to evaluate how differences in surrogate model accuracy affect the resulting posterior distributions of the reaction parameters. As emphasised in Fig. 3, validating the updated nitrogen model is cru-

cial to assess its robustness. To this end, only half of the experimental conditions listed in Tab. III were used for the calibration. The remaining conditions were retained exclusively for model validation. Specifically, the configurations G1_100_400, G1_200_370, G8_100_370, and G8_200_400 were selected as the new calibration points.

Due to the large number of parameters involved in the calibration process, this manuscript shows solely the posterior distributions of the gas-phase dependent nitridation reaction. This specific reaction, identified as reaction (15) in Tab. II, was introduced to improve model accuracy under high-pressure conditions. The pre-exponential factor and the activation energy are shown, respectively, in Figs. 9 and 10. In each figure, we also show the deterministic value proposed by Prata et al. [17], as well as the posterior distribution previously obtained by Capriati et al. [15], providing a direct comparison with the results of the current study.

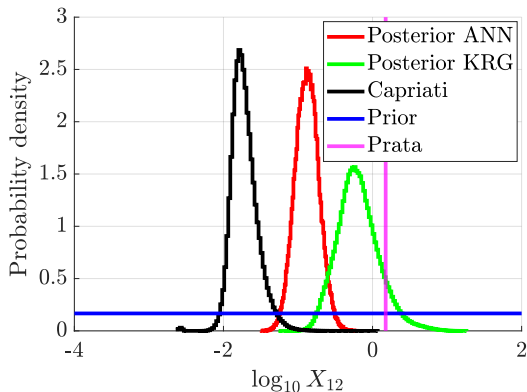


FIG. 9: Pre-exponential factor of the gas-dependent nitridation reaction: $N+N(s)+C(b) \rightarrow CN+N+(s)$. Posterior distribution obtained using ANN in red. Posterior distribution obtained using Kriging in green. Capriati's posterior distribution in black. Prior distributions in blue. Value from Prata et al. in magenta.

The first notable observation is that, although the Kriging and ANN surrogate models differ in prediction accuracy by less than 4%, their impact on the resulting posterior distributions is far from negligible. Not only do the most probable values shift between the two models, but the posterior distribution obtained with Kriging exhibits a greater uncertainty. This is reflected in the larger area under the Kriging-based posterior curve. Given these differences and the higher accuracy of the ANN model, it is reasonable to consider the ANN-based posterior as the reference solution.

Capriati et al. observed that, when combining MB and Plasmatron experimental data, the inferred activa-

tion energy for the gas-phase nitridation reaction deviated from the value originally proposed by Prata et al. This deviation suggested a potential incompatibility in the model's ability to reproduce the MB nitridation efficiencies, particularly at low-temperatures [15]. This time, incorporating additional measurements at higher pressures, the new posterior distributions tend towards Prata's deterministic values, but still do not agree with the original value. This trend holds consistently across all three reaction parameters analysed in this study.

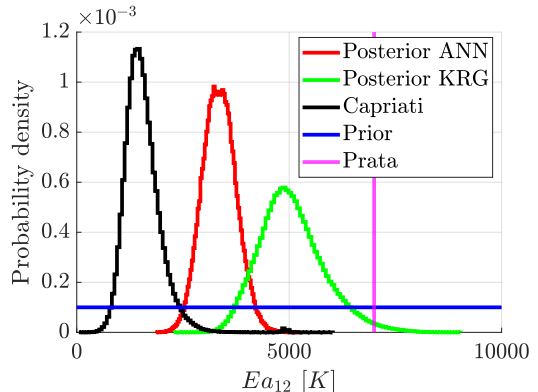


FIG. 10: Activation energy of the gas-dependent nitridation reaction: $N+N(s)+C(b) \rightarrow CN+N+(s)$. Posterior distribution obtained using ANN in red. Posterior distribution obtained using Kriging in green. Capriati's posterior distribution in black. Prior distributions in blue. Value from Prata et al. in magenta.




FIG. 11: Total active site density. Posterior distribution obtained using ANN in red. Posterior distribution obtained using Kriging in green. Capriati's posterior distribution in black. Prior distributions in blue. Value from Prata et al. in magenta.

One of the main results of this calibration is the complete characterisation of the total active site density,

shown in Fig. 17. Unlike the results reported by Capriati et al., where the posterior distribution of parameter B tended toward the lower bound of the prior, spanning a relatively broad region, the current calibration shows an increase in B by approximately one order of magnitude, bringing it closer to the deterministic value proposed by Prata et al. The Kriging-based posterior remains in disagreement with the ANN-based curve.

To conclude, the two Bayesian inversions are performed by considering a total of 50 MCMC chains, each characterised by $3 \cdot 10^5$ samples. The first 20% of post-samples generated by each chain is discarded as Burn-in. Gelman-Rubin diagnostic is employed to assess the convergence in all chains. For more details, the interested reader can refer to [33].

The posterior distributions of the remaining model parameters are provided in the Appendix, which also includes an additional calibration performed using only the high-pressure measurements from the VKI Plasmatron facility. The results clearly indicate that including the MB experiments is essential to obtain reliable and physically consistent predictions.

Posterior distribution propagation and Validation

The posterior distributions obtained from the calibration of the model were propagated through both the 0D reactor model and its ANN-surrogate counterpart. This step aimed to analyse the behaviour of the new calibrated finite-rate nitrogen model at both low- and high-pressure regimes. In addition to assessing the model’s predictive capability, the propagation also served to evaluate whether the calibration points were informative enough.

The low-pressure analysis enabled us to evaluate whether the calibrated model offers any improvement over Prata’s original formulation in reproducing the MB experimental from Murray et al. [18]. The corresponding posterior predictions are shown in Figs. 12 and 13.

The results clearly indicate that the finite-rate nitrogen model calibrated in this work outperforms the original deterministic model in capturing both nitridation and recombination phenomena at low pressure. The confidence intervals of the predictions result narrow or wide in agreement with the associated experimental uncertainties. Note that, in the case of nitrogen recombination, the model uncertainty retrieved from the posterior propagation is smaller than the experimental uncertainty, suggesting a high level of confidence in model’s predictions for that regime.

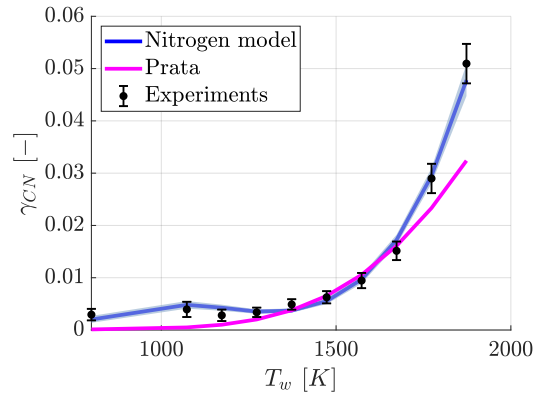


FIG. 12: Nitridation reaction efficiency predicted by the calibrated model. The blue solid line represents the mean of the prediction. The shadow area refers to the confidence interval (95%) around the mean. Prata’s prediction in magenta. Experimental calibration points from Murray et al. in black.

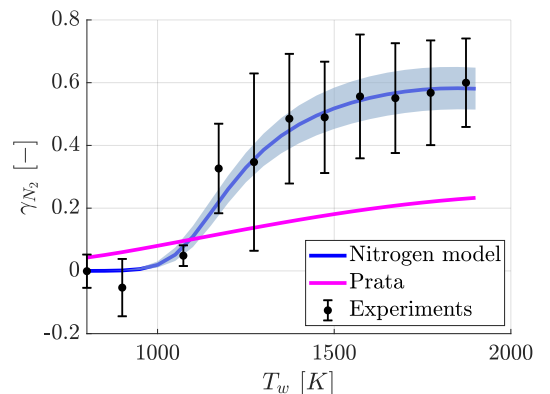


FIG. 13: Nitrogen recombination reaction efficiency predicted by the calibrated model. The blue solid line represents the mean of the prediction. The shadow area refers to the confidence interval (95%) around the mean. Prata’s prediction in magenta. Experimental calibration points from Murray et al. in black.

To conclude the analysis of the finite-rate nitrogen model before proceeding to the calibration of the extended air-carbon GSI model, we focus on the validation of the model itself. Instead of propagating the posterior through the high-pressure calibration points, we use unseen experimental data to assess the model’s predictive capabilities at higher pressures. This approach allows us to assess whether the information extracted during the calibration process is sufficient to describe new conditions and to verify the robustness of the inferred model. In particular, the validation is carried out on those points from Tab. IV that were excluded from the calibration set,

namely G1_100_370, G1_200_400, G8_100_400, and G8_200_370. The results of the validation process are shown in Fig. 14.

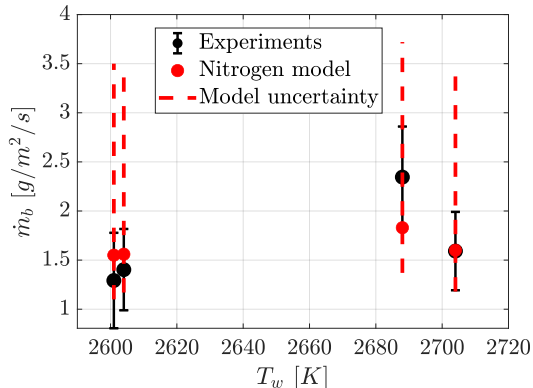


FIG. 14: Posterior prediction of the stagnation mass blowing rate on the validation set. Red markers represent the most probable predictions. Red dashed lines indicate the model uncertainty. Black markers refer to the experimental data with their uncertainty.

What is shown in the figure is the result of the posterior propagation, along with the entire support on the edge temperature T_e , which was not measured during the experiments. The model predictions are represented by the most probable value and the corresponding uncertainty interval. Note that, since the relationship between the model inputs and outputs is nonlinear, and the propagated posterior distributions are not necessarily Gaussian, the resulting model predictions will also, in general, not follow a Gaussian distribution.

Although the predictions do not fully capture all the experimental observations, they still provide a good approximation. Both the third and fourth conditions are accurately predicted, while for the other cases, the experimental value lies only partially within the model's uncertainty range. However, the most probable predictions remain in very good agreement with the measured data. Finally, the predicted trend of the stagnation mass blowing rate as a function of surface temperature shows a less pronounced increase than that observed experimentally, as already noticed by Capriati et al. [15]. This discrepancy could be attributed either to a possible model inadequacy or to a measurement bias in the experiments.

C. Air-Cabon Model Calibration

In this section we present our preliminary results of the ACA model calibration. Since atomic nitrogen and oxygen coexist in the post-shock layer, a joint calibration

of the corresponding reaction rates is essential to ensure a more physically accurate GSI model.

This time, we refer to the reactions listed in both Tab. I and II. Therefore, the set of model parameters to infer becomes:

$$\mathbf{x} = [\mathbf{x}_O + \mathbf{x}_N] \quad (34)$$

where,

$$\begin{aligned} \mathbf{x}_O = & [\log_{10} X_1, \log_{10} X_3, \log_{10} X_5, \log_{10} X_7, \\ & \log_{10} X_8, \log_{10} X_9, E_{a,2}, E_{a,3}, E_{a,4}, \\ & E_{a,6}, E_{a,7}, E_{a,8}, E_{a,9}], \end{aligned} \quad (35)$$

and

$$\begin{aligned} \mathbf{x}_N = & [\log_{10} X_{12}, \log_{10} X_{13}, \log_{10} X_{14}, \\ & \log_{10} X_{15}, E_{a,10}, E_{a,12}, E_{a,13}, \\ & E_{a,14}, E_{a,15}, \log_{10} B]. \end{aligned} \quad (36)$$

Note that, $\log_{10} B$ is included in \mathbf{x}_N only for consistency with the previous calibration procedure. However, the total active site density B is inherently a material property. Both nitrogen and oxygen atoms compete for the same free reactive sites on the surface, meaning that B plays a role in both reaction mechanisms.

The set of experimental data is enlarged as follows:

$$\begin{aligned} \mathbf{y}_{MB,N} &= [\gamma_{CN}^{\text{exp}}, \gamma_{N_2}^{\text{exp}}], \\ \mathbf{y}_{MB,O} &= [\gamma_{CO}^{\text{exp}}, \gamma_{CO_2}^{\text{exp}}, \gamma_{O_2}^{\text{exp}}], \\ \mathbf{y}_{PL} &= [\dot{m}_b^{\text{exp}}]. \end{aligned} \quad (37)$$

The distinction between nitrogen and oxygen has been made for clarity. Compared to the previous calibration, this time we also employ the oxygen reaction efficiencies from Murray et al. [18]. No experimental data at high pressure from Tab. V will be used as calibration points, they will be employed as unseen points to validate the ACA model in future works. Therefore, there is no need to train a new surrogate model on STAGLINE simulations coupled with the ACA model. The 0D reactor is enough to describe numerically the experimental campaign with atomic oxygen. However, all models used in the calibration of the nitrogen model must also be used.

By definition, the prior distributions correspond to our prior beliefs about the parameters we want to infer. While for the nitrogen model parameters we used the same priors as Capriati, now it was necessary to carry out a preliminary study for the oxygen ones. We opted for a preliminary calibration considering only \mathbf{x}_O and $\mathbf{y}_{MB,O}$,

maintaining a very large support around Prata’s deterministic value. As a result, the prior distributions assigned to each new parameter are reported in Tab. VII. The very wide bounds, e.g., the pre-exponential factor and activation energy of reaction (8), reflect the fact that the preliminary calibration did not provide sufficient information to constrain these parameters. These wide priors were kept intentionally large, with the expectation that additional insights might be gained in future calibration incorporating experimental data from the air Plasmatron campaign [22].

Parameter	Min	Max	Parameter	Min	Max
$\log_{10} X_1$	-1	1	$E_{a,2}$ [K]	25000	100000
$\log_{10} X_3$	1	6	$E_{a,3}$ [K]	0	15000
$\log_{10} X_5$	-2	1	$E_{a,4}$ [K]	0	5000
$\log_{10} X_7$	0	7	$E_{a,6}$ [K]	15000	150000
$\log_{10} X_8$	-7	4	$E_{a,7}$ [K]	0	10000
$\log_{10} X_9$	-6	4	$E_{a,8}$ [K]	0	100000
			$E_{a,9}$ [K]	0	35000

TABLE VII: Non-informative uniform prior distribution for atomic oxygen reactions of the ACA model.

The parameters associated with high-pressure oxygen conditions remain poorly identified even after the calibration combining oxygen and nitrogen, as the current calibration relied exclusively on low-pressure oxygen measurements, which lack the sensitivity needed to constrain them. The activation energy of reaction (7) is shown in Fig. 15 as proof.

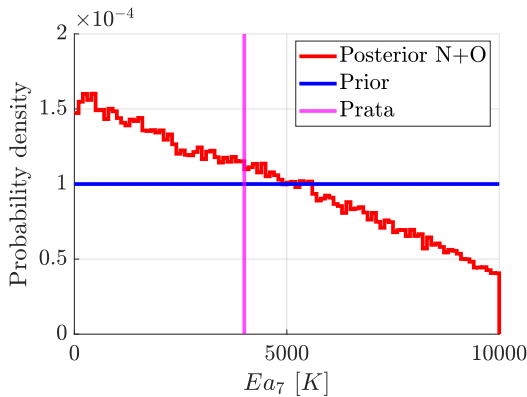


FIG. 15: Activation energy of oxidation reaction at high pressures: $O+O^*(s)+C(b) \rightarrow CO+O+(s)$. Posterior distribution combining nitrogen and oxygen experiments in red. Prior distribution in blue. Value from Prata et al. in magenta.

On the contrary, the reactions governing low-pressure conditions are well characterised. Combined with a well-

identified total active site density, this results in an ACA model that accurately reproduces the experimental data of Murray et al. [18], as shown below. As an example, the pre-exponential factor of the recombination reaction (9) through O(s) is shown in Fig. 16.




FIG. 16: Pre-exponential factor of gas-independent recombination reaction through O(s): $O(s)+O(s) \rightarrow O_2 + 2(s)$. Posterior distribution combining nitrogen and oxygen experiments in red. Prior distribution in blue. Value from Prata et al. in magenta.

More interestingly, the posterior distribution of B obtained by combining nitrogen measurements at both low and high pressures, shown in Fig. 17, exhibits the same trend as the one obtained when oxygen data are also included. This clearly indicates that, in this context, the low-pressure oxygen measurements do not contribute meaningful information to the inference of B .

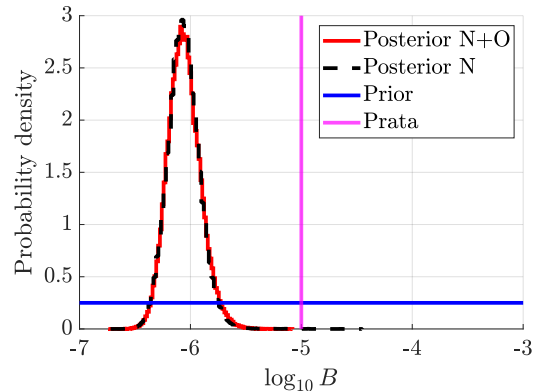


FIG. 17: Total active site density. Posterior distribution combining nitrogen and oxygen experiments in red. Posterior distribution obtained employing only nitrogen experiments in black. Prior distribution in blue. Value from Prata et al. in magenta.

Posterior distribution propagation

The posterior distributions obtained from the calibration were propagated only through the 0D reactor.

As for the finite-rate nitrogen model, the calibrated ACA model offers an improvement over the deterministic model proposed by Prata et al. in reproducing the oxygen low-pressure experiments from Murray et al. [18].

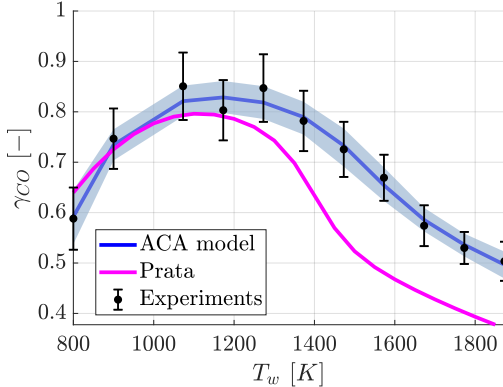


FIG. 18: Oxidation reaction efficiency predicted by the calibrated ACA model. The blue solid line represents the mean of the prediction. The shadow area refers to the confidence interval (95%) around the mean. Prata's prediction in magenta. Experimental calibration points from Murray et al. in black.




FIG. 19: Oxygen recombination efficiency predicted by the calibrated ACA model. The blue solid line represents the mean of the prediction. The shadow area refers to the confidence interval (95%) around the mean. Prata's prediction in magenta. Experimental calibration points from Murray et al. in black.

The results shown in Fig. 18, 19, and 20 are very promising. For the first time, the reaction rates of atomic

oxygen have been successfully calibrated, at least under low-pressure conditions.

To conclude, as in the case of B , when oxygen and nitrogen are combined, no significant improvement is observed in the posterior distributions of the nitrogen model parameters. This indicates that, even though the parameters are highly correlated, the measurements $\mathbf{y}_{MB,N}$ and \mathbf{y}_{PL} already provided sufficient information. In fact, propagating the updated posterior distribution through the ANN surrogate model for the high-pressure nitrogen campaign yielded only marginal gains. The remaining posterior distributions are reported in the Appendix.

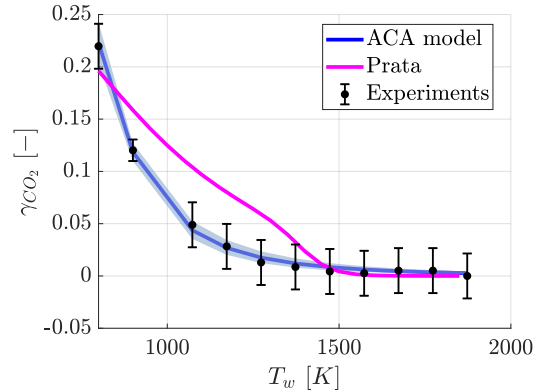


FIG. 20: Carbon dioxide reaction efficiency predicted by the calibrated ACA model. The blue solid line represents the mean of the prediction. The shadow area refers to the confidence interval (95%) around the mean. Prata's prediction in magenta. Experimental calibration points from Murray et al. in black.

VI. CONCLUSIONS

This work showed a robust methodology to improve the predictivity of the ACA model by calibrating its reaction rates through Bayesian inference. The proposed strategy combines low-pressure MB and high-pressure Plasmatron experiments, along with dedicated numerical models capable of reproducing the corresponding experimental campaign. MB are described by a 0D reactor, while Plasmatron experiments are described by a CFD solver coupled with the GSI model.

Two different calibrations were performed. The first focused solely on a subset of nitrogen chemical reactions, and the resulting finite-rate nitrogen model predictions show a good agreement with the counterpart experiments on both calibration and validation sets. In fact, we have not only enhanced the robustness of the nitrogen model under high-pressure conditions but also validated its predictive capability against experimental data not included

in the calibration process. In this way, the limitations encountered by Bandera when attempting to validate Capriati’s model were overcome. However, the predicted stagnation mass blowing rate exhibits a less pronounced increase with surface temperature compared to what is observed experimentally. This discrepancy could be either a potential inadequacy in the model or a bias in the experimental measurements. By model inadequacy, we refer to an irreducible error between model predictions and observations, which persists even when the best possible parameter values are used, due to incomplete or imperfect representation of the underlying physical phenomena. Up to this point, the ACA model has been calibrated by considering only experimental uncertainty, under the assumption that the model is perfect. Nevertheless, in the Appendix we show that the model still failed to accurately capture the experimental data used for calibration. This persistent mismatch indicates the presence of model error. In the literature, calibration frameworks that incorporate model error have proven to be a promising technique for identifying both experimental biases and model inadequacy. To speed up the MCMC sampling, both ANN and Kriging were trained on a dataset of CFD simulations of the Plasmatron experimental campaigns. The ANN resulted in being more suitable than a Kriging surrogate model in a high-dimensional input space. Although the difference between ANN and Kriging in terms of accuracy is less than 4%, the corresponding posterior distributions disagree.

The second calibration marks the real novelty of this work: for the first time, the reaction rates involving atomic oxygen were incorporated into the calibration process. Since atomic nitrogen and oxygen coexist in the post-shock layer, a joint calibration of the corresponding reaction rates is essential to ensure a more physically accurate GSI model. Low- and high-pressure nitrogen data were combined with low-pressure oxygen measurements. Although high-pressure oxygen reaction rates could not be fully characterised, the calibrated ACA model outperforms Prata’s original deterministic formulation at low pressures. This is evident in both nitrogen and oxygen experimental campaigns, where the stochastic model better captures the measured reaction efficiencies.

Recently, an experimental campaign was carried out at VKI, providing the high-pressure oxygen data required to perform model validation and, if necessary, a new calibration. Since the oxygen reaction rates included in the ACA model to ensure accuracy at high pressure are not well defined, it is highly likely that a new calibration will be required. The final goal is to extend the calibration to the full ACA model, incorporating molecular oxygen reactions as well. This extension would increase the number of parameters to be inferred, thereby raising the dimensionality of the problem. Consequently, a sensitivity analysis will be essential to reduce the problem’s dimensionality before carrying out the final Bayesian inversion.

-
- [1] A. del Val, D. Luís, and O. Chazot, Experimental methodology for the accurate stochastic calibration of catalytic recombination affecting reusable spacecraft tps, *Chemical Physics* **559**, 111528 (2022).
- [2] B. Helber, A. Turchi, and T. E. Magin, Determination of active nitridation reaction efficiency of graphite in inductively coupled plasma flows, *Carbon* **125**, 582 (2017).
- [3] G. Bellas-Chatzigeorgis, *Development of advanced gas-surface interaction models for chemically reacting flows for re-entry conditions*, Ph.D. thesis, Politecnico di Milano (2018).
- [4] M. Capriati, *Multi-fidelity Bayesian inference of hypersonic flow free-stream conditions and heterogeneous chemistry model parameters*, Ph.D. thesis, Institut Polytechnique de Paris (2024).
- [5] A. Turchi, J. J. Matesanz Saiz, T. E. Magin, and O. Chazot, Duplication of hypersonic stagnation-region aerothermochemistry and gas-surface interaction in high-enthalpy ground testing, *Experiments in Fluids* **62**, 238 (2021).
- [6] A. F. Cortesi, P. M. Congedo, T. E. Magin, B. Van Hove, and O. Karatekin, Rebuilding freestream atmospheric conditions using surface pressure and heat flux data, in *8th AIAA Atmospheric and Space Environments Conference* (American Institute of Aeronautics and Astronautics, Washington, D.C., 2016).
- [7] A. F. Cortesi, P. G. Constantine, T. E. Magin, and P. M. Congedo, Forward and backward uncertainty quantification with active subspaces: Application to hypersonic flows around a cylinder, *Journal of Computational Physics* **407**, 109079 (2020).
- [8] J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*, Applied Mathematical Sciences No. volume 160 (Springer, New York, NY [Heidelberg], 2005).
- [9] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*, Vol. 2 (MIT press Cambridge, MA, 2006).
- [10] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016) <http://www.deeplearningbook.org>.
- [11] R. Upadhyay, K. Miki, O. Ezekoye, and J. Marschall, Uncertainty quantification of a graphite nitridation experiment using a bayesian approach, *Experimental Thermal*

- and Fluid Science **35**, 1588 (2011).
- [12] A. del Val, O. P. Le Maître, T. E. Magin, O. Chazot, and P. M. Congedo, A surrogate-based optimal likelihood function for the bayesian calibration of catalytic recombination in atmospheric entry protection materials, *Applied Mathematical Modelling* **101**, 791 (2022).
- [13] A. del Val, O. P. Le Maître, P. M. Congedo, and T. E. Magin, Stochastic calibration of a carbon nitridation model from plasma wind tunnel experiments using a bayesian formulation, *Carbon* **200**, 199 (2022).
- [14] B. Bottin, O. Chazot, M. Carbonaro, V. V. der Haegen, and S. Paris, The vki plasmatron characteristics and performance (2000).
- [15] M. Capriati, A. del Val, T. E. Schwartzentruber, T. K. Minton, P. M. Congedo, and T. E. Magin, Bayesian calibration of a finite-rate nitridation model from molecular beam and plasma wind tunnel experiments, in *Aerospace Europe Conference 2023–10th EUCASS–9th CEAS* (2023).
- [16] J. Marschall and M. MacLean, Finite-rate surface chemistry model, i: Formulation and reaction system examples, in *42nd AIAA thermophysics conference* (2011) p. 3783.
- [17] K. Prata, T. Schwartzentruber, and T. Minton, English (US) Air-carbon ablation model for hypersonic flight from molecular-beam data, *AIAA journal* **60**, 627 (2022).
- [18] V. J. Murray, P. Recio, A. Caracciolo, C. Miossec, N. Balucani, P. Casavecchia, and T. K. Minton, Oxidation and nitridation of vitreous carbon at high temperatures, *Carbon* **167**, 388 (2020).
- [19] A. J. Lutz, *Experimental investigation and analysis of high-enthalpy nitrogen flow over graphite* (The University of Vermont and State Agricultural College, 2015).
- [20] M. Bandera, Stagnation line aerothermochemistry. focus on uncertainty quantification in nitrogen plasma-graphite interactions in vki plasmatron, (2023).
- [21] T. E. Magin, M. Capriati, G. Kale, G. Bellas, A. Turchi, B. Helber, and O. Chazot, *PLASMUT - Nonequilibrium Gas-Surface Interactions at High Temperature, VKI Plasmatron and Mutation++Library*, Tech. Rep. (von Karman Institute for Fluid Dynamics, 2023).
- [22] T. E. Magin, M. Capriati, F. Bariselli, B. Helber, and O. Chazot, *DRAGROUND FLIGHT - Disruptive research approach for GSI-model consolidation through on-ground and in-flight analyses*, Tech. Rep. (von Karman Institute for Fluid Dynamics, 2025).
- [23] R. Goulard, On catalytic recombination rates in hypersonic stagnation heat transfer, *Journal of Jet Propulsion* **28**, 737 (1958).
- [24] A. Kolesnikov, The concept of local simulation for stagnation point heat transfer in hypersonic flows - applications and validation, in *21st Aerodynamic Measurement Technology and Ground Testing Conference*, <https://arc.aiaa.org/doi/pdf/10.2514/6.2000-2515>.
- [25] A. Munafò and T. E. Magin, Modeling of stagnation-line nonequilibrium flows by means of quantum based collisional models, *Physics of Fluids* **26**, 097102 (2014).
- [26] B. McBride, M. Zehe, and S. Gordon, *NASA Glenn coefficients for calculating thermodynamic properties of individual species* (Technical report, NASA, 10 2002).
- [27] T. E. Magin and G. Degrez, Transport algorithms for partially ionized and unmagnetized plasmas, *Journal of Computational Physics* **198**, 424 (2004).
- [28] A. Klomfass and S. Müller, Calculation of stagnation streamline quantities in hypersonic blunt body flows, *Shock Waves* **7**, 13 (1997).
- [29] P. Roe, Approximate riemann solvers, parameter vectors, and difference schemes, *Journal of Computational Physics* **135**, 250 (1997).
- [30] J. B. Scoggins, V. Leroy, G. Bellas-Chatzigeorgis, B. Dias, and T. E. Magin, Mutation++: Multi-component thermodynamic and transport properties for ionized gases in c++, *SoftwareX* **12:100575**, 10.1016/j.softx.2020.100575 (2020).
- [31] C. Lataniotis, S. Marelli, and B. Sudret, *UQLab user manual – Kriging (Gaussian process modelling)*, Tech. Rep. (Chair of Risk, Safety & Uncertainty Quantification, ETH Zurich, 2015) report UQLab-V0.9-105.
- [32] S. Marelli and B. Sudret, Uqlab: A framework for uncertainty quantification in matlab, in *Vulnerability, Uncertainty, and Risk* (American Society of Civil Engineers, 2014) pp. 2554–2563.
- [33] P.-R. Wagner, J. Nagel, S. Marelli, and B. Sudret, *UQLab user manual – Bayesian inversion for model calibration and validation*, Tech. Rep. (Chair of Risk, Safety and Uncertainty Quantification, ETH Zurich, Switzerland, 2022) report UQLab-V2.0-113.
- [34] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, Pytorch: An imperative style, high-performance deep learning library (2019), arXiv:1912.01703 [cs.LG].
- [35] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, Optuna: A next-generation hyperparameter optimization framework, in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2019).
- [36] J. Wakefield *et al.*, *Bayesian and frequentist regression methods* (Springer, 2013).
- [37] K. Y. Chan, B. Abu-Salih, R. Qaddoura, A. M. Al-Zoubi, V. Palade, D.-S. Pham, J. D. Ser, and K. Muhammad, Deep neural networks in the cloud: Review, applications, challenges and research directions, *Neurocomputing* **545**, 126327 (2023).
- [38] D. P. Kingma, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).
- [39] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *The journal of chemical physics* **21**, 1087 (1953).
- [40] F. Homann, *The effect of high viscosity on the flow around a cylinder and around a sphere*, Tech. Rep. (1952).