

BAYESIAN IDENTIFICATION OF PYROLYSIS MODEL PARAMETERS FOR THERMAL PROTECTION MATERIALS USING AN ADAPTIVE GRADIENT-INFORMED SAMPLING ALGORITHM WITH APPLICATION TO A MARS ATMOSPHERIC ENTRY

Joffrey Coheur,^{1,2,3,*} Thierry E. Magin,² Philippe Chatelain,³ & Maarten Arnst¹

¹Aerospace & Mechanical Engineering, Université de Liège, Allée de la Découverte 9, 4000 Liège, Belgium

²Aeronautics and Aerospace Department, von Karman Institute for Fluid Dynamics, Chaussée de Waterloo 72, 1640 Rhode-Saint-Genèse, Belgium

³Institute of Mechanics, Materials and Civil Engineering (iMMC), Université catholique de Louvain, Place du Levant 2, 1348 Louvain-la-Neuve, Belgium

*Address all correspondence to: Joffrey Coheur, Aerospace & Mechanical Engineering, Université de Liège, Allée de la Découverte 9, 4000 Liège, Belgium, E-mail: joffrey.coheur@uliege.be

Original Manuscript Submitted: 1/26/2022; Final Draft Received: 9/11/2022

For space missions involving atmospheric entry, a thermal protection system is essential to shield the spacecraft and its payload from the severe aerothermal loads. Carbon/phenolic composite materials have gained renewed interest to serve as ablative thermal protection materials (TPMs). New experimental data relevant to the pyrolytic decomposition of the phenolic resin used in such carbon/phenolic composite TPMs have recently been published in the literature. In this paper, we infer from these new experimental data an uncertainty-quantified pyrolysis model. We adopt a Bayesian probabilistic approach to account for uncertainties in the model identification. We use an approximate likelihood function involving a weighted distance between the model predictions and the time-dependent experimental data. To sample from the posterior, we use a gradient-informed Markov chain Monte Carlo method, namely, a method based on an Itô stochastic differential equation, with an adaptive selection of the numerical parameters. To select the decomposition mechanisms to be represented in the pyrolysis model, we proceed by progressively increasing the complexity of the pyrolysis model until a satisfactory fit to the data is ultimately obtained. The pyrolysis model thus obtained involves six reactions and has 48 parameters. We demonstrate the use of the identified pyrolysis model in a numerical simulation of heat-shield surface recession in a Martian entry.

KEY WORDS: Bayesian inference, chemical kinetics, carbon/phenolic composite, Markov chain Monte Carlo, Itô stochastic differential equation, thermal protection system

1. INTRODUCTION

Ablative porous thermal protection materials (TPMs) are used in heat shields for atmospheric entries of spacecraft. Low-density carbon/phenolic composite materials, such as the Phenolic Impregnated Carbon Ablator (PICA) material, have gained renewed interest to serve as ablative TPMs [1–6]. These carbon/phenolic composite materials can accommodate the high heating rates and heat loads encountered during the entry, at hypersonic velocities, by absorbing part of the incoming heat through physico-chemical transformations [7]. One of the main endothermic processes

is the pyrolysis of the phenolic compound, whereby volatile products are released and a carbonaceous residue is left on the carbon fibers.

New experimental studies have recently been reported relevant to the pyrolytic decomposition of the phenolic resin used in carbon/phenolic composites, namely, in Wong et al. [8] and Torres-Herrador et al. [9], under several isothermal conditions, and in Bessire et al. [10] for linearly increasing temperatures. In these experimental studies, the mass of a sample is monitored as a function of temperature either by using thermogravimetric analysis (TGA) or by collecting the volatile decomposition products by means of a mass spectrometer (in [8]) or gas chromatography (in [10] and [9]).

The objective of this paper is to infer from the new experimental data of Bessire et al. [10] an uncertainty-quantified pyrolysis model. The class of pyrolysis models that we adopt is the class of multi-component single-step models [11,12], which represent the decomposition mechanisms in terms of a system of chemical kinetic equations with reaction rates that obey an Arrhenius-like generic model. In previous work by Torres-Herrador et al. [13], the new experimental data of Bessire et al. [10] have already been used to estimate parameters of such a multi-component single-step pyrolysis model, but the model identification was formulated as a deterministic inverse problem in [13], and an estimation of associated uncertainties is thus still required.

There have been several efforts to infer kinetic parameters of Arrhenius-like chemical reaction models from experimental data with a probabilistic approach, especially for combustion applications; see, for instance, [14–17]. Investigations of pyrolysis experiments by using Bayesian inference have been performed for high-impact polystyrene, bisphenol-A polycarbonate, and polyvinyl chloride at two different heating rates for fire applications in [18]. In recent work, Torres-Herrador et al. [19] have presented a probabilistic parameter inference for the decomposition of the PICA material, but the focus of the research in [19] was on establishing the relevance of accounting for competitive mechanisms in the pyrolysis model.

Several factors make it challenging to formulate in a probabilistic setting the identification of a multi-component single-step model with Arrhenius-like reaction rates from the new experimental data of Bessire et al. [10]. Notably, the experimental data are subject to measurement uncertainty and sample variability, but Bessire et al. [10] provide only a limited characterization of these sources of uncertainty: they repeated their experimental study a small number of times (three or five) under nominally identical conditions and then used the differences in the observed species productions between the repeated experimental runs to deduce error bars for the experimental data. In addition, significant modeling errors can be expected to be present: significant physical mechanisms responsible for species production in the experiment can be expected to be either unmodeled or represented imperfectly in the multi-component single-step model with Arrhenius-like reaction rates. For example, in the initial stage of the experiment, physical mechanisms other than pyrolysis, such as outgassing of gases like water that had been absorbed by the sample, may cause significant species production. Furthermore, expressing the reaction rates with Arrhenius laws may imperfectly represent the physical behavior. Because of the lack of knowledge of the measurement uncertainty and sample variability and because of the likely presence of significant modeling errors, the model identification cannot be readily formulated as a standard problem of Bayesian statistical inference of a deterministic model augmented with a stochastic representation of experimental noise.

To account for modeling errors in model identification, there exist approaches that seek to incorporate a representation of the model discrepancy explicitly into the model, such as by incorporating additional stochastic terms into the governing equations [20,21]. However, such approaches significantly modify the model: they convert a model in the form of a system of differential equations into a stochastic model in the form of a system of stochastic differential equations. It is then this stochastic model that is identified from the data. In the application that we study, the purpose of the identified pyrolysis model is to serve in material response codes for numerical simulations of heat-shield recession during atmospheric entries. Such approaches would thus require the identified stochastic pyrolysis model, with the additional stochastic terms inherent to it, to be implemented within such numerical tools, thus turning these numerical simulations themselves into stochastic simulations.

In this work, we adopt an alternative approach whereby we keep the structure of the pyrolysis model unchanged (we do not introduce additional stochastic terms) and we seek instead to incorporate a representation of the uncertainties into the model parameters. We formulate the model identification as a Bayesian probabilistic inverse problem, in which we use an approximate likelihood function involving a weighted distance between the model predictions and

the experimental data. This formulation provides a probability distribution over the model parameters, the posterior, which assigns higher probabilities to parameter values that correspond to predictions that better fit the experimental data. We introduce a free parameter in this approximate likelihood function that allows us to control how significantly the probability of parameter values is penalized in proportion to the misfit between the corresponding predictions and the experimental data. This approach allows the free parameter to be tuned to seek to introduce into the model parameters an amount of uncertainty that can serve as an overall representation of the impact of various sources of uncertainty, including measurement uncertainty, sample variability, and modeling errors. In this work, we rely on an *ad hoc*, visual, criterion to tune the free parameter.

The number of model parameters of the pyrolysis model is high, and the posterior exhibits strong dependencies, so that sampling from the posterior is challenging. Notably, pre-exponential factors and activation energies involved in Arrhenius laws exhibit significant dependencies because of their complementary role in the expression of Arrhenius's law. This high correlation is called the kinetic compensation effect and has already been observed in numerous kinetics studies [18,19,22,23] (also in deterministic inference, see, for instance, [24,25]). Bruns [18] and Najm [14] applied adaptive random-walk Metropolis–Hastings algorithms, with a reparameterization of the pre-exponential factor using a logarithmic change of variables, to sample efficiently from posteriors over parameters of Arrhenius laws. In this work, we use a gradient-informed Markov chain Monte Carlo method, namely, a method based on an Itô stochastic differential equation [26–28]. This method, at least in the variant used in this paper, requires an estimate of the covariance matrix of the posterior, which we propose to learn adaptively from previous samples in the chain, similarly to adaptive Metropolis–Hastings algorithms.

Because we do not introduce additional stochastic terms in the pyrolysis model, our approach allows the identified pyrolysis model to be used in material response codes for numerical simulations of heat-shield recession in atmospheric entries without requiring stochasticity to be implemented explicitly inside these codes. Notably, our approach allows the uncertainties associated with the identified pyrolysis model to be propagated through such a numerical simulation of heat-shield recession in an atmospheric entry by using a Monte Carlo approach implemented as a wrapper around the code. The deterministic pyrolysis model is thereby implemented inside a material response code, which is then run for multiple values of the model parameters drawn from the posterior. At the end of this paper, we demonstrate such a use of the identified pyrolysis model in a numerical simulation of heat-shield surface recession during a Martian entry.

This paper is organized as follows. In Section 3, the experiments and the multi-component single-step model with Arrhenius-like reaction rates are first reviewed. In Section 4, the Bayesian probabilistic formulation, the expression for the distance in the approximate likelihood function, and the choice of the weights are described. The gradient-informed algorithm is then presented in Section 5, along with the adaptive method used to approximate the covariance matrix. Next, in Section 6, numerical results are provided, including a study of computational aspects and efficiency of the employed adaptive sampling algorithm. In Section 6, the results from the inference are also used as input uncertainty for an uncertainty propagation through a numerical simulation of the impact on the recession of the surface of an ablating TPS during an atmospheric entry on Mars.

2. SYSTEM OF NOTATION

The system of notation in this work is inspired from the tensorial notation from Kolda and Bader [29]. Vectors (tensors of order one) are denoted by boldface lowercase letters, e.g., $\mathbf{a}$. Matrices (tensors of order two) are denoted by capital letters, e.g., A . Higher-order tensors are denoted by boldface Euler script letters, e.g., $\mathcal{W}$. Scalars are denoted by lowercase letters, e.g., a . The i th entry of a vector $\mathbf{a}$ is denoted by a_i , the element (i, j) of a matrix A is denoted by a_{ij} , the element (i, j, k) of the third-order tensor $\mathcal{W}$ is denoted w_{ijk} , etc. Subarrays are formed when a subset of the indices is fixed and we use a colon to denote all elements of the fixed dimension. For instance, $\mathbf{a}_{i:}$ refers to the i th row of the matrix A and $\mathbf{a}_{:,j}$ refers to the j th column of A .

3. PYROLYSIS EXPERIMENTS AND MODELING

The materials considered in this study are composites made up of short carbon fibers infused with a phenolic resin that are used in the thermal protection system of spacecraft. In particular, we consider the phenolic-impregnated

carbon ablator (PICA) that has been successfully used in previous space missions for its excellent thermal protection performance [2,30].

Studies on ablative materials date back to the 60's with the development of early material response codes [31,32] based on the experiments of Sykes [33] and Goldstein [34]. Several experiments have been carried out more recently to get new insight into the pyrolysis of the resin [8,10,35–38] and improve models for numerical solvers. Among the experimental studies on the pyrolysis of TPMs, the pyrolysis experiments of Bessire and Minton [10] on the PICA material provide a measure of both the mass loss and the species production. The measurement of the species that are produced during the pyrolysis allows us to have a better idea of the finite-rate chemical reactions and in which proportion they occur, which could ultimately help identifying the underlying decomposition mechanisms occurring within the material.

In this section, the pyrolysis experiments carried out in Bessire and Minton [10] are first described and the choice of the data used in this work is motivated. The model and the system of ordinary differential equations that characterize the pyrolysis are then summarized in Section 3.2.

3.1 Pyrolysis Experiments

The heating rates encountered during an actual re-entry application largely vary within the heat shield. In addition, they may be significantly higher than current ground-testing techniques. In the experiments of Bessire et al. [10], the pyrolysis of a 0.441 g PICA sample was performed at significantly different heating-rate values, namely 186, 366, 762, and 1500 K min⁻¹. In these experiments, the PICA sample is heated from 375 to 1400 K in a vacuum chamber by means of an electrical current that passes through the sample. *In situ* measurements of PICA pyrolysis product yields are performed using mass spectrometry, which provides the product yields of 14 gaseous species as a function of the material temperature.

The experimental results at the heating rate of 366 K min⁻¹ for the mass loss and the yield of gaseous species are shown in Fig. 1. The x axis represents the temperature measured inside the material, which is here a function of time as $T(t) = \beta t$ for $t_0 < t < t_{\text{end}}$, where β denotes the heating rate in K min⁻¹.

We formalize the results of the experiments as follows. We denote by n_r the number repetitions of the experiment. The data are acquired at given time values that depend on the acquisition frequency, which is fixed by the experimental device, and results in n_t data points for each observation (for instance, at $\beta = 366$ K min⁻¹, $n_t = 69$, see Fig. 1). The different experimental observations (species, mass loss) as a function of time, or temperature, with n_{obs} the total number of experimental observations, can be specified in a third-order tensor of observations $\mathcal{Y}^{\text{obs}}$ of dimensions $n_{\text{obs}} \times n_t \times n_r$. The corresponding k th matrix of observations is thus denoted

$$Y_{::k}^{\text{obs}} = \begin{bmatrix} y_{11k}^{\text{obs}} & y_{12k}^{\text{obs}} & \cdots & y_{1n_tk}^{\text{obs}} \\ y_{21k}^{\text{obs}} & y_{22k}^{\text{obs}} & \cdots & y_{2n_tk}^{\text{obs}} \\ \vdots & \vdots & \ddots & \vdots \\ y_{n_{\text{obs}}1k}^{\text{obs}} & y_{n_{\text{obs}}2k}^{\text{obs}} & \cdots & y_{n_{\text{obs}}n_tk}^{\text{obs}} \end{bmatrix}, \quad (1)$$

with $1 \leq k \leq n_r$. In particular, in Bessire and Minton [10], the experiments were repeated three times ($n_r = 3$) and a mean and a standard deviation as a function of temperature were estimated for each observation (see Fig. 1):

$$\bar{y}_{ij}^{\text{obs}} = \frac{1}{n_r} \sum_{k=1}^{n_r} y_{ijk}^{\text{obs}}, \quad (2)$$

$$s_{ij}^2 = \frac{1}{n_r - 1} \sum_{k=1}^{n_r} (y_{ijk}^{\text{obs}} - \bar{y}_{ij}^{\text{obs}})^2. \quad (3)$$

with $\bar{Y}^{\text{obs}}$ the matrix of mean observations and S the matrix of the standard deviations of the observations.




FIG. 1: Total mass loss and the 14 mass yields at a heating rate of 366 K min^{-1} from the experiments of [10]. The markers are the estimated means [Eq. (2)] and the error bars represent the empirical standard deviations [Eq. (3)]. (a) Total mass loss, (b) permanent gases and water, (c) phenol and derivatives, and CO_2 , and (d) alcohol and aromatics.

3.2 Multi-Component Single-Step Pyrolysis Model

Multi-component single-step pyrolysis models are widely used for the modeling of the pyrolysis of thermal protection materials and are implemented in most state-of-the-art material response codes [39]. The global pyrolysis reaction mechanism is typically encapsulated within phenomenologically coherent steps describing the elementary unknown processes that are responsible for the production peaks observed in Fig. 1. Multiple advancement of reaction factors $0 \leq \alpha_i \leq 1$ are defined, where $1 \leq i \leq N_p$ indexes independent pyrolysis reactions occurring in parallel and N_p is the number of pyrolysis decomposition processes present in the solid resin (or the number of reactants, or the number of steps). The evolution of the advancement of reaction factors is given by the expression

$$\frac{d\alpha_i}{dt}(t) = \mathcal{A}_i \exp\left\{\frac{-\mathcal{E}_i}{\mathcal{R}T(t)}\right\} (1 - \alpha_i(t))^{r_i}, \quad 1 \leq i \leq N_p, \quad (4)$$

where $\mathcal{R} = 8.3145 \text{ (J K}^{-1} \text{ mol}^{-1})$ is the universal gas constant, $\mathcal{A}_i \text{ (s}^{-1})$ is the pre-exponential factor, $\mathcal{E}_i \text{ (J mol}^{-1})$ is the activation energy, and r_i is the reaction order for the i th reactant. The total mass of the sample as a function of time can be linked to the advancement of reaction factors as

$$m(t) = m_0 - m_0 \sum_{i=1}^{N_p} F_i \alpha_i(t), \quad (5)$$

where m_0 is the initial mass of the sample and F_i is the mass fraction of the i th reaction, which is the ratio of the initial mass that is converted into gas during the pyrolysis process due to the i th reaction.

The production of an observed gaseous species $1 \leq j \leq n_G$ is obtained from the derivative of the mass loss as

$$\nu_j(t) = m_0 \sum_{i=1}^{N_p} \gamma_{ij} \frac{d\alpha_i}{dt}(t), \quad (6)$$

where γ_{ij} is the mass fraction of species j produced from reaction i . The mass m as a function of time and the mass fractions F_i can then be retrieved by assuming that all the sample mass loss is converted into gas such that we can write $m(t) = m_0 - \sum_{j=1}^{n_G} \int_0^t \nu_j(s) ds$, which results in $\sum_{j=1}^{n_G} \gamma_{ij} = F_i$.

The vector of unknown parameters, or the vector of the calibration parameters, is denoted $\mathbf{m} = \{\mathcal{A}_i, \mathcal{E}_i, r_i, \gamma_{ij}, 0 \leq i \leq N_p, 1 \leq j \leq n_G\}$. The number of unknowns in the inference problem, denoted n_u , can be high due to the introduction of the γ_{ij} , up to $n_u = (3 + n_G)N_p$. In order to reduce the number of unknowns, we can proceed by inspection of the pyrolysis mechanism from the experimental curves and set to zero the γ_{ij} that are thought not to contribute to the production of species j due to reaction i . We denote by $I_{i,G}$ the subset of indexes from $1 \leq j \leq n_G$ corresponding to the values of γ_{ij} that participate to the production of species j from reaction i , and $I_G = \{I_{i,G}, 0 \leq i \leq N_p\}$. The number of unknowns is therefore $n_u = 3N_p + |I_G| \leq (3 + n_G)N_p$.

When the temperature is a linear function of time (constant heating rate), an analytical solution can be found for the multi-component single-step pyrolysis model [Eq. (6)] and its gradients [12]. We will use the analytical expressions in the gradient-informed sampling method which will significantly speed up the computational time taken by each evaluation of the model and its gradients and avoid discretization errors in the solution of the system of ODEs.

One issue of the present multi-component single-step pyrolysis model is that it is suitable only to heating rates close to which the parameters are calibrated. An other model with the calibration of its parameters might be needed to describe the decomposition at other heating rates. To include in a single model multiple heating rates, a more general pyrolysis model that takes into account competitive reactions and successive dependent steps was presented in [19] for aerospace applications. The uncertainty quantification methodology proposed next can be readily extended to the inference using the competitive mechanisms, at the cost of a more expensive computation of the solution and its gradients as there is no analytical expression of the solution available in this case.

3.3 Model Representation

Let $\mathbf{g}(\cdot, \mathbf{m}) : \mathbb{R}^+ \rightarrow \mathbb{R}^{n_{\text{obs}}}$ be the time-dependent solution to the computational model representing the dynamic physical process, with the vector of model outputs $\mathbf{g}(\cdot, \mathbf{m}) = (g_1(\cdot, \mathbf{m}), g_2(\cdot, \mathbf{m}), \dots, g_{n_{\text{obs}}}(\cdot, \mathbf{m}))$ with values in $\mathbb{R}^{n_{\text{obs}}}$, provided the model parameters $\mathbf{m} \in \mathbb{R}^{n_u}$. The value predicted by the model for a given observable indexed by i , $1 \leq i \leq n_{\text{obs}}$, at time $t = t_j$, $1 \leq j \leq n_t$, is denoted by

$$d_{ij} = g_{ij}(\mathbf{m}) = g_i(t_j, \mathbf{m}). \quad (7)$$

The different observables can be the different species produced by the pyrolysis reaction, e.g., $n_{\text{obs}} = n_G$. We refer, for instance, to Torres-Herrador et al. [19] for an example where the observable is the derivative of the mass loss, or Bruns [18] where the observable is the mass loss. The choice here to use the individual species production is because we expect that it provides more information on the different reaction steps and the reaction mechanisms.

4. BAYESIAN APPROACH TO INVERSE PROBLEM FOR PARAMETER CALIBRATION

We adopt a Bayesian approach to inverse problems in which the solution to the inverse problem is the probability distribution of the unknown parameters $\mathbf{m}$ given the value of the experimental observations $\mathbf{y}^{\text{obs}}$ with the following density:

$$\pi(\mathbf{m} | \mathbf{y}^{\text{obs}}) \propto \exp(-\Phi(\mathbf{y}^{\text{obs}}, G(\mathbf{m}))) \pi_0(\mathbf{m}), \quad (8)$$

where the symbol $\propto$ is interpreted to mean “proportional up to a normalization factor independent of $\mathbf{m}$.” The exponential factor is an approximate likelihood function and depends on a potential function Φ , or misfit function, that

seeks to gauge the goodness of fit between the actual observations and the model output. The matrix G represents the error-free matrix of observations that corresponds to the quantities defined in Eq. (7). The second factor $\pi_0 : \mathbb{R}^{n_u} \rightarrow \mathbb{R}^+$ is the prior and acts as a regularization that can incorporate any prior knowledge about the parameters $\mathbf{m}$.

4.1 Prior Density Function

The only prior knowledge we have about the parameters is that $\mathcal{A}_i$, $\mathcal{E}_i$, and r_i are positive and the mass fractions γ_{ij} are bounded between 0 and 1. We choose the following prior:

$$\pi_0(\mathbf{m}) = \prod_{i=1}^{N_p} 1_{\mathbb{R}^+}(\mathcal{A}_i) 1_{[L_{\mathcal{E}_i}, U_{\mathcal{E}_i}]}(\mathcal{E}_i) 1_{[1, 10]}(r_i) \prod_{j \in I_G} 1_{[0, 1]}(\gamma_{ij}), \quad (9)$$

where $1_{\mathbb{S}}$ is the indicator function of $\mathbb{S} \subseteq \mathbb{R}$. The activation energy is strictly positive, but in the case where several pyrolysis reactions are considered, the values of $\mathcal{E}_i$ are further limited to the bounded domain $[L_{\mathcal{E}_i}, U_{\mathcal{E}_i}]$ in order to avoid the different pyrolysis reactions to mix together, which would make the potential function more complex to explore. The maximum values for r_i are limited to 10, as higher values were not found in previous similar pyrolysis models [13], and first-order reactions ($r_i = 1$), for which the analytical solution is different, are not considered here [12].

4.2 Choice of the Approximate Likelihood Function

In Eq. (8), the goodness of fit is through the approximate likelihood function $\exp(-\Phi(\mathbf{y}^{\text{obs}}, G(\mathbf{m})))$. We define the potential function using the sum over all the observations of the squares of the weighted norms of the differences between the model output and the observations, similarly to a weighted least-squares formulation in optimization,

$$\Phi(\mathbf{y}^{\text{obs}}, G(\mathbf{m})) = \frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \phi_i(\bar{\mathbf{y}}_{i:}^{\text{obs}}, \mathbf{g}_{i:}(\mathbf{m})) = \frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \|\bar{\mathbf{y}}_{i:}^{\text{obs}} - \mathbf{g}_{i:}(\mathbf{m})\|_{W_{::i}}^2, \quad (10)$$

where $\|\mathbf{x}\|_W = \mathbf{x}^\top W \mathbf{x}$ denotes the weighted norm of a vector $\mathbf{x} \in \mathbb{R}^{n_t}$ and $W_{::i}$ is the $\mathbb{R}^{n_t \times n_t}$ symmetric positive-definite matrix of weights, formed from the i th subarray of the third-order tensor of weights $\mathcal{W}$. If we assume further that the weighting matrices $W_{::i}$ are diagonal, i.e., $w_{kji} = 0, \forall k \neq j, 1 \leq k \leq n_t, 1 \leq j \leq n_t$, then we arrive at the following function for the likelihood

$$\exp(-\Phi(\mathbf{y}^{\text{obs}}, G(\mathbf{m}))) = \exp\left(-\frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \sum_{j=1}^{n_t} w_{jji} (\bar{y}_{ij}^{\text{obs}} - g_{ij}(\mathbf{m}))^2\right), \quad (11)$$

where $\bar{y}_{ij}^{\text{obs}}$ is computed from Eq. (2).

In classical optimization problems, the weights w_{jji} can be used to rescale the different physical quantities, or they can be used to give more weight to certain data rather than other. Equal importance is specified with $w_{jji} = 1, \forall i, j$. In the Bayesian approach, larger values of the weights result in a posterior that is more concentrated around the best-fit parameter values.

To choose the values of the w_{jji} , we will introduce two normalizations. The first normalization takes into account the observed dispersion [the s_{ij} in Eq. (3)] by choosing the weights w_{jji} inversely proportional to s_{ij}^2 . The second normalization takes into account the number of data acquired in the time series n_t and is chosen proportional to $(\Delta t_j / t_f)$, where t_f is the total time span of the experiment and Δt_j are the data acquisition time, and, for a constant acquisition time, the ratio $\Delta t_j / t_f = n_t^{-1}$. This second normalization prevents the sum in the exponential from becoming very high if the number of data acquired is high, as in this case Δt_j tends to 0. The presence of this factor proportional to the number of terms in the sum yields a posterior that is less dependent on the data acquisition frequency.

We set $w_{ij} = \vartheta s_{ij}^{-2} (\Delta t_j / t_f)$ with s_{ij} computed from Eq. (3) and ϑ a free (tuning) parameter that must be chosen. This choice finally leads to the following expression for the likelihood function:

$$\exp(-\Phi(\mathbf{y}^{\text{obs}}, G(\mathbf{m}))) = \exp\left(-\frac{1}{2} \sum_{i=1}^{n_{\text{obs}}} \sum_{j=1}^{n_t} \vartheta \frac{\Delta t_j}{t_f} \left(\frac{\bar{y}_{ij}^{\text{obs}} - g_{ij}(\mathbf{m})}{s_{ij}} \right)^2\right). \quad (12)$$

The free parameter ϑ controls all the weights at the same time and will influence the width of the resulting posterior. Later in this work, we will rely on an *ad hoc*, visual, criterion to tune the free parameter, so that we introduce in the model parameters an amount of uncertainty that represents an overall impact of measurement uncertainty, sample variability, and modeling errors.

5. MARKOV CHAIN MONTE CARLO METHOD

5.1 Algorithm Based on an Itô Stochastic Differential Equation

The gradient-informed MCMC algorithm considered is based on the solution of an Itô stochastic differential equation (ISDE) [26–28]. The ISDE introduced by Soize [26] describes a dissipative Hamiltonian system characterized by its position $\mathbf{q}$ and momentum $\mathbf{p}$ with a stochastic forcing term. Soize [27] and Arnst et al. [40,41] applied a generalized version of the ISDE that takes into account the scales and correlations between the different parameters. The following stochastic differential equation describes a damped second-order system with a stochastic excitation:

$$d\mathbf{Q} = M^{-1}\mathbf{P}dt, \quad (13)$$

$$d\mathbf{P} = -\nabla_{\mathbf{q}}\mathcal{U}(\mathbf{Q})dt - \frac{1}{2}f_0\mathbf{P}dt + \sqrt{f_0}L_M d\mathbf{W}, \quad (14)$$

where $M = L_M L_M^T$ is a symmetric, positive-definite matrix with L_M the lower triangular matrix resulting from the Cholesky decomposition, f_0 is a free parameter ([41]) that controls the amplitude of the deterministic damping and the stochastic forcing. We denote by $\{\mathbf{Q}(t), t \geq 0\}$ and $\{\mathbf{P}(t), t \geq 0\}$ two $\mathbb{R}^{n_u}$ -valued stochastic processes indexed by $\mathbb{R}^+$ and $\{\mathbf{W}(t), t \geq 0\}$ the $\mathbb{R}^{n_u}$ -valued normalized Wiener process whose increment $\Delta\mathbf{W}_{st} = \mathbf{W}(t) - \mathbf{W}(s)$ is an $\mathbb{R}^{n_u}$ -valued centered Gaussian random vector with covariance matrix $(t - s)I_{n_u}$. In the limit of zero damping, the system is similar to a (deterministic) conservative Hamiltonian system with potential energy $\mathcal{U}(\mathbf{q})$ and kinetic energy $(1/2)\mathbf{p}^T M^{-1}\mathbf{p}$.

Under regularity conditions regarding the potential function $\mathcal{U}$, the stationary density of $\mathbf{Q}, \mathbf{P}$ is independent of the damping parameter and is given by

$$\pi(\mathbf{q}, \mathbf{p}) \propto \exp\left(-\mathcal{U}(\mathbf{q}) - \frac{\mathbf{p}^T M^{-1}\mathbf{p}}{2}\right). \quad (15)$$

By letting the position variable $\mathbf{q}$ be the variable of interest $\mathbf{m}$ and choosing a potential function that corresponds to $\mathcal{U}(\mathbf{m}) = -\log(\pi(\mathbf{m}|\mathbf{y}^{\text{obs}}))$, the marginal stationary density for $\mathbf{Q}$ will be exactly the distribution we wish to sample from (here, the Bayesian posterior). The dynamics of the ISDE is not reversible due to the dissipative behavior of the system and the accept-reject step from the Metropolis–Hastings cannot be applied. In order to limit the error induced by the numerical discretization, the time step τ will have to be chosen carefully. The numerical solution of the stochastic irreversible (non-Hamiltonian) system is performed by applying the Störmer–Verlet scheme [28].

The implementation of the ISDE algorithm assumes that the support of the target pdf is the entire space $\mathbb{R}^{n_u}$. It is not directly applicable in the case of a bounded or semi-bounded parameter set as the algorithm may reach parameter values where the potential is not defined. To avoid this, Soize [28] has proposed to use regularization or rejection techniques. Here, we will use a change of variable as in Arnst et al. [40] that will map the original variables into a transformed parameter space where the ISDE algorithm is applicable. We distinguish two cases. First, the case where $\mathbb{S}$ is bounded from above and below with $\mathcal{S} = [L_b, U_b]$, we choose a change of variables of the following form:

$$\tilde{m}_i = \tan\left\{\left(\frac{m_i - L_b}{U_b - L_b} - \frac{1}{2}\right)\pi\right\}. \quad (16)$$

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

5.2 Change of Scales: Approximation of the Covariance

In Eq. (15), the momentum variable is a Gaussian random vector whose covariance matrix is the symmetric, positive-definite matrix M ; thus, the linear transformation $d\mathbf{Q}/dt = M^{-1}\mathbf{P}$ is a Gaussian random vector with covariance matrix $M^{-1}MM^{-T} = M^{-T} = M^{-1}$. Hence, the matrix $M^{-1} \equiv C$ is the covariance matrix of the position variable $\mathbf{Q}$, and can be interpreted as a representation of the scales and correlations of the problem [40]. It is crucial for the mixing and convergence of the Markov chain toward the posterior distribution to have a good approximation of this covariance matrix.

For the tuning of the proposal in Metropolis–Hastings algorithms, Girolami and Calderhead [42] use a mechanism that learns the local structure of the log posterior distribution at the current state of the Markov chain for two other gradient-informed sampling algorithms, namely the Metropolis-adjusted Langevin and the Hamiltonian Monte Carlo algorithms. Law [43] uses proposals that include a local estimation of the value of the Hessian of the sign-reversed logarithm of the posterior, which is representative of the local scales and correlations. However, this position-dependent approach requires modifications to the differential equations of the dynamical system of the proposal mechanism [40,42]. For the ISDE algorithm, references [40,41] use a fixed covariance approximation $\hat{C}$ based on the value of the Hessian of the sign-reversed logarithm of the posterior at the maximum likelihood estimate. This estimate of the covariance matrix can sometimes lack accuracy, for instance, if the Hessian is estimated at a parameter value that is far from the maximum likelihood value (because, for instance, we do not know it) and where the approximation of the covariance may result in a matrix that is not positive-definite. To avoid this, the initial estimate $\hat{C}$ can be set to the identity matrix, $\hat{C} = I$ (we do not provide any information about scales and correlations), or we can also use the diagonal of the Hessian estimate that can be easily made positive-definite by taking the absolute value of the diagonal elements (it represents only the local scales and not the correlations). Depending on the choice of the covariance matrix estimation, the value for the time step τ will have to be chosen carefully. The adaptation of the parameters of the discretization scheme can also be performed [44,45], but this is not investigated here.

In order to have a more global estimation of the scales and correlations, which is also robust independently of the covariance matrix estimate, we propose to use an adaptive procedure in order to refine the initial covariance estimation. Adaptation of the covariance matrix for gradient-informed algorithms was already used in Atchadé [46] and Marshall and Roberts [45] for the Metropolis-adjusted Langevin algorithm. Similarly, we use here an adaptive procedure for the covariance matrix that we apply to the ISDE algorithm. The covariance matrix is estimated from past sampled values, as in the adaptive Metropolis–Hastings from [47–51]. The adaptation starts after an initial burn-in phase and the whole sample history is used in the covariance computation. The adaptation updates the proposal covariance matrix based on previous iterates at given increments Δn . The covariance matrix is set initially to $\hat{C}_0$ during the first n_0 iterations before being updated. We denote by n_0 the iteration at which we start the adaptive procedure and by Δn_{adapt} the period at which the covariance is adapted. The first adaptation is based on the covariance estimation of the n_0 iterations. Then we apply a recursive formula for estimating the covariance as a function of each iteration to avoid saving all the sample values [47]. Considering that we have just sampled a new value $\mathbf{m}^{n+1}$ from the posterior distribution with the covariance estimate $\hat{C}_n$, we compute the new covariance matrix estimate $\hat{C}_{n+1}$ using the $(n+2)$ samples as follows:

$$\hat{C}'_{n+1} = \frac{n+1}{n+2}\hat{C}_n + \frac{1}{n+2}((n+1)\bar{\mathbf{m}}_n\bar{\mathbf{m}}_n^T - (n+2)\bar{\mathbf{m}}_{n+1}\bar{\mathbf{m}}_{n+1}^T + \mathbf{m}^{n+1}(\mathbf{m}^{n+1})^T), \quad (17)$$

$$\bar{\mathbf{m}}_{n+1} = \frac{1}{n+2}((n+1)\bar{\mathbf{m}}_n + \mathbf{m}^{n+1}), \quad (18)$$

$$\hat{C}_{n+1} = s_d\hat{C}'_{n+1} + s_d\epsilon I_{n_u}, \quad (19)$$

where s_d is an additional scaling parameter that depends only on the dimension n_u of the parameter space, $\epsilon > 0$ is a correcting parameter and I_{n_u} is the n_u -dimensional identity matrix. The parameter s_d is used to improve the acceptance rate in Metropolis–Hastings algorithms and it was shown in [52] that the optimal scaling for a Gaussian distribution was $s_d = 2.38^2/n_u$. We also keep this parameter as we also observed that the posterior distribution was sampled more efficiently using it. The correcting parameter $\epsilon > 0$ is used to ensure that the covariance matrix estimate will not become singular, but it can be set to zero in practice [47].

Since the update of the covariance matrix depends on previous iterations, adaptive algorithms are by essence non-Markovian and the development of such algorithms requires to show that the chain has the correct ergodic properties (converges toward the right distribution and is stationary) [47,48,53]. Such a proof for the ISDE is left as an interesting direction for future work.

Instead, we discard the samples drawn during the adaptation phase. Once the adaptation phase has provided a satisfactory estimation of the covariance matrix, we keep the latter fixed and the subsequent iterations will be used for posterior estimations. The adaptive phase can be more or less long, depending on the initial approximation of the covariance matrix, the parameters of the algorithms, and the time for the MCMC chain to learn from the whole posterior density. We denote by n_f the last iteration at which the covariance matrix is updated and before which the samples are discarded for posterior estimations. The ISDE algorithm with an initial adaptive phase is presented in Algorithm 1.

Algorithm 1: ISDE algorithm with an initial adaptive phase of the covariance matrix

```

1 Set  $\tau, f_0$ ;
2 Initialize  $\mathbf{m}^0, \hat{C} = \hat{C}_0$ ;
3 while  $n < n_{\text{iterations}}$  do
4   Sample  $\Delta \mathbf{W}^{n+1} \sim \mathcal{N}(0, \tau)$ ;
5   Solve Eqs. (13)–(14) for  $\mathbf{Q}^{n+1}, \mathbf{P}^{n+1}$ ;
6    $\mathbf{m}^{n+1} = \mathbf{Q}^{n+1}$ ;
7   Update covariance estimate from Eqs. (17)–(19);
8   if  $n \geq n_0$  and  $n \leq n_f$  and  $\text{mod}(n, \Delta n_{\text{adapt}}) = 0$  then
9     | Update  $\hat{C} = \hat{C}_{n+1}$ ;
10  end
11 end

```

5.3 Discussion on the Choice of the Sampling Algorithm

We investigated different sampling algorithms to explore efficiently the whole posterior, in particular, the RWMH, HMC and ISDE, as well as the adaptation of covariance matrix and a reparameterization of the equations. The comparison of these methods is beyond the scope of this paper, but the results are documented in [54] and they are briefly summarized below.

The comparison of the sampling of the posterior shows that a reparameterization based on the structure of the equations is an efficient method to improve the convergence of the RWMH algorithm. However, this reparameterization is particular to the present pyrolysis equations and requires the tuning of scaling factors, which can be challenging in high dimensions where the number of scaling factors increases. Therefore, gradient-informed sampling algorithms are used and it is showed that the adaptation of the covariance matrix based on the previous sample history improves by far their convergence rate, and can even fail if not performed. The adaptation of the covariance matrix only for a few iterations already improves significantly the convergence rate of both the HMC and ISDE, with similar performances for the two methods. At higher dimensions, it was observed that in the present case the HMC algorithm failed to converge if the reparameterization of the pyrolysis equations was not performed, whereas the ISDE algorithm showed a better robustness in higher dimensions. In the case where the HMC algorithm is used with the reparameterization based on the pyrolysis equations, it shows very good convergence properties, but as for the RWMH algorithm the tuning of the scaling factors can be particularly tedious.

5.4 Computational Considerations

This adaptive version of ISDE is implemented within pybitup, an in-house solver written in python. A version is made available on GitHub.[†]

[†] <https://github.com/jcoheur/pybitup/tree/v1.1>

6. RESULTS

First, in Section 6.1, different sampling strategies are applied to a two-parameter pyrolysis model in order to compare their efficiency, along with a discussion of the challenge of incorporating modeling errors. Then, in Section 6.2 we perform the parameter calibration and uncertainty quantification on a pyrolysis model for the experiment presented in Section 3.

6.1 Uncertainty Analysis for the Production of CH₄

We analyze the production of gaseous CH₄ that can be mainly characterized by a single peak in the production curve [Fig. 1 (top right)]. In order to feature the strong kinetic compensation effect that results in a posterior distribution with large scales and correlations and a complex dependence structure, we consider two uncertain parameters that are the pre-exponential factor and the activation energy, denoted by $\mathcal{A}$ and $\mathcal{E}$, respectively. Furthermore, this two-parameter model allows for the comparison of the results of the sampling methods with the direct evaluation of the posterior distribution from formula Eq. (8) with prior from Eq. (9) and likelihood function from Eq. (12).

Prior to the analysis, we perform a regularization of the initial estimation of the standard deviations s_j with $1 \leq j \leq n_t$ (to simplify the notation, we omit the single i index because we consider only CH₄) which appear in the weights w_j^{-1} . Some values of the standard deviations obtained from the experiments, denoted by $s_j^{(e)}$, have a zero value where the production rate is zero, for instance, at the beginning or at the end of the pyrolysis process, and the resulting weights are not defined properly. We thus apply the following regularization:

$$s_j = \begin{cases} s_{\min} & \text{if } s_j^{(e)} = 0, \\ s_j^{(e)} & \text{else,} \end{cases} \quad \text{with } 1 \leq j \leq n_t, \quad (20)$$

where $s_{\min} = \min_j s_j^{(e)}$ for $1 \leq j \leq n_t$ with strictly positive values of $s_j^{(e)}$.

6.1.1 Further Modifications of the Weights

We assume first that the production rate for CH₄ can be reasonably reproduced using a one-equation model that simulates the single production peak. The initial parameter guess is obtained by manually tuning the parameters and visual inspection of the simulated production rate. The two other parameters of the pyrolysis model, namely γ and r , are kept fixed to 0.0011 and 4.96, respectively. From the experiments, the time interval between two acquisitions is constant, hence, we have $\Delta t/t_f = n_t^{-1}$, and we set $\vartheta = n_t$. The calibration is then performed using 10^4 iterations with the ISDE algorithm.

The evaluations of the one-equation model at the MCMC sample values are shown in Fig. 2(a). At every temperature, the 2.5–97.5 percentile range [shaded area in Fig. 2(a)] is computed using all the model evaluations. It can be seen that the resulting range misses the maximum of production rate and favors a good fitting of the extremities of the curve, which is to us unwanted. This is due to a local maximum in the parameter posterior density that has a higher density than the one at the parameter values that provides the good fitting around the maximum production peak. The presence of this spurious local maximum in the posterior appears because of the high values of the weights at the extremities of the curves [even if those values were already regularized in Eq. (20), the values of the s_j are still low]. Any discrepancy of the simulated curve compared to the experimental data at these locations will have a significant influence on the posterior compared to the other data close to the peak, to which lower values of weights are assigned.

In order to suppress this undesirable outcome, a first possibility is to relax the weights at the extremities of the curves, that is, at low and high temperatures. Let $s_{\max} = \max_j s_j^{(e)}$ for $1 \leq j \leq n_t$. In order to reduce the influence of the small values of the weights that are away from the peak of maximum production, we assign

$$s_j = \begin{cases} cs_{\max} & \text{if } j \text{ s.t. } s_j^{(e)} \leq cs_{\max}, \\ s_j^{(e)} & \text{else,} \end{cases} \quad (21)$$




FIG. 2: Influence of the weights on the calibration of the production rate of CH₄ for the one-equation model with uncertain parameters $\mathcal{A}$ and $\mathcal{E}$. The shaded areas represent the region where 95% of the model evaluations lie and the error bars represent the values of the inverse of the weights. If no relaxation of the weights is included (a), we need to increase the number of equations in the model to better capture the maximum production peak. With the weights relaxed (b), a one-equation model is enough.

where c is a constant that needs to be fixed. The constant c is a trade-off between providing flexibility to the model by increasing the model error where it is expected to be present, and constraining the model to fit these values. The experimental data with the relaxed weights with $c = 0.2$ and the resulting model evaluations for the 10^4 samples using the ISDE algorithm are shown in Fig. 2(b). The higher values of the $s_j^{(c)}$ can be observed at both extremities of the curve. As expected, the maximum production peak is now captured by the one-equation model.

A second possibility is to improve the overall fitting by using a more flexible model for the production of CH₄. From Fig. 2(a), a first increase in the production rate of CH₄ can be observed around 575 K and a residual production can be observed above 1000 K. Thus, two pyrolysis equations are added to the previous one-equation model to improve the fitting in these temperature ranges. A manual tuning and a few pre-simulations are performed in order to find the additional parameters introduced by the three-equation model, which are then fixed to the values that are summarized in Table 1. The two parameters $\mathcal{A}_2$ and $\mathcal{E}_2$ (we will drop the indexes in the following) control the kinetics of the second equation that is mainly responsible for the maximum production peak and are therefore considered to be uncertain. Finally, the ISDE algorithm is run for 10^4 iterations in order to infer $\mathcal{A}$ and $\mathcal{E}$ and the resulting evaluations of the three-equation model are shown in Fig. 2(a). It can be seen that the three-equation model fits reasonably the production of CH₄ over the whole temperature range, which allows us to use the initial weights from Fig. 2(a).

The choice of the model and the value of the weights have thus an important influence on the result of the calibration of the model parameters. Because we mainly want to capture the main production peak, we will both favor the choice of relaxing the weights and providing flexibility to the model, as it will be done in the example of Section 6.2, where the model will feature more parameters.

TABLE 1: Nominal parameter values for the three-reaction pyrolysis model for the production of CH₄. The two uncertain parameters are the ones for which the values are inferred

Reaction	F	$\log(\mathcal{A})$	$\mathcal{E}$ (kJ mol ⁻¹)	ro_i
1	0.000369	27.83	164.55	14.37
2	0.003789	Uncertain	Uncertain	3.34
3	0.003593	27.83	284.722	67.37

6.1.2 Influence of ϑ as a Parameter for Controlling the Posterior Uncertainty

Let us now consider the three-equation pyrolysis model with the two uncertain parameters $\mathcal{A}$ and $\mathcal{E}$. The result of the propagated samples obtained with $\vartheta = n_t$ was shown in Fig. 2(a) and the corresponding computation of the posterior distribution is shown in Fig. 3. The evaluation of the posterior from Eq. (9) is evaluated on a uniform 200×200 logarithmic grid in order to improve the accuracy in the region of high probability density in the original parameter set. As expected, we can observe a strong correlation between $\mathcal{A}$ and $\mathcal{E}$ emphasized by the large values they take and a non-Gaussian shape of the posterior. Because of the difficulty of visualizing the whole posterior in the original parameter set [Fig. 3(a)], the posterior is also shown on a log-log scale in Fig. 3(b).

The results for $\vartheta = 2n_t$, $\vartheta = n_t$, and $\vartheta = 0.8n_t$ are shown in Fig. 4. As it can be seen on Fig. 4(b), the decrease of ϑ results in a posterior that is wider and thus results in more uncertainty in the parameter values and reduces the kinetic compensation effect between $\mathcal{A}$ and $\mathcal{E}$. The 95% ranges for the model evaluations are therefore larger as seen in Fig. 4(a) when ϑ decreases. The increase in the uncertainty when ϑ decreases is because ϑ appears in the numerator of the exponential term of the posterior.

As a summary, we observed that, increasing the value of theta decreases the parameter uncertainty and one can expect that it will continue to do so by increasing further the value of ϑ . The presence of modeling error may also be




FIG. 3: Bivariate posterior densities for the parameters $\mathcal{A}$ and $\mathcal{E}$ in the original parameter space (a) and logarithmic scales (b). Fourteen density contours are shown with intensity represented on the colorbars on the right of each plot.




FIG. 4: Influence of the parameter ϑ on the result of the inference for the three-equation model with two uncertain parameters $\mathcal{A}$ and $\mathcal{E}$. The production rate of CH_4 (a) and the 95% posterior density probability (b) for three different values of ϑ are shown. In (a), the shaded areas represent regions where 95% of the model evaluations lie and the error bars represent the values of the inverse of the weights. In (b), the inner part of the curves represents a region of parameter values within which 95% of the probability density lies.

partly responsible for the reduction of the uncertainty in the parameter values. Decreasing the value of θ increases the uncertainty in the parameter values. Decreasing too much for the value of ϑ may trigger solutions in the posterior that were previously far from the current maximum density, and render the sampling of this part of the posterior possible (already observed for $\vartheta = 0.6n_t$). If unwanted, such solutions can be removed, for instance, by incorporating more restrictive bounds on the parameters defined in the prior.

6.1.3 Comparison of the Sampling Strategies and Parametric Study of f_0

The ISDE algorithms with and without the adaptive phase of the covariance matrix are compared together with $\vartheta = n_t$ and again the three-equation pyrolysis model with the same two uncertain parameters $\mathcal{A}$ and $\mathcal{E}$. For the reproducibility of the path of the Markov chains, we control the initial state of the random generator by setting the seed equal to zero. The free parameters in the ISDE also need to be set. The time step τ should neither be too small, in which case the sampling is inefficient, nor be too large, in which case the numerical scheme is not stable. It is chosen by a few trial-and-error simulations and the time step is set to $\tau = 0.1$ s. The damping parameter is set to $f_0 = 4$ and its influence will be studied. The sampling algorithms are run for 1.1×10^5 iterations.

In order to assess the efficiency of the different algorithms, the ergodic mean and variance for the marginal posterior distributions as a function of the number of iterations n are computed for $\mathcal{A}$ and $\mathcal{E}$ as

$$\mu_i(n) = \frac{1}{n} \sum_{k=1}^n m_i^k, \quad (22)$$

$$\text{var}_i(n) = \frac{1}{n} \sum_{k=1}^n (m_i^k - \mu_i(n))^2. \quad (23)$$

The translated mean and normalized variance can be computed as $\tilde{\mu}_i(n) = \mu_i(n) - \bar{\mu}_i$ and $\tilde{\text{var}}_i(n) = \text{var}_i(n) / \bar{\text{var}}_i$ such that their values tend to zero and one, respectively. Good approximations of the reference mean $\bar{\mu}_i$ and reference variance $\bar{\text{var}}_i$ can be obtained from the direct evaluation of the posterior distribution, which in the case shown in Fig. 3 are equal to (for $\mathcal{A}$ and $\mathcal{E}$ in logarithmic values) $\bar{\mu}_{\mathcal{A}} = 28.53$, $\bar{\text{var}}_{\mathcal{A}} = 12.22$, $\bar{\mu}_{\mathcal{E}} = 3.74$, and $\bar{\text{var}}_{\mathcal{E}} = 0.0032$.

First, from Fig. 5, it can be noted that when using the diagonal covariance matrix estimation, the ISDE (ISDE-D) is far from convergence. However, it shows good performance when the covariance matrix estimation results from the adaptation phase (ISDE-CA), which is performed during the initial burn-in phase. During this initial phase, the algorithm is run for 10^3 iterations before the covariance matrix starts to update using the recurrence formulas from Eqs. (17)–(19). The covariance is then updated every 100 iterations until the 10^4 -th iteration, which corresponds




FIG. 5: Convergence of the mean (a) and variance (b) for $\log(\mathcal{A})$ as a function of the number of function evaluations (after the burn-in phase). ISDE with the covariance estimated from the diagonal of the inverse of the Hessian and change of variables based on the support of the posterior (ISDE-D); ISDE with the covariance resulting from the adaptation process (ISDE-CA).

to the end of the burn-in phase. This shows that the adaptive ISDE is robust even if the estimation for the initial covariance matrix was representative of the local scales and not the correlation. Adaptation may also succeed with a less informative covariance matrix, e.g., the identity matrix, but a larger adaptive phase would be required.

The 2×10^4 first sample values from the MCMC chain for $\mathcal{E}$ are shown in Fig. 6. Similar graphs are obtained for $\mathcal{A}$. The first part on the left are the samples obtained during the 10^3 first iterations of the ISDE using the fixed diagonal covariance matrix, obtained from the diagonal of the inverse of the Hessian computed at the initial parameter values. From iterations 10^3 to 10^4 , the covariance is estimated from the recurrence formula and updated in the algorithm every 10^2 iterations. The effect of the adaptation is clearly observed when compared to the first part of the chain. After 10^4 , the covariance matrix is kept fixed.

In Fig. 7, we compare the effect of the number of iterations during which the adaptation is performed on the convergence of the mean and variance. The total number of iterations performed is 1.1×10^5 for all cases. Figure 7 shows the results when the adaptation is stopped after the first 5×10^3 , 1×10^4 , and 2×10^4 iterations. Once again, the mean and variance are estimated after discarding the samples from these adaptation phases. Figure 7 also shows the result when the adaptation is not stopped and is run during all the 1.1×10^5 iterations, in which case we use all the sample history to compute the mean and variance.

It can be observed that the results converge toward zero for the mean and one for the variance in all cases. The best result is for $n_{\text{end}} = 10^4$, which motivates our choice in the previous section for stopping the adaptation at this value. The result when the adaptation is performed during all the iterations shows that the algorithm may also work with a continuous adaptation of the covariance matrix, although we do not have the theoretical support.




FIG. 6: Iterations of the Markov chain for $\mathcal{E}$ (shown every 10 iterations for clarity) obtained using the ISDE algorithm. The first (green) dotted vertical line represents the iteration at which the adaptation starts. The second (red) dotted line represents the last iteration of the adaptation. Above 10^4 iterations, the covariance is fixed to the estimated value and the MCMC chain continues.




FIG. 7: Convergence of the mean (a) and variance (b) for $\log(\mathcal{A})$ as a function of the number of function evaluations and comparison of the number of iterations in the adaptation phase for the adaptive algorithm

Similarly, the effect of the period at which the covariance matrix is updated in the algorithm (Δn_{adapt}) on the convergence rate of the mean and variance is also investigated. When the covariance matrix is effectively updated in the algorithm, it requires the computation of its inverse for computing L_M in Eq. (14). This may take some time when the dimensionality increases if the covariance matrix is adapted at every iteration. Conversely, a long adaptation period Δn_{adapt} may require a longer adaptive phase for having the benefit of adapting the covariance matrix from samples. Results were obtained for $\Delta n_{\text{adapt}} = 10$, $\Delta n_{\text{adapt}} = 100$, and $\Delta n_{\text{adapt}} = 1000$, and they showed that $\Delta n_{\text{adapt}} = 100$ was a good compromise between the accuracy of the estimation of the mean and variance and the computational time. Having a short adaptation period did not show a significant improvement in the convergence rate, and neither a large value provided bad convergence results.

Finally, we provide a numerical analysis for the influence of the damping factor f_0 on the convergence of the mean and variance for the parameter $\mathcal{A}$. The parameter f_0 introduces dissipation in the stochastic equation and will influence the convergence speed for the estimation of the mean and variance. We consider again the case with parameters from Table 1 with $\vartheta = n_t$, the covariance matrix results from the adaptation of the first 10^4 iterations and is fixed for the 10^5 iterations.

Results for the estimation of the mean and variance of $\mathcal{A}$ are shown in Fig. 8. For low values of f_0 ($f_0 = 0.01$ and $f_0 = 0.1$) we observe an erratic convergence, in particular for the variance. In this case, because the stochastic term is proportional to the square root of f_0 , it is predominant compared to the damping term. The system is closer to a Hamiltonian system and the algorithm generates highly correlated samples with rapid local variations that can be observed on the curves. For values of the order of unity (from $f_0 = 0.5$ to $f_0 = 10$) the convergence is faster, with the closest results to zero mean and unit variance achieved for $f_0 = 0.5$ and $f_0 = 1.5$. For larger values of f_0 ($f_0 = 50$), the damping factor is predominant and the convergence is again slower, but the local oscillations are not observed. Therefore, it seems that values of f_0 ranging from 0.5 to 10 can all achieve a good result for the convergence of the mean and variance, while values outside of this range should be avoided.

Almost identical graphs are obtained for the parameter $\mathcal{E}$ due to the high correlation between $\mathcal{A}$ and $\mathcal{E}$ and the results are shown in the Supplementary Material for completeness (Figs. S1–S3).

6.2 Bayesian Parameter Inference for PICA Devolatilization at High Heating Rate

In this section, we consider the whole experimental data set of Bessire and Minton [10] at the heating rate $\beta = 360$ K min⁻¹ for the pyrolysis of PICA (Fig. 1) and we perform the Bayesian calibration of the parameters. The data for the 1-propanol and 2-propanol are discarded, as they are scarce and might be subjected to high experimental errors, and therefore $n_{\text{obs}} = 12$.




FIG. 8: Convergence of the mean (a) and variance (b) for $\log(\mathcal{A})$ as a function of the number of function evaluations with different values for the damping factor f_0





6.2.1 Iterative Procedure for the Construction of the Pyrolysis Model and Posterior Analysis

There are two main issues that need to be solved: the first issue is that the structure of the pyrolysis model is unknown; the second issue is that we do not have an initial estimate of the parameters for a pyrolysis model once its structure is given.

In Torres-Herrador et al. [13], a six-equation ($N_p = 6$) pyrolysis model was identified by visual inspection of the number of maximum species production peaks. A model-fitting procedure was then performed using a multi-objective genetic algorithm to infer the values of the 35 parameters. However, by performing the Bayesian parameter inference from the model structure identified in [13], similar observations as the ones from Fig. 2(a) were obtained, that is, some production peaks were not appropriately simulated due to the influence of the weights. It does not mean that the model is wrong, but that some misfits in the model that are not accounted for in the weights degrade the overall fitting of the production peaks.

Because of the complexity of the pyrolysis model structure, identifying which values of the weights are degrading the goodness of fit is not obvious. This leads us to the need of rebuilding the pyrolysis model structure in an iterative procedure, in which models of increasing complexity are successively inferred. We start by inferring a one-equation model using a single experimental curve, as it was done for CH_4 in Section 6.1. Calibrating this model is simpler as it is defined by fewer parameters and the initial guess can be easily obtained, for instance, by running the sampling algorithms for a few iterations to converge toward the region of higher probability. Once we have a satisfactory solution for the parameters $\mathcal{A}_1$, $\mathcal{E}_1$, r_1 , and γ_{11} , we add one experimental curve in the set of calibration data (we add one species j). If a new reaction maximum is observed in the additional curve, we increase the number of equations by one and we need to calibrate the new set of parameters $\mathcal{A}_2$, $\mathcal{E}_2$, r_2 , and γ_{2j} together with the previous set of parameters for which we have an initial guess. If a reaction already exists at the new peak, then we just need to calibrate the new mass fractions γ_{2j} . This procedure is repeated until the model includes all the species from the data set. As already mentioned in Section 6.1, issues may arise when the inverses of the weights approach zero where the production rate is zero. Therefore, we apply once again the regularization of the weights from Eq. (20) and the relaxation of the weights at the extremities from Eq. (21) with $c = 0.6$.

The following six-equation pyrolysis model is found:



where R refers to the unknown reactants of the reaction. Using the model from Eqs. (24)–(29), the number of parameters to infer is equal to 48. The prior is the product distribution from the marginal priors. The marginal priors for pre-exponential factors $\mathcal{A}_i$ are semi-infinite uniform priors. The activation energies $\mathcal{E}_i$ are bounded in given ranges, provided in Table 2, in order to separate the different peaks of production. Evidence from some of our previous simulations shows that the solution can converge slowly or can sample unwanted regions in the posterior if the activation energies are not limited within a given range. The bounds from Table 2 on the activation energies were obtained by defining small intervals around the initial guesses, then by progressively increasing them until no truncation of the marginal posteriors for the activation energies were observed *a posteriori*. Activation energy ranges overlap for the reaction close to the main production peak, but they are clearly separated from the low and high temperature reactions. The values of the reaction orders are also limited in the range $]1, 10]$ and all mass fractions range between 0 and 1.

TABLE 2: Values of the bounds for the activation energies


	$\mathcal{E}_0$ (kJ mol ⁻¹)	$\mathcal{E}_1$ (kJ mol ⁻¹)	$\mathcal{E}_2$ (kJ mol ⁻¹)	$\mathcal{E}_3$ (kJ mol ⁻¹)	$\mathcal{E}_4$ (kJ mol ⁻¹)	$\mathcal{E}_5$ (kJ mol ⁻¹)
L_b	65	165	165	160	230	400
U_b	85	185	195	190	260	410

The ISDE algorithm is run for 10^6 iterations with $\tau = 0.002$ s and $f_0 = 4.0$. The initial covariance matrix is estimated using the diagonal matrix of the inverse of the Hessian computed at the initial guess value. Then, the covariance is updated every 100 iterations from $n_0 = 10^4$ to $n_f = 10^5$. Due to the higher dimensionality, the algorithm needs more iterations to learn the local scales and correlations and the adaptation phase is therefore longer. The solution to the pyrolysis model and its derivatives with respect to the parameters are all computed analytically [12].

The results for $\vartheta = n_t$ are shown in Figs. 9–11. In Fig. 9, the Markov chains are provided for the kinetic parameters of the reaction R_1 and the corresponding mass fraction of CH_4 . The 10^5 first iterations from the adaptation phase are discarded. The other chains representing the samples drawn from the 48-dimensional posterior show similar behavior of mixing but are not shown here.

In Fig. 10 the bivariate and marginal posteriors for the reaction R_1 are shown. In the bivariate posterior for $\mathcal{A}_1$ (in log) and $\mathcal{E}_1$, the correlation between the two parameters can be observed. This correlation is reduced compared to the previous section because of the restricted range on $\mathcal{E}$ and is necessary to limit the potential interactions between the different reactions. The region where $\mathcal{A}$ varies is therefore reduced. Regarding the marginal parameter distributions, they feature a single mode distribution.

The production rate curves at the sample values are shown in Fig. 11. We represent the regions where 95% of the model evaluations lie with the mean value in between. From this figure, we can directly assess the goodness

**FIG. 9:** MCMC chains for selected parameters after a burn-in of the firsts 10^5 samples




FIG. 10: Marginal and bivariate posterior distributions for selected parameters. $\mathcal{A}_1$ is represented on a log scale (unit of $\mathcal{A}_1$ is in s^{-1}). The probability densities have the inverse units of their variable(s).

of fit of the model, and we can conclude that the model structure chosen provides an overall fit that is good. This model underestimates a bit the production of trimethyl phenol, xylene, and toluene and overestimates slightly the maximum of CO_2 . In general, the model represents well the observations within the range [500, 1000] K. Above 1000 K, the production almost reaches a plateau for CO , CO_2 , H_2O , and H_2 , which is difficult to account for using the multi-component approach from Eq. (4).




FIG. 11: Mass yields as a function of temperature for each species obtained after the propagation of the MCMC samples with $\vartheta = n_t$. The symbols with the error bars represent the experimental results from [10]. The continuous lines are the results of the propagation for the maximum likelihood estimate and the shaded areas represent all the possible outcomes of the propagation using every sample from the posterior distribution.

6.2.2 Discussion of the Choice of the Value of ϑ

Selecting the value of ϑ based on the overall goodness of fit not obvious as there is significant model inadequacy at low and high temperatures. However, we can compare the fitting by visual inspection based on the results in the range from [500, 1000] K, which is fairly good, and we can compare the overlap of the computed intervals with the experimental intervals. From that point of view, the computed intervals from Fig. 11 might still be underestimating the uncertainty ranges and the posterior distribution be over-confident.

In order to increase the uncertainty in the simulated curves and be more in agreement with the experimental ranges, the inference is also performed with $\vartheta = 0.7n_t$. This value was obtained after a few trials, and the resulting

simulated curves are shown in Fig. 12. The increase in the uncertainty range can be clearly noticed, which is due to the posterior distribution that is less concentrated, showing that the parameter ϑ allows to control the overall uncertainty range on the predictions. The maximum *a posteriori* values for the whole set of parameters are provided in the Supplementary Material (Table S1). For instance, the graph for H₂O and Cresol better incorporates the uncertainty bars in Fig. 12, whereas the simulated uncertainty ranges were narrower in Fig. 11 compared to the experimental data.

Yet, some observations are still outside of the computed uncertainty ranges, particularly below 500 K and above 1000 K. This issue highlights limitations of the approach of seeking to incorporate the effect of modeling errors as an uncertainty in the model parameters. If unmodeled or poorly represented physical behavior results in model




FIG. 12: Mass yields as a function of temperature for each species obtained after the propagation of the MCMC samples with $\vartheta = 0.7n_t$. The symbols with the error bars represent the experimental results from [10]. The continuous lines are the results of the propagation for the maximum likelihood estimate and the shaded areas represent all the possible outcomes of the propagation using every sample from the posterior distribution.

predictions that are not sufficiently close to the experimental data, then such model inadequacy cannot be compensated by introducing uncertainty in the model parameters.

6.2.3 Summary of the Procedure for Pyrolysis Model Parameter Inference

We summarize in this section the procedure used in this work for the inference of the pyrolysis model parameters and we identify perspectives for improvements. The procedure is expected to be similar in the case of other experimental data of the same kind, namely, time-dependent data defined by a set of uncoupled ordinary differential equations.

1. Select an appropriate pyrolysis model structure. The iterative procedure described above can be followed at first. As a future work, an automated procedure for building such reaction model would be highly valuable.
2. Select a value for ϑ . Its value depends on the capability of the pyrolysis model to simulate the experimental curves. As a starting value, select ϑ equal to the number of sample points, such that $\vartheta/n_t = 1$ and the likelihood is similar to a Gaussian function. Next, adjust its value based on the observed noise by visual inspection (increasing its value results in reducing the uncertainty and inversely). As a future work, it would be valuable to introduce a measure for a more robust criterion.
3. Compute the weights from Eq. (3), regularize their value using Eq. (20) to avoid regions where the production rate is zero and relax their value using Eq. (21) to avoid the influence of data points at the extremities of the production curve. For the latter, two values of the parameter c were used, $c = 0.2$ ($n_u = 2$) and $c = 0.6$ ($n_u = 48$), respectively.
4. Define the prior function bounds, in particular for the activation energies, which need to be tuned. As a rule of thumb, activation energies were found to range from 10 to 30 kJ mol⁻¹ around the initial guess.
5. Define the ISDE numerical parameters f_0 and τ . Values of f_0 ranging from 0.5 to 10 were found adequate in this work and are confirmed in the literature. A common choice for the step size is $\tau\omega = 1/2$ ([55]), where ω is the largest eigenfrequency (square root of an eigenvalue) of the Hessian matrix $\nabla^2\mathcal{U}(\mathbf{q})$. The effect of the change of scales is that of a clustering of these eigenvalues close to unity ([41]); yet we only have an approximation of the covariance matrix and in practical applications one can still be far from unity, even after adapting the value of the covariance matrix. In this work, in the 2-dimensional example the largest eigenfrequency (after the change of scales) is $\mathcal{O}(10^1)$ and in the 48-dimensional case is $\mathcal{O}(10^3)$. These orders of magnitudes can be slightly increased to provide a faster convergence rate ([56]), but not too much to avoid numerical errors. The final values used were in this case $\tau = 0.1$ s for $n_u = 2$ and $\tau = 0.002$ s for $n_u = 48$.

6.3 Propagation of the Uncertainty in Pyrolysis Model Parameters for a Martian Atmospheric Re-Entry

We will now demonstrate the use of the identified pyrolysis model in a numerical simulation of heat shield recession during a Martian entry.

The porous-material analysis toolbox based on OpenFOAM (PATO) [57], developed at NASA Ames Research Center, is used to perform 2D axisymmetric simulations of the full heat shield during an atmospheric entry and to analyze the recession of the heat-shield surface. PATO solves the conservation equations of solid mass, gas mass, gas momentum, and total energy for a porous reactive material using a volume-averaged formulation. Assuming local thermal equilibrium at the pore scales, material pyrolysis decomposition, production of gases from the pyrolysis reactions and ablation, and homogeneous chemical reactions in the gas phase are modeled. The chemistry of the gas phase is assumed to be at equilibrium, and only the mass fractions of the elements produced by the pyrolysis have to be provided, which are carbon, hydrogen, and oxygen. The equilibrium chemistry assumption allows to reduce the number of independent variables that are solved by the numerical solver compared to a finite-rate chemistry model and reduces the computational time of the simulations. The open source library Mutation++, developed at the von Karman Institute for Fluid Dynamics, is used to compute equilibrium chemistry compositions and thermodynamic

and transport properties [58]. Gas-surface interactions are modeled using equilibrium chemistry models that are preferred for design due to the lack of reliability of available finite-rate chemistry models and data [59]. As in [59], the Theoretical Ablative Composite for Open Testing (TACOT) database, based on the PICA material and defined in [60], is used for the porous material properties except from the pyrolysis reactions and model parameters.

We consider a re-entry in the atmosphere of Mars and a thermal protection system similar to the Mars Science Laboratory (MSL) heat shield. Deterministic simulations of an atmospheric entry on Mars using PATO were already performed in [59] and we consider a similar numerical setup. The MSL heat shield is a rigid 4.5 m diameter spherically blunted half-angle cone made up of 113 PICA tiles [61,62]. The mesh, shown in Fig. 13, is composed of 103 K cells with 64 cells in depth and is a monolithic simplification of the tiled configuration of the manufactured heat shield. The physical time interval takes place between 50 and 100 s along the re-entry path, during which the pressure peak and the peak heating on the heat shield occur. The aerothermal boundary conditions are provided from hypersonic CFD simulations, which are considered to be deterministic [59].

From the Markov chains computed in Section 6.2 for each case of the free parameter ϑ , we select 500 evenly spaced samples to be propagated. Selecting all the samples would not be realistic in terms of computational time and would result in a poor convergence of the statistics of interests due to the correlation between the samples, inherent to the MCMC procedure. The values for the γ_{ij} provided in Table S1 need to be slightly adapted because of the difference between the pyrolysis model implemented in PATO (see [63]) and the model described in Eqs. (5) and (6). In PATO, the production rates are computed with respect to the mass (or density in PATO) of the initial resin content of the material, while in Eqs. (5) and (6), the production is proportional to the total mass of the sample. Therefore, each γ_{ij} is divided by a factor equal to 0.43 which represents the ratio of the mass (density) of the initial resin with the total mass (density) of the sample.

Each single PATO simulation runs on 34 cores from the NIC5 cluster from the University of Liège, which features AMD EPYC Rome 7524 32-Core cpus (2.9 GHz), and each simulation lasts around 1 h 20 min. Allocating more than 34 cores does not further improve significantly the computational time.

The results of the simulations are probed at the four locations on the heat shield given in Table 3. To represent the uncertainty on the quantity of interest, 95% probability regions are computed using the results of the 500 samples. The recession of the surface of the heat shield at the probe locations is shown in Fig. 14.

The recession as a function of time [Fig. 14(a)] for $\vartheta = 0.7n_t$ shows an increase of the uncertainty as time increases. At 100 s after the beginning of the re-entry, the median recession is around 0.922 mm at the nose region, 0.961 mm at the wind side, 0.531 mm at the middle side, and 0.195 mm at the lee side, and the results for the median are similar for $\vartheta = n_t$ (not shown on the graph for clarity). In Fig. 14(b), we provide the estimation of the probability density functions at $t = 100$ s using a Gaussian kernel-density estimation method for the two values of ϑ . The PDFs feature one mode and are almost centered around the median. For $\vartheta = 0.7n_t$, the supports of the PDFs are slightly wider, showing that the uncertainty is larger. Indeed, the thickness of the 95% probability regions are around 0.038




FIG. 13: Illustration of the mesh for the MSL heat shield. View from the front (left) and the side (right).

TABLE 3: Heat-shield probe coordinates

Name	x (m)	y (m)	z (m)
Nose	0.0	0.0	0.0
Wind side	0.72346	-0.0025	-2.184
Lee side	0.72346	0.0025	2.184
Middle side	0.72346	2.184	0.0




FIG. 14: Results of the propagation of uncertainty in pyrolysis model parameters on the recession of the material at four different locations on the heat shield. (a) Recession as a function of time shown for $\vartheta = 0.7n_t$, and (b) estimation of the probability density functions at $t = 100$ s for the nose region (continuous line) and wind side (dashed line) for $\vartheta = n_t$ (no markers) and $\vartheta = 0.7n_t$ (circle markers).

mm (the ratio of this thickness and the median is equal to 4.14%) and 0.034 mm (3.58%) at the nose and the wind side, respectively, while for $\vartheta = n_t$ these values are equal to 0.03 mm (3.27%) and 0.026 mm (2.7%), respectively. At the lee side, the uncertainty is relatively higher for both values of ϑ , which is around 0.019 mm (9.70%) and 0.023 mm (11.86%) for $\vartheta = n_t$ and $\vartheta = 0.7n_t$, respectively.

These results show that the value of the free parameter ϑ does not affect significantly the median and the mean of the recession of the surface of the heat shield. However, as expected, decreasing the value of ϑ increases the resulting uncertainty on the recession. In general, we note that the level of uncertainty on the recession due to pyrolysis model parameters is relatively low. Recession is related to pyrolysis model parameters only through effective material properties, themselves related to the reaction rates through the porosity of the material. Hence, other parameters like the intrinsic material thermal conductivity, the heat capacity, or the initial porosity, are foreseen to increase further the general uncertainty level on the recession.

The current pyrolysis model was calibrated at a single heating rate (366 K min^{-1}) and more complex models could be developed in order to account for the various heating rates encountered over the heat shield [19]. The model error on the pyrolysis model at 366 K min^{-1} was partly included in the input parameters by relaxing the weights in the likelihood function, but the error due to the extrapolation of this model to other heating rates was not accounted, which could also have an influence on the uncertainty on recession.

7. CONCLUSIONS

This work presents the characterization of input parameter uncertainty using new experimental data for thermal protection material decomposition, as well as the propagation of the uncertainty for an atmospheric entry problem. A

Bayesian approach involving an approximate likelihood function is applied to infer the model parameters of chemical reactions relevant to the pyrolysis. The parameters of the model feature strong correlations and the Bayesian posterior is high-dimensional, and, in order to sample efficiently from the posterior, it is proposed to use a gradient-informed sampling algorithm based on the solution of an Itô stochastic differential equation, along with a method for the adaptation of the covariance matrix.

The results show that the adaptation of the covariance matrix based on the previous sample history improves the convergence rate of the sampling algorithm, even with a few samples, as compared to when there is no adaptation. A six-equation pyrolysis model is found by iteratively increasing the number of equations in the model until the experimental data are sufficiently accurately reproduced. The resulting uncertainty-quantified model allows us to reproduce the data set together with the estimated observed uncertainty ranges, although with some exceptions at low and high temperatures where model inadequacy is too significant due to limitations in the ability of the model to represent the physics. Results are shown for two values of the free parameter ϑ , that is introduced in the approximate likelihood to control the concentration of the posterior distribution. The choice of a value of $\vartheta = 0.7n_t$ is motivated by the quality of the overlapping of the uncertainty intervals and the observations.

The impact of the uncertainty in the pyrolysis model parameters on the recession of the surface of an ablating TPS is finally assessed for a Mars atmospheric entry with the configuration of the Mars Science Laboratory mission heat shield. The propagated uncertainty (2σ) on the recession is equal to 0.922 ± 0.019 mm (with 4.14% the ratio of the thickness of the 95% probability interval and the median) at the nose and up to 11.85%, 0.195 ± 0.023 mm, at the lee side, with the resulting distributions that are almost Gaussian. The comparison of the values obtained using either $\vartheta = n_t$ or $\vartheta = 0.7n_t$ shows that the resulting uncertainty range is about 1% higher for the latter, thus providing more conservative margins on the recession.

As future work, it would be interesting to study, in order for the iterations in the adaptive phase of the Markov chains to be used in the posterior distribution, the ergodicity of the chains generated by the gradient-informed algorithm with the adaptive procedure. Secondly, the evaluations of the model and its gradients are not always available analytically and may be computationally expensive, in particular in high dimension. Efficient gradient computations could be made, for instance, by using automatic differentiation [64] or adjoint-based methods [65,66]. Methods that account for the discretization errors, such as the multilevel Monte Carlo method, could also be considered for the estimation of the model [67–70]. Thirdly, building a suitable pyrolysis model and obtaining an initial guess was found to be particularly difficult and the automatization of this procedure would be valuable, for instance, by evaluating a set of possible models with a genetic optimization (as in [13]) for finding the best initial guess and applying model inference criteria to evaluate the model plausibility [71]. The selection of the model could then be performed using methods for model comparison and model selection, for instance, using Akaike's criterion for statistical model inference ([71]) or in a Bayesian framework, e.g., see in [71–73] or in [74]. Building an unknown reaction model in a Bayesian inference framework starting from a general reaction model as performed in [75] is also an interesting research direction to improve the model inference performed in this work. Finally, the approach of encapsulating all the uncertainties in the model parameters has its limitations. Indeed, the reference deterministic model needs to be close enough to the observations, otherwise model discrepancy cannot be represented only by accommodating uncertainty in the model parameters. Further studies could include more sophisticated mechanisms to simulate the pyrolysis reactions [19], and alternative formulations for the kinetic rates [76], which could help to account for the plateau in the production curves, could be sought.

ACKNOWLEDGMENTS

The work of J. Coheur is supported by the Fund for Research Training in Industry and Agriculture (FRIA) 1E05418F provided by the Belgian Fund for Scientific Research (F.R.S.-FNRS) and by the Wallonie-Bruxelles International (WBI) fellowship for excellence (WBI.World). J. Coheur is indebted to Dr. Nagi Mansour and to the NASA Ames Research Center Advanced Supercomputing Division, where part of this work was carried out, and to J. Meurisse for the discussions and for generating the meshes for the MSL simulations. The authors would like to acknowledge K. Bulthuis for a thorough reading of the full paper and discussions, and F. Torres-Herrador for the review of parts concerning the pyrolysis modeling.

REFERENCES

- Williams, S.D. and Curry, D.M., Thermal Protection Materials: Thermophysical Property Data, NASA Reference Publication RP-1289, 1992.
- Stackpoole, M., Sepka, S., Cozmuta, I., and Kontinos, D., Post-Flight Evaluation of Stardus Sample Return Capsule Forebody Heatshield Material, *J. Thermophys. Heat Transf.*, **24**(4):694–707, 2010.
- Wright, M., Cozmuta, I., Laub, B., Chen, Y.K., and Wilcoxson, W.H., Defining Ablative Thermal Protection System Margins for Planetary Entry Vehicles, *42nd AIAA Thermophysics Conf.*, American Institute of Aeronautics and Astronautics, 2011.
- Seedhouse, E., Dragon Design, Development, and Test, in *SpaceX's Dragon: America's Next Generation Spacecraft*, Chichester, UK: Springer, pp. 23–44, 2016.
- Natali, M., Puri, I., Rallini, M., Kenny, J., and Torre, L., Ablation Modeling of State of the Art EPDM Based Elastomeric Heat Shielding Materials for Solid Rocket Motors, *Computat. Mater. Sci.*, **111**:460–480, 2016.
- Reimer, T., Zuber, C., Rieser, J., and Rothermel, T., Determination of the Mechanical Properties of the Lightweight Ablative Material Zuram, in *Ceramic Transactions Series*, Hoboken, NJ: Wiley, pp. 311–326, 2018.
- Duffa, G., *Ablative Thermal Protection Systems Modeling*, Reston, VA: AIAA, Inc., 2013.
- Wong, H.W., Peck, J., Assif, J., Panerai, F., Lachaud, J., and Mansour, N.N., Detailed Analysis of Species Production from the Pyrolysis of the Phenolic Impregnated Carbon Ablator, *J. Anal. Appl. Pyrolys.*, **122**:258–267, 2016.
- Torres-Herrador, F., Eschenbacher, A., Coheur, J., Blondeau, J., Magin, T.E., and Geem, K.M.V., Decomposition of Carbon/Phenolic Composites for Aerospace Heatshields: Detailed Speciation of Phenolic Resin Pyrolysis Products, *Aerospace Sci. Technol.*, **119**:107079–107089, 2021.
- Bessire, B.K. and Minton, T.K., Decomposition of Phenolic Impregnated Carbon Ablator (PICA) as a Function of Temperature and Heating Rate, *ACS Appl. Mater. Int.*, **9**(25):21422–21437, 2017.
- Vyazovkin, S., Burnham, A.K., Criado, J.M., Pérez-Maqueda, L.A., Popescu, C., and Sbirrazzuoli, N., ICTAC Kinetics Committee Recommendations for Performing Kinetic Computations on Thermal Analysis Data, *Thermochim. Acta*, **520**(1-2):1–19, 2011.
- Coheur, J., Torres-Herrador, F., Chatelain, P., Mansour, N.N., Magin, T.E., and Arnst, M., Analytical Solution for Multi-Component Pyrolysis Simulations of Thermal Protection Materials, *J. Mater. Sci.*, **56**:6845–6860, 2021.
- Torres-Herrador, F., Meurisse, J.B., Panerai, F., Blondeau, J., Lachaud, J., Bessire, B.K., Magin, T.E., and Mansour, N.N., A High Heating Rate Pyrolysis Model for the Phenolic Impregnated Carbon Ablator (PICA) Based on Mass Spectroscopy Experiments, *J. Anal. Appl. Pyrol.*, **141**:104625–104635, 2019.
- Najm, H., Berry, R., Safta, C., Sargsyan, K., and Debusschere, B., Data-Free Inference of Uncertain Parameters in Chemical Models, *Int. J. Uncertainty Quantif.*, **4**(2):111–132, 2014.
- Cheung, S.H., Miki, K., Prudencio, E., and Simmons, C., Uncertainty Quantification and Robust Predictive System Analysis for High Temperature Kinetics of HCN/O₂/Ar Mixture, *Chem. Phys.*, **475**:136–152, 2016.
- Khalil, M. and Najm, H.N., Probabilistic Inference of Reaction Rate Parameters from Summary Statistics, *Combust. Theory Model.*, **22**(4):635–665, 2018.
- Urzay, J., Kseib, N., Constantine, P.G., Davidson, D.F., and Iaccarino, G., Uncertainty-Quantifying Models for Chemical-Kinetic Rates, *Center for Turbulence Research Annual Briefs*, 2012.
- Bruns, M.C., Inferring and Propagating Kinetic Parameter Uncertainty for Condensed Phase Burning Models, *Fire Technol.*, **52**(1):93–120, 2015.
- Torres-Herrador, F., Coheur, J., Panerai, F., Magin, T.E., Arnst, M., Mansour, N.N., and Blondeau, J., Competitive Kinetic Model for the Pyrolysis of the Phenolic Impregnated Carbon Ablator, *Aerospace Sci. Technol.*, **100**:105784, 2019.
- Beck, J.L. and Katafygiotis, L.S., Updating Models and Their Uncertainties. I: Bayesian Statistical Framework, *J. Eng. Mech.*, **124**(4):455–461, 1998.
- Kennedy, M.C. and O'Hagan, A., Bayesian Calibration of Computer Models, *J. R. Stat. Soc.*, **63**(3):425–464, 2001.
- Koga, N., A Review of the Mutual Dependence of Arrhenius Parameters Evaluated by the Thermoanalytical Study of Solid-State Reactions: The Kinetic Compensation Effect, *Thermochim. Acta*, **244**:1–20, 1994.
- Galwey, A.K. and Brown, M.E., *Thermal Decomposition of Ionic Solids*, Amsterdam, The Netherlands: Elsevier, 1999.
- Rodionova, O.E. and Pomerantsev, A.L., Estimating the Parameters of the Arrhenius Equation, *Kinet. Catal.*, **46**(3):305–308, 2005.

25. Pomerantsev, A.L., Kutsenova, A.V., and Rodionova, O.Y., Kinetic Analysis of Non-Isothermal Solid-State Reactions: Multi-Stage Modeling without Assumptions in the Reaction Mechanism, *Phys. Chem. Chem. Phys.*, **19**(5):3606–3615, 2017.
26. Soize, C., Construction of Probability Distributions in High Dimension Using the Maximum Entropy Principle: Applications to Stochastic Processes, Random Fields and Random Matrices, *Int. J. Numer. Methods Eng.*, **76**(10):1583–1611, 2008.
27. Soize, C., Polynomial Chaos Expansion of a Multimodal Random Vector, *SIAM/ASA J. Uncertainty Quantif.*, **3**(1):34–60, 2015.
28. Soize, C., *Uncertainty Quantification*, Cham, Switzerland: Springer International Publishing, 2017.
29. Kolda, T.G. and Bader, B.W., Tensor Decompositions and Applications, *SIAM Rev.*, **51**(3):455–500, 2009.
30. Wright, M.J., Beck, R.A.S., Edquist, K.T., Driver, D., Sepka, S.A., Slimko, E.M., and Willcockson, W.H., Sizing and Margins Assessment of Mars Science Laboratory Aeroshell Thermal Protection System, *J. Spacecraft Rockets*, **51**(4):1125–1138, 2014.
31. Kendall, R.M., Barlett, E.P., Rindal, R.A., and Moyer, C.B., An Analysis of the Coupled Chemically Reacting Boundary Layer and Charring Ablator: Part I, Tech. Rep., NASA CR-1060, 1968.
32. Chen, Y.K. and Milos, F.S., Ablation and Thermal Response Program for Spacecraft Heatshield Analysis, *J. Spacecraft Rockets*, **36**(3):475–483, 1999.
33. Sykes, G.F., Decomposition Characteristics of a Char-Forming Phenolic Polymer Used for Ablative Composites, Tech. Rep., NASA TN D-3810, National Aeronautics and Space Administration, 1967.
34. Goldstein, H.E., Pyrolysis Kinetics of Nylon 6-6, Phenolic Resin, and Their Composites, *J. Macromolec. Sci.*, **3**(4):649–673, 1969.
35. Trick, K.A. and Saliba, T.E., Mechanisms of the Pyrolysis of Phenolic Resin in a Carbon/Phenolic Composite, *Carbon*, **33**(11):1509–1515, 1995.
36. Trick, K.A., Saliba, T.E., and Sandhu, S.S., A Kinetic Model of the Pyrolysis of Phenolic Resin in a Carbon/Phenolic Composite, *Carbon*, **35**(3):393–401, 1997.
37. Wong, H.W., Peck, J., Bonomi, R.E., Assif, J., Panerai, F., Reinisch, G., Lachaud, J., and Mansour, N.N., Quantitative Determination of Species Production from Phenol-Formaldehyde Resin Pyrolysis, *Polymer Degradat. Stab.*, **112**:122–131, 2015.
38. Bessire, B.K., Lahankar, S.A., and Minton, T.K., Pyrolysis of Phenolic Impregnated Carbon Ablator (PICA), *ACS Appl. Mater. Inter.*, **7**(3):1383–1395, 2015.
39. Lachaud, J., Magin, T.E., Cozmuta, I., and Mansour, N.N., A Short Review of Ablative-Material Response Models and Simulation Tools, *7th European Symposium on Aerothermodynamics*, L. Ouwehand, Ed., Brugge, Belgium, ESTEC-ESA, SP-692, pp. 91–98, 2011.
40. Arnst, M., Álvarez, B.A., Ponthot, J.P., and Boman, R., Itô-SDE MCMC Method for Bayesian Characterization of Errors Associated with Data Limitations in Stochastic Expansion Methods for Uncertainty Quantification, *J. Comput. Phys.*, **349**:59–79, 2017.
41. Arnst, M. and Soize, C., Identification and Sampling of Bayesian Posteriors of High-Dimensional Symmetric Positive-Definite Matrices for Data-Driven Updating of Computational Models, *Comput. Methods Appl. Mech. Eng.*, **352**:300–323, 2019.
42. Girolami, M. and Calderhead, B., Riemann Manifold Langevin and Hamiltonian Monte Carlo Methods, *J. R. Stat. Soc.*, **73**(2):123–214, 2011.
43. Law, K., Proposals Which Speed Up Function-Space MCMC, *J. Comput. Appl. Math.*, **262**:127–138, 2014.
44. Hoffman, M.D. and Gelman, A., The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.*, **15**:1593–1623, 2014.
45. Marshall, T. and Roberts, G., An Adaptive Approach to Langevin MCMC, *Stat. Comput.*, **22**(5):1041–1057, 2011.
46. Atchadé, Y.F., An Adaptive Version for the Metropolis Adjusted Langevin Algorithm with a Truncated Drift, *Methodol. Comput. Appl. Probab.*, **8**(2):235–254, 2006.
47. Haario, H., Saksman, E., and Tamminen, J., An Adaptive Metropolis Algorithm, *Bernoulli*, **7**(2):223–242, 2001.
48. Atchadé, Y.F. and Rosenthal, J.S., On Adaptive Markov Chain Monte Carlo Algorithms, *Bernoulli*, **11**(5):815–828, 2005.
49. Andrieu, C. and Thoms, J., A Tutorial on Adaptive MCMC, *Stat. Comput.*, **18**:343–374, 2008.
50. Roberts, G.O. and Rosenthal, J.S., Examples of Adaptive MCMC, *J. Comput. Graph. Stat.*, **18**(2):349–367, 2009.
51. Rosenthal, J.S., *Optimal Proposal Distributions and Adaptive MCMC*, Boca Raton, FL: Chapman and Hall, 2010.

52. Gelman, A.G., Roberts, G.O., and Gilks, W.R., *Efficient Metropolis Jumping Rules*, Oxford, UK: Oxford University Press, pp. 599–608, 1996.
53. Andrieu, C. and Moulines, É., On the Ergodicity Properties of Some Adaptive MCMC Algorithms, *Annals Appl. Probab.*, **16**(3):1462–1505, 2006.
54. Coheur, Uncertainty Quantification of Aerothermal Flow-Material Simulations of Low-Density Ablative Thermal Protection Systems, PhD, University of Liege, University of Louvain, 2021.
55. Hairer, E., Lubich, C., and Wanner, G., *Geometric Numerical Integration*, Berlin, Germany: Springer-Verlag, 2006.
56. Guillemot, J. and Soize, C., Itô SDE-Based Generator for a Class of Non-Gaussian Vector-Valued Random Fields in Uncertainty Quantification, *SIAM J. Sci. Comput.*, **36**(6):A2763–A2786, 2014.
57. Lachaud, J. and Mansour, N.N., Porous-Material Analysis Toolbox Based on OpenFOAM and Applications, *J. Thermophys. Heat Transf.*, **28**(2):191–202, 2014.
58. Scoggins, J.B., Leroy, V., Bellas-Chatzigeorgis, G., Dias, B., and Magin, T.E., Mutation++: MULTicomponent Thermodynamic and Transport Properties for IONized Gases in C++, *SoftwareX*, **12**:100575–100583, 2020.
59. Meurisse, J., Lachaud, J., Panerai, F., Tang, C., and Mansour, N.N., Multidimensional Material Response Simulations of a Full-Scale Tiled Ablative Heatshield, *Aerospace Sci. Technol.*, **76**:497–511, 2018.
60. Lachaud, J., Martin, A., Van Eekelen, T., and Cozmuta, I., Ablationtest-Caseseries#2, *5th Ablation Workshop*, Lexington, Kentucky, KY, pp. 10–18, 2012.
61. Edquist, K., Dyakonov, A., Wright, M., and Tang, C., Aerothermodynamic Design of the Mars Science Laboratory Heatshield, *41st AIAA Thermophysics Conf.*, AIAA, 2009.
62. White, T.R., Mahzari, M., Bose, D., and Santos, J.A., Post-Flight Analysis of Mars Science Laboratory's Entry Aerothermal Environment and Thermal Protection System Response, *44th AIAA Thermophysics Conf.*, AIAA, 2013.
63. Lachaud, J., Scoggins, J., Magin, T., Meyer, M., and Mansour, N., A Generic Local Thermal Equilibrium Model for Porous Reactive Materials Submitted to High Temperatures, *Int. J. Heat Mass Transf.*, **108**:1406–1417, 2017.
64. Griewank, A. and Walther, A., *Evaluating Derivatives*, Philadelphia, PA: SIAM, 2008.
65. Cao, Y., Li, S., Petzold, L., and Serban, R., Adjoint Sensitivity Analysis for Differential-Algebraic Equations: The Adjoint DAE System and Its Numerical Solution, *SIAM J. Sci. Comput.*, **24**(3):21, 2003.
66. Bosco, A., Bayesian Inference for the Identification of Model Parameters in Atmospheric Entry Problems, Master's, University of Liège, 2019.
67. Dodwell, T.J., Ketelsen, C., Scheichl, R., and Teckentrup, A.L., A Hierarchical Multilevel Markov Chain Monte Carlo Algorithm with Applications to Uncertainty Quantification in Subsurface Flow, *SIAM/ASA J. Uncertainty Quantif.*, **3**(1):1075–1108, 2015.
68. Beskos, A., Jasra, A., Law, K., Tempone, R., and Zhou, Y., Multilevel Sequential Monte Carlo Samplers, *Stochastic Proc. Their Appl.*, **127**(5):1417–1440, 2017.
69. Heng, J., Jasra, A., Law, K.J.H., and Tarakanov, A., On Unbiased Estimation for Discretized Models, *Stat. Comput.*, arXiv:2102.12230, 2021.
70. Jasra, A., Law, K.J.H., and Lu, D., Unbiased Estimation of the Gradient of the Log-Likelihood in Inverse Problems, *Stat. Comput.*, **31**(3):21, 2021.
71. Burnham, K.P. and Anderson, D.R., *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*, 2nd ed., New York: Springer, 2002.
72. MacKay, D., *Information Theory, Inference, and Learning Algorithms*, Cambridge, UK: Cambridge University Press, 2003.
73. Link, W.A. and Barker, R.J., Model Weights and the Foundations of Multimodel Inference, *Ecology*, **87**(10):2626–2635, 2006.
74. Strong, M. and Oakley, J.E., When Is a Model Good Enough? Deriving the Expected Value of Model Improvement via Specifying Internal Model Discrepancies, *SIAM/ASA J. Uncertainty Quantif.*, **2**(1):106–125, 2014.
75. Galagali, N. and Marzouk, Y.M., Bayesian Inference of Chemical Kinetic Models from Proposed Reactions, *Chem. Eng. Sci.*, **123**:170–190, 2015.
76. Málek, J. and Criado, J., Empirical Kinetic Models in Thermal Analysis, *Thermochim. Acta*, **203**:25–30, 1992.