Fire Safety Journal 144 (2024) 104116

Contents lists available at ScienceDirect

Fire Safety Journal

journal homepage: www.elsevier.com/locate/firesaf

Sensitivity analysis for an effective transfer of estimated material properties from cone calorimeter to horizontal flame spread simulations

Tássia L.S. Quaresma [a], Tristan Hehnen [b], Lukas Arnold [a,b,*]

[a] Institute for Advanced Simulation, Forschungszentrum Jülich, Wilhelm-Johnen-Straße, 52428 Jülich, Germany
[b] Chair of Computational Civil Engineering, University of Wuppertal, Pauluskirchstraße 7, 42285 Wuppertal, Germany

ARTICLE INFO

Dataset link: https://doi.org/10.5281/zenodo.7618897

Keywords:
Global sensitivity analysis
Sobol indices
Horizontal flame spread simulation
Cone calorimeter simulation
Polymethyl methacrylate (PMMA)
Material property estimation
Inverse modelling
Optimisation
Fire Dynamics Simulator (FDS)

ABSTRACT

Predictive flame spread models based on temperature dependent pyrolysis rates require numerous material properties as input parameters. These parameters are often derived by optimisation and inverse modelling using data from bench scale experiments such as the Cone Calorimeter. The estimated parameters are then transferred to flame spread simulations, where self-sustained propagation is expected. A fundamental requirement for this transfer is that the simulation model used in the optimisation is sufficiently sensitive to the input parameters that are important to flame spread. Otherwise, the estimated parameters will have an increased associated uncertainty that will be transferred to the flame spread simulation. This is investigated here using a variance-based global sensitivity analysis method, the Sobol indices. The sensitivities of a Cone Calorimeter and a horizontal flame spread simulation to 15 effective properties of polymethyl methacrylate (PMMA) are compared. Results show significant differences between the setups: the Cone Calorimeter is dominated by strong interaction effects between two temperature dependent specific heat values, whereas the flame spread is influenced by several parameters. Furthermore, the importance of some parameters for the Cone Calorimeter is found to be time-varying, suggesting that single-value cost functions may not be sufficient to account for all sensitive parameters during optimisation.

1. Introduction

The development of fire simulation models that are capable of predicting the behaviour of flame spread from material properties is essential to overcome limitations of prescriptive design fires. By accounting for the heat transfer inside the solid and the coupling between the pyrolysis rates and the material temperatures, the heat release rate (HRR) can be predicted, rather than prescribed. This way, predictive models allow for the HRR to respond to changes in the surrounding conditions, such as reduced oxygen concentrations or activation of fire suppression systems. However, one major constraint in the development of predictive models is the difficulty in experimentally measuring all the required material properties, which are taken as model input parameters. Especially in the high temperature ranges where pyrolysis takes place, behaviours such as melting, bending, and bubbling make the determination of physical parameters extremely challenging under these conditions. During the burning of the sample, intermediate materials may be formed, which are difficult to isolate and analyse. Some of the phenomena, like bending or bubble formation can also not be directly reproduced in commonly used simulation software, such as the Fire Dynamics Simulator (FDS) [1].

As an alternative strategy, optimisation algorithms have been applied to parameter estimation in a so-called inverse modelling process (IMP) [2–12]. In the IMP, data from bench-scale experiments are used as a target for determining the set of input parameters that lead to the closest fit between the simulation output and the experimental data. Typically, the HRR measured from Cone Calorimeter experiments is taken as target, and the fitness is evaluated by a cost function. The result is a set of effective material properties, whose performance is subsequently validated in simulation setups of different scales, where a self-sustained flame spread is expected to occur.

This approach assumes that the estimated parameters are transferable across inherently different setups. Further, it is implied that the sensitivities of the models to the input parameters are equivalent, both in the bench-scale and in configurations involving flame spread. However, this assumption might not hold, particularly when the estimated parameter set is to be validated in flame spread simulations of increased scales. One possible reason is that the bench-scale experiments are performed under conditions (e.g. small sample size, uniform heating) where the flame spread is either negligible or simply does not occur.

* Corresponding author at: Institute for Advanced Simulation, Forschungszentrum Jülich, Wilhelm-Johnen-Straße, 52428 Jülich, Germany.
E-mail addresses: t.quaresma@fz-juelich.de (T.L.S. Quaresma), hehnen@uni-wuppertal.de (T. Hehnen), l.arnold@fz-juelich.de, arnold@uni-wuppertal.de (L. Arnold).

https://doi.org/10.1016/j.firesaf.2024.104116
Received 4 October 2023; Received in revised form 5 February 2024; Accepted 8 February 2024
Available online 9 February 2024

In addition, the bench scale experiments are difficult to be modelled sufficiently well because of unknown boundary conditions, limitations of existing sub-models, and difficulties in achieving high grid resolutions due to the prohibitive computational cost. As a consequence, the estimated parameter set becomes model-dependent, posing a challenge to its transferability. Thus, it is not possible to affirm that an estimated parameter set, that performs well in the Cone Calorimeter, will achieve an equivalent satisfactory result when transferred to the flame spread simulation [2,8]. It should be noted that while the focus of the current study is limited to investigating the uncertainties introduced in flame spread simulations by the estimated parameters, other sources of errors in the model can also contribute to inaccurate results. These may include errors related to the inadequate representation of physical mechanisms or numerical errors.

In this context, we focus on answering the following question: how sensitive is the Cone Calorimeter simulation setup to the parameters that are important to the flame spread simulation? This is crucial because parameters with low importance to the Cone Calorimeter simulation will be estimated with an enlarged degree of uncertainty, which is then carried over to the flame spread simulation. This happens because their effect on the cost function used in the optimisation is expected to be comparably low, leading the optimiser to freely choose any value from the pre-defined sampling ranges. If such parameters, on the other hand, are important to the flame spread simulation, their uncertainty is propagated, potentially compromising the reliability of the simulated results when compared to experimental data.

In this regard, we carry out sensitivity analyses (SAs) on two simplified simulation setups: a Cone Calorimeter; and a horizontal flame spread, to evaluate the models' sensitivities to a set of 15 input parameters. The parameter set consists of effective material properties of polymethyl methacrylate (PMMA), of its pyrolysis residue, and of an insulation material. The parameter set investigated here is taken from a previous study by Hehnen and Arnold [2]. In said study, the parameter set is determined, using two independent inverse modelling steps. In the first step, the pyrolysis reaction scheme is designed, for which the kinetic parameters and heats of gasification are estimated. After its successful completion, thermophysical parameters are determined in the second step, using a simplified Cone Calorimeter simulation setup. Since the pyrolysis reaction scheme is considered fixed at this point, we focus here only on this second step. To ensure compatibility to the previous study, the same Cone Calorimeter simulation setup is used here. The flame spread setup represents as well a simplified horizontal configuration, in which a self-sustained spread develops over a PMMA sample in still air. To the author's best knowledge, only a limited number of studies have investigated the flame spread behaviour over PMMA samples in such configuration, either experimentally [13,14] or numerically [15]. For this reason, this configuration is specifically chosen. All the simulations discussed in the present work were conducted with a self-compiled version of FDS (version FDS6.7.6-810-ge59f90f-HEAD) [1]. It is the same FDS version as was used in the previous study [2], to ensure compatibility.

Past studies [6,7,12,16,17] have addressed the importance of running SAs in order to improve the efficiency of strategies for material property estimation, which are based on multi-objective optimisation. In general, the SA is aimed for model simplification, by determining the relative importance of input parameters, such that the unimportant ones can be filtered out from the optimisation, saving computing time. More recently, Ding et al. [18] conducted one-at-a-time SAs on large-scale upward flame spread simulations to identify which input parameters affect the spread the most, so that their measurement/estimation can be improved. However, literature lacks comprehensive investigations into the importance of parameters across different simulation setups. Specifically, there is a shortage of research addressing whether the influence of a fixed set of estimated parameters changes when transitioning from the bench-scale simulation, with which they were derived, to scenarios involving flame spread. Our study seeks to fill this gap by not only identifying the most influential parameters, but also examining how these parameters' importance may vary from the Cone Calorimeter to the horizontal flame spread setup.

The SAs performed in this contribution are mainly discussed in terms of the Sobol indices [19,20], a robust global SA method based on the decomposition of variances. By varying input parameters simultaneously and not one at a time, the method is capable of quantitatively capturing interaction effects between input parameters on the model output of interest. It is therefore suitable for determining the sensitivities of non-linear and high-dimensional models, such as the complex pyrolysis models in question. Sensitivities are determined based on the degree of contribution that a certain input parameter has to the uncertainty (variance) of the model output. It has been applied in the field of fire safety science to investigate the influence of inputs on environmental fire spread models [21], and on the mass loss rate (MLR) calculated by the Arrhenius equation [17]. In this contribution, the Sobol indices are estimated to express the effects of the input parameters on different simulation outputs:

- the temporal development of the HRR, which is convenient to assess how the influence of a certain input parameter varies throughout the course of the simulation;
- the root mean square error (RMSE), calculated between the simulated HRRs and the measured HRR, commonly used as a cost function in the IMP;
- and the rate of spread (ROS), calculated in an additional post-processing step, based on the derivative of the flame front position with respect to time.

In addition, the relation between the RMSE and the ROS and their two respective most influential input parameters are qualitatively discussed on the basis of scatterplots, providing a visual representation of their responses to changes in the inputs.

This article is accompanied by a publicly available data repository on Zenodo [22], containing the simulation data, the Python scripts used for data analysis, and supplementary material data.

2. Methods

2.1. Cone Calorimeter simulation

The simulation setup of a simplified Cone Calorimeter considered as reference case in this work stems from freely available previous studies [2]. The model was initially developed to be used in an IMP for estimating thermophysical properties of PMMA, based on data of Cone Calorimeter experiments of black cast PMMA. The simulations were conducted with the same FDS version used in the scope of this work for consistency. The experimental data was provided by the Aalto University to the publicly available MaCFP database [23]. In the Cone Calorimeter experiments, a square sample of PMMA with 10 cm edge length and 6 mm thickness is exposed to a radiative heat flux of 65 kW/m^2. The insulation material of equal surface area is positioned below the PMMA sample for insulation, and is 2 cm thick.

Amongst the different Cone Calorimeter simulation models presented previously [2], we chose the one labelled as "Cone_04". Cone_04 uses an increased resolution when compared to similar approaches in the field, in which Cone Calorimeter models were also employed to estimate material properties. Viitanen et al. [3], Hehnen et al. [5], McCoy et al. [24] and Beji and Merci [8] considered fluid cell sizes of 5 cm or more, whereas in Cone_04, 3.33 cm fluid cells are defined, see Fig. 1(a). The coarse fluid cell resolution of 3.33 cm is necessary, due to the large amount of simulations conducted during the IMP [2].

In Cone_04, the radiative heat flux of 65 kW/m^2 is assigned to the sample surface, thus avoiding the need for modelling the heater. Despite this simplification, the non-uniform heating of the sample surface induced by the conical heating element is accounted for. This

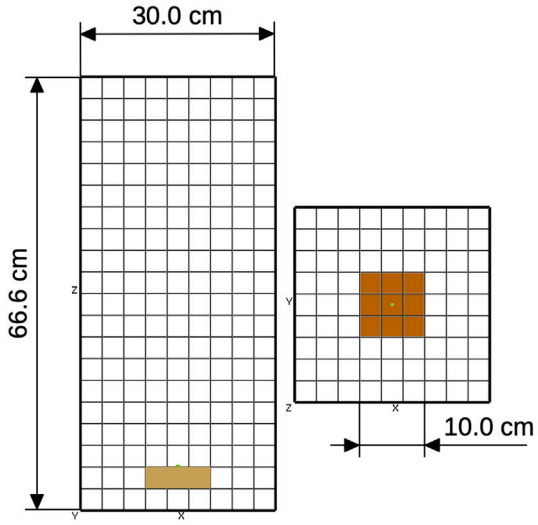

(a) Dimensions and fluid phase resolution.

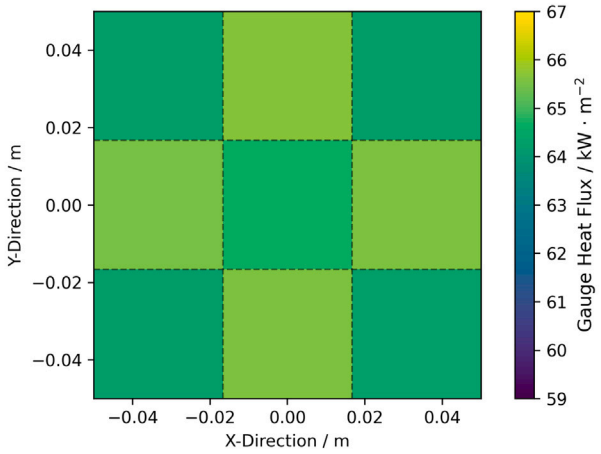

(b) Gauge heat flux mapping applied to the sample surface.

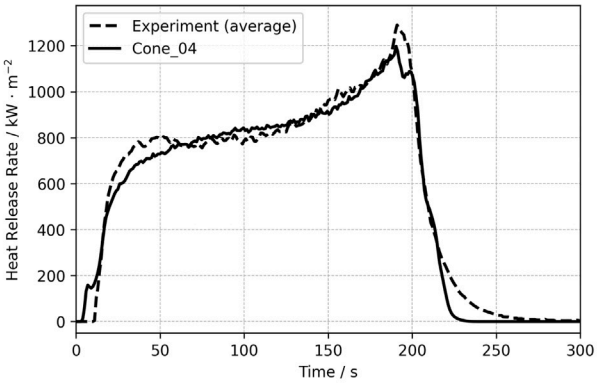

(c) Best fit of simulated HRR to experimental data.

Fig. 1. Simulation setup of the simplified Cone Calorimeter, "Cone_04", used for estimating thermophysical parameters of black cast PMMA [2], taken as reference case.

is accomplished, by defining multiple surfaces with slightly different values of heat fluxes, as presented in Fig. 1(b). In addition, the model stood out for providing a set of thermophysical properties, which lead to the best fit to the experimental HRR data in the IMP, see Fig. 1(c). For those reasons, Cone_04 is expected to lead to an improved representation of the sample heating and thermal decomposition, and it is therefore taken here as reference case.

In the following subsections, we provide only a concise description of the model that is relevant to the scope of this contribution. For further details and additional resources, the reader should refer to the original work [2] and to the FDS User's Guide [1], both freely available.

2.1.1. Pyrolysis and solid phase

In this approach, it is assumed that PMMA pyrolysis can be described by an elementary first-order reaction of the form:

$$\text{PMMA (solid)} \longrightarrow 0.99\,\text{FuelMixture (gas)} + 0.01\,\text{Residue (solid)}$$

where the fuel mixture is composed by methane, ethylene and carbon dioxide, and the residue is an inert solid product. The rates of pyrolysis are dependent on the local temperatures of the solid, and are calculated with the Arrhenius equation, as implemented in FDS [1]. As a strategy to achieve a better fit to the experimental data, it was assumed that different fractions of the total PMMA mass decompose, each, at a different rate, as described in [2]. The kinetic parameters (pre-exponential factor, activation energy), heats of gasification, and corresponding PMMA mass fractions were determined in the first IMP step, taking data of Microscale Combustion Calorimetry (MCC) and Thermogravimetric Analysis (TGA) as target [2]. In this same step, the volume fractions of the fuel mixture components were also determined. Once the first step is concluded successfully, the second IMP step follows.

In this contribution, the focus lies on the second IMP step, in which the kinetic parameters, heats of gasification, and mass fractions will not be adjusted anymore. In this second step, thermophysical properties of PMMA, of the insulation material, and of the pyrolysis residue are determined, using the Cone Calorimeter. For PMMA, these properties are: emissivity, absorption coefficient, refractive index, specific heat and thermal conductivity. The thermal conductivity and the specific heat of PMMA were both defined as temperature-dependent values, following a piecewise linear function with reference points at 150 °C, 480 °C and 800 °C. This implies that values falling between the specified temperature points are determined through linear interpolation. In FDS, material property values below or above the given temperature range are assumed constant, equal to the first or last specified value for the property. The temperature-dependent variation of the estimated values of the thermal conductivity and the specific heat of PMMA is depicted in Fig. 2. The experimental MCC data is presented in the plots in a secondary y-axis to underscore the temperature range associated with PMMA decomposition. Notably, in Fig. 2, it should be observed that, although 800 °C surpasses the upper temperature limit indicated

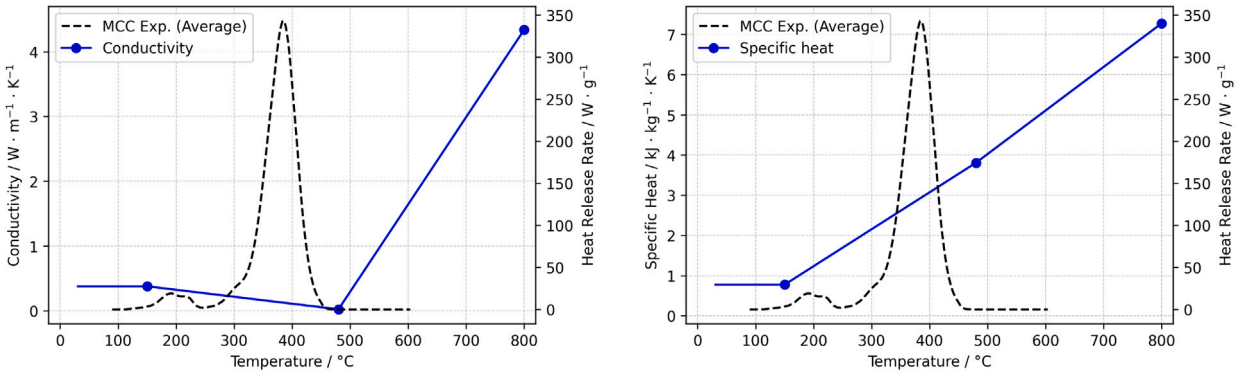

(a) Estimated conductivity and MCC data.

(b) Estimated specific heat and MCC data.

Fig. 2. Temperature dependence of the piecewise functions defined for the estimated conductivity and specific heat of PMMA. The experimental MCC data is included to highlight the temperature range at which the material pyrolyses.

Table 1
Set of effective properties describing the solid phase materials. The parameters were estimated in the work of Hehnen and Arnold [2] using the reference case Cone_04.

	Material	Parameter	Estimated IMP value	Unit
1	PMMA	Emissivity	0.941	–
2		Absorption coefficient	7978.8	m^{-1}
3		Refractive index	2.854	–
4		Conductivity at 150 °C	0.379	W m^{-1} K^{-1}
5		Conductivity at 480 °C	0.024	W m^{-1} K^{-1}
6		Conductivity at 800 °C	4.337	W m^{-1} K^{-1}
7		Specific heat at 150 °C	0.774	kJ kg^{-1} K^{-1}
8		Specific heat at 480 °C	3.808	kJ kg^{-1} K^{-1}
9		Specific heat at 800 °C	7.275	kJ kg^{-1} K^{-1}
10	Residue	Emissivity	0.552	–
11		Conductivity	4.509	W m^{-1} K^{-1}
12		Specific heat	5.893	kJ kg^{-1} K^{-1}
13	Backing	Emissivity	0.441	–
14		Conductivity	2.408	W m^{-1} K^{-1}
15		Specific heat	4.067	kJ kg^{-1} K^{-1}

by the MCC data (600 °C), the 800 °C point contributes to estimating values within the 480 to 600 °C range. However, given that the MCC data indicates nearly complete decomposition below 500 °C, it is expected that the 800 °C point has minimal impact on the simulations, if any. Conversely, as the MCC data shows that most of the PMMA decomposes between 150 and 500 °C, values of material properties at these temperatures are anticipated to have a more significant influence.

For the insulation material and the residue, the parameters are: emissivity, thermal conductivity and specific heat. In total, the material properties estimated in an IMP using the Cone Calorimeter count 15 input parameters. Their respective estimated values are presented in Table 1.

The value of PMMA density was directly calculated from reported mass and dimensions of the sample. Density of the residue was fixed to an arbitrary value due to lack of information, and density of the insulation material was taken from the MaCFP database [23].

The default one-dimensional heat conduction model in FDS was used [1]. In this model, heat conduction is calculated only in the direction normal to the sample surface. The solid phase solution is updated at every time step and the node spacing of the PMMA layer is set to uniform. For the layer of the insulation material, the default stretched node spacing is considered.

The FDS default grid resolution in both layers is increased by a factor of 10, by setting the `CELL_SIZE_FACTOR` (CSF) to 0.10. This modification leads the PMMA layer to be discretised in 96 equally spaced cells, and the layer of the insulation material in 11 stretched cells. Nonetheless, the cell size of the PMMA layer is automatically re-defined during the simulation, as the temperature-dependent parameters affect the thermal diffusivity, as well as changing layer thickness due to sample consumption.

2.1.2. Combustion and gas phase

Combustion of the fuel mixture is assumed mixing-controlled and the soot yield is set to 0.022 g/g [25]. The Large Eddy Simulation (LES) is chosen as simulation mode. All default settings related to the models that accompany LES in FDS are kept unchanged. This means that sub-grid scales are modelled using the Deardorff model for the eddy viscosity, and the wall-adapting local eddy-viscosity (WALE) is used as near-wall turbulence model. The LES default radiation model considers the grey gas assumption and the flame to be optically thick. A specified radiative fraction controls the portion of the total heat that is released as thermal radiation, and it is set to 0.35 for the defined fuel mixture. The used radiative fraction is the default value in FDS for unspecified species. The initial ambient temperature is set to 30.85 °C, consistent with what was reported from the experiment.

An overview of the geometry, domain and mesh resolution can be seen in Fig. 1(a). The computational domain extends in the x- and y-directions from −15.0 cm to 15.0 cm, and in the z-direction from −6.6 cm to 60.0 cm. A uniform grid is defined by assigning 9 x 9 x 20 cells respectively in the x-, y-, z-directions, resulting in cells of 3.33 cm edge length. A single mesh is used for the whole domain, whose boundaries are set to open conditions. The centre of the sample is positioned at the origin (0,0,0). The sample holder is considered an inert obstruction of 10 cm edge length in the x- and y-directions and 3.33 cm in the z-direction.

In order to evaluate grid independence, three additional simulation cases are built from the reference case Cone_04, each considering a different fluid cell size. The fluid cell sizes were determined based on the feasibility of implementing the resolutions in a model suitable for optimisation, in terms of computing times [2]. The goal of this analysis is to verify whether the sensitivities to the input parameters are maintained across different resolutions that could be used in the context of an IMP. Following the terminology proposed earlier [2], Cone_04 is here labelled as “C3”, as a reference to the 3 × 3 number of divisions of the sample surface, see Figs. 1(a) and 1(b). Thus, the first case is called “C2” (2 × 2 divisions), referring to 5 cm cells, the second case is called “C5” (5 × 5 divisions), with 2 cm cells in the fluid phase, and the third is “C7” (7 × 7 divisions) with 1.43 cm cells. The heat flux mapping applied to the surface of the sample is adjusted in each case to conform with the C2, C5 and C7 resolutions. It is important to note, that the C5 and C7 cases represent a level of refinement that would introduce great computational effort in an IMP. For this reason, resolutions above C7 are not explored here.

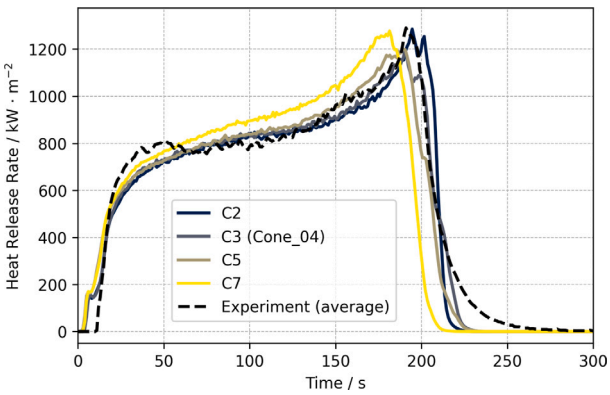

Fig. 3. Effect of different fluid phase resolutions on the simulated HRR of the reference case Cone_04.

From Fig. 3, it is possible to observe that the simulated HRR does not significantly change with grid refinement in the fluid phase. The overall shape of the curve is maintained, and the main impact is on the region of the HRR peak, which is only slightly decreased and shifted to the left, as the fluid phase resolution increases. It is therefore expected, that the relative importance of the input parameters to the simulated HRR is conserved across the tested fluid cell sizes.

2.2. Horizontal flame spread simulation

The flame spread simulation considered in the SA represents a simplified horizontal configuration, in which a self-sustained spread develops over a slab of PMMA. No influence of ventilation or wind conditions is considered, such that the flame is not disturbed by any external changes in the flow field.

The solid phase modelling used in the flame spread simulation is transferred from the reference case of the Cone Calorimeter simulation. This means that the PMMA sample, pyrolysis, the insulation material, and the solid phase resolution are the same as described in Section 2.1.1. This is important since the main goal of this contribution is to compare the responses of both simulations to variations in the same set of material properties presented in Table 1.

In the flame spread setup, the PMMA sample is positioned on top of an inert ochre obstruction which serves as a sample holder, see Figs. 4(a) and 4(b). The sample dimensions are 23 cm x 9.5 cm x 0.6 cm, and an external heat flux of 65 kW/m^2 is prescribed for 100 s to an area of 2.5 cm x 9.5 cm to start ignition. The ignition area is located at the left end of the sample, and it is represented by the dark brown patch in Figs. 4(a) and 4(b).

The dimensions of the computational domain are 26 cm x 12.5 cm x 11 cm, and it is divided in 26 meshes of 2.0 cm x 12.5 cm x 5.5 cm each, to allow parallel computation. The fluid cell size used in the flame spread simulation is set to 0.5 cm. The fluid phase modelling differs from the Cone Calorimeter setup only by the cell size. The simulation mode, combustion and radiation modelling are given also according to what was used in the Cone Calorimeter, as described in Section 2.1.2. An overview of the horizontal flame spread setup is presented in Fig. 4.

Given its dimensions, the horizontal flame spread simulation can be seen as a simplified small-scale setup of the same scale as the Cone Calorimeter simulation. In addition, the heat flux applied to the dark brown patch is intentionally set to the value used in the Cone Calorimeter, 65 kW/m^2. Such similarities allow for the flame in the flame spread simulation to transition from the ignition stage, where the sample heating is modelled according to the Cone Calorimeter, to a self-sustained stage of spread.

In order to build a reference case also for the flame spread simulation, the set of material properties used in Cone_04 was transferred to the horizontal flame spread setup. The fluid cell size was set to 0.5 cm because it was the maximum size to allow a self-sustained spread to occur over the sample. Several other attempts with reasonably larger fluid cell sizes (3 cm, 2 cm, 1 cm) were carried out for even larger samples, but the flame would not spread much farther than the ignition zone, extinguishing shortly after the end of the applied external heat flux. Testing an increased grid resolution with 0.25 cm cells revealed higher HRR and rate of spread compared to the 0.5 cm case, indicating that no grid independence was achieved. The grid dependence in flame spread simulations based on material pyrolysis appears to be a current limitation of FDS. This limitation has been attributed to the implemented empirical boundary layer correlations. However, it is important to highlight that in FDS there are a number of other (sub)model parameters that influence the spread, for example, parameters related to the radiation model. Determining the sensitivity of these other parameters, however, falls outside the scope of the present work, which is focused on the estimated material properties.

Fig. 5 shows slices containing data of the HRR per unit volume (HRRPUV) taken at the central plane of the domain ($y = 0$) along the x-axis, and at different points in time. The patches coloured in magenta in the slices show the location of the flame front, which is here defined as the cell containing the maximum value of HRRPUV in the one-dimensional row of fluid cells touching the surface of the sample. The HRRPUV slices shown in Fig. 5 were generated using the fdsreader version 1.9.9, an open source Python module developed to read FDS output data [26].

2.2.1. Rate of spread (ROS)

In this section, a methodology is introduced to determine the ROS, since it is not a direct output of the simulation software, and requires an additional post-processing step. The ROS is an important quantity to represent the flame spread phenomenon in the context of fire safety, because it describes how fast a fire can develop in a compartment, impacting the degree of damage and the time to reach flashover. For this reason, the ROS is taken as output of interest in the SA, such that the influence of the material properties can not only be evaluated on the temporal development of the simulated HRR, but also on a single value representing the whole phenomenon. Furthermore, the ROS could be used for validation or as target in future IMPs that aim at estimating input parameters from flame spread experiments in bench-scale dimensions.

The ROS is determined by the rate at which the position of the flame front changes with respect to time. At every time step, the maximum HRRPUV is tracked in the fluid cells ahead of the last recorded flame front position, in order to distinguish between leading and trailing edges of the flame. This is done in order to prevent a maximum HRRPUV located in the back of the flame from being mistakenly accounted as the front. In the post-processing, the values of HRRPUV can be read either from slice files, as shown in Fig. 5, or from multiple devices that are positioned along the centre line of the sample, as indicated by the green dots in Figs. 4 and 6(a).

The recorded positions can then be plotted against time, as demonstrated in Fig. 6(b) for the reference case. The small plateaus in the black curve correspond to the periods when the position of the flame front does not change. The vertical blue dashed line at 100 seconds indicates the time instant at which the applied external heat flux over the ignition area ceases. Three zones of spread are identified, as shown in Figs. 6(a) and 6(b): (1) ignition zone; (2) self-sustained spread; and (3) extinction zone, influenced by the end of the sample. A linear relation between position and time is fitted to the region where a steady self-sustained spread develops without the influence of the ignition and extinction zones. Finally, the ROS is taken as the slope of the red linear curve in Fig. 6(b). In this case, the ROS calculated with the presented method resulted in 0.13 mm/s, which is a value of the same order of magnitude (0.10 mm/s) as those obtained from horizontal flame spread experiments over cast PMMA samples of similar thickness and under the same ventilation conditions, as reported elsewhere [13].

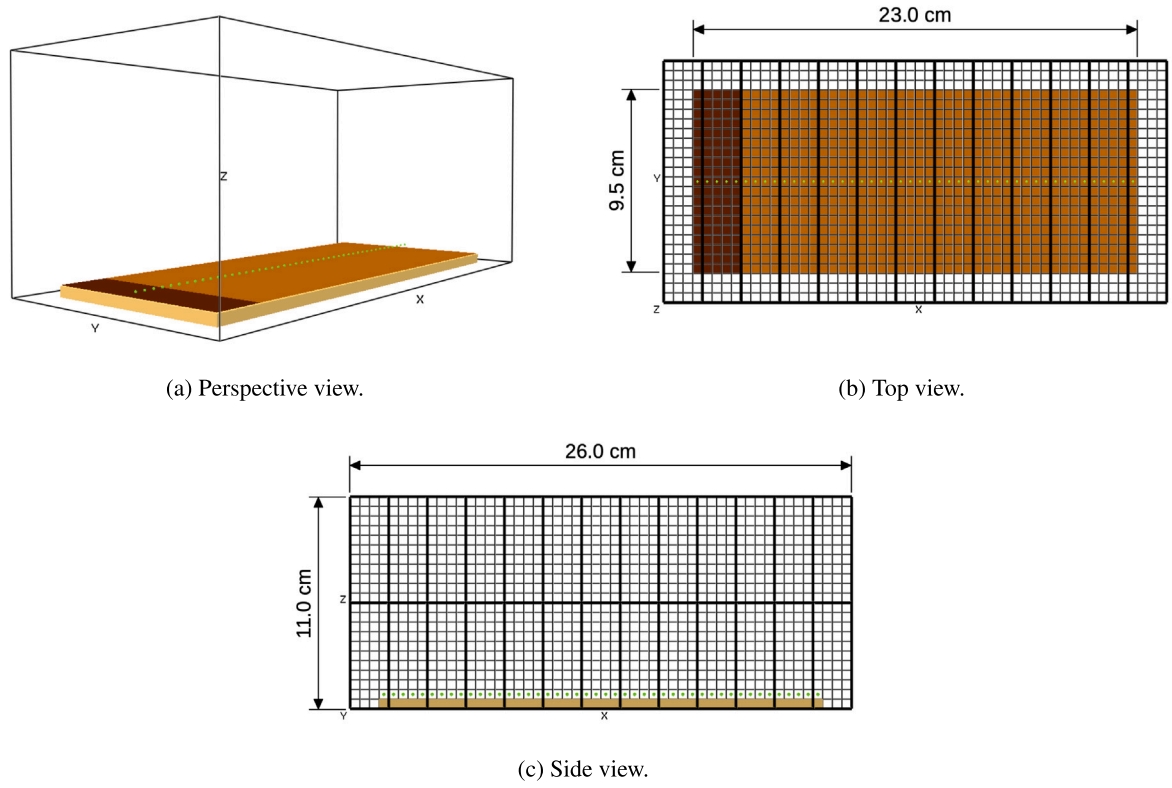

(a) Perspective view.

(b) Top view.

(c) Side view.

Fig. 4. Overview of the horizontal flame spread simulation setup used in the SA, showing the fluid phase resolution, meshes and dimensions. Bold lines indicate mesh borders.

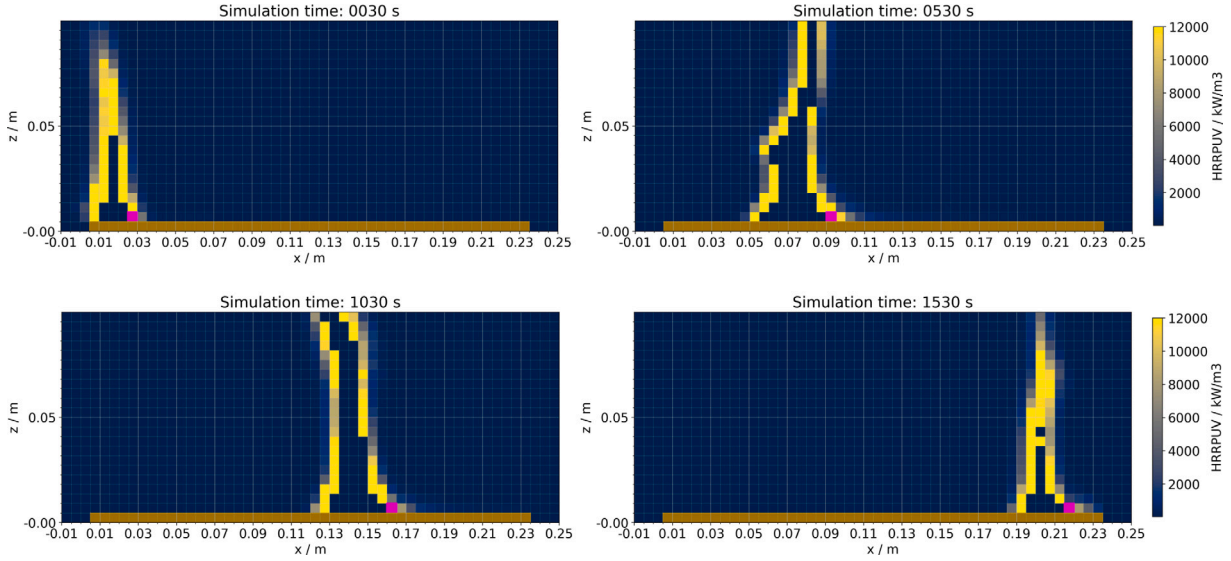

Fig. 5. Slices at $y = 0$ showing the HRRPUV of the reference case of flame spread simulation, at different points in time. The set of material properties used in the solid phase modelling was transferred from Cone_04. The fluid cells coloured in magenta indicate the position of the flame front.

2.3. Sensitivity analysis (SA)

The SAs conducted in this study follow the methodology of the Sobol sensitivity indices, a global SA method based on the decomposition of variances [19]. Unlike local SA methods, which vary input variables one at a time, the Sobol indices can efficiently explore multi-dimensional parameter spaces, accounting for the effects of all possible combinations of input parameters. Such feature makes it well-suited for quantifying sensitivities of non-linear models, which cannot be adequately assessed by one-at-a-time methods [27].

2.3.1. Mathematical formulation

In the mathematical formulation of the method proposed by Sobol [19], the model can be described as a multivariable function $Y = f(\mathbf{X})$, where Y is a scalar output and $\mathbf{X} = (X_1, X_2, \ldots, X_k)$ is a point in a k-dimensional parameter space. If f satisfies the according requirements, it can be decomposed into terms of increasing dimensions:

$$f = f_0 + \sum_i f_i(X_i) + \sum_i \sum_{j>i} f_{ij}(X_i, X_j) + \cdots + f_{12\ldots k}. \quad (1)$$

Each individual term is a function only of the inputs in its index, that is, f_0 corresponds to the constant part of the function, $f_i = f_i(X_i)$

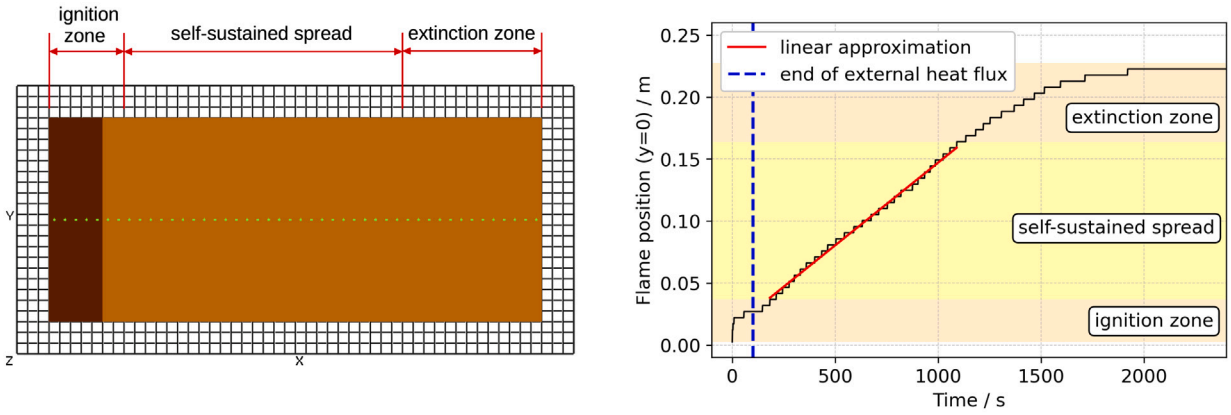

(a) Regimes of spread.

(b) Flame front position *vs* time.

Fig. 6. Determining the rate of spread (ROS) for the reference case of the horizontal flame spread simulation.

depends only on one component (here X_i) of the parameter vector $\mathbf{X}$, $f_{ij} = f_{ij}(X_i, X_j)$ depends on two components and so on. The total number of terms is equal to 2^k, out of each k terms are called first-order function f_i, f_{ij} are second-order functions, and so forth. Eq. (1) can be squared and integrated to generate the decomposition of variances:

$$V(Y) = \sum_i V_i(X_i) + \sum_i \sum_{j>i} V_{ij}(X_i, X_j) + \cdots + V_{12\ldots k} \tag{2}$$

where the total variance of the output $V(Y)$ is split down into $2^k - 1$ different partial variances, each accounting for fractions of the output variance that is induced by the corresponding input, or combinations of inputs. For example, $V_i(X_i)$ is the induced variance on the output when X_i is varied alone, while $V_{ij}(X_i, X_j)$ is the induced variance when X_i and X_j are varied together. Two parameters are said to interact when their combined effect on the output is different from the sum of their single effects.

The Sobol indices are obtained by dividing each partial variance in Eq. (2) by the total variance of the output $V(Y)$, giving:

$$\sum_i S_i + \sum_i \sum_{j>i} S_{ij} + \cdots + S_{123\ldots k} = 1 \tag{3}$$

where the indices S_i, S_{ij}, etc, are ratios varying from 0 to 1. Following a similar terminology used for the decomposition in Eq. (1), different types of sensitivity indices are defined:

- S_i: first-order indices, provide a measure of main effects, i.e. the fractional contribution of X_i to the total variance of Y;
- $S_{ij..k}$: higher-order indices, measure interaction effects between the inputs indicated in their subscripts;
- ST_i: total-order indices, account for all the effects due to variations in X_i, i.e. first-order effects and interaction effects.

The indices provide a measure of importance of each input parameter by quantifying how much the variance of the output could be reduced if a given input parameter, or combination of parameters, could be fixed. For example, $S_i = 0.10$ means that 10% of the variance of the output could be reduced if X_i is fixed to a known value. Similarly, $ST_i = 0$ implies that X_i is non-influential and can be fixed anywhere in its distribution without affecting the variance of the output [28]. When interaction effects exist, the sum of all total-order effect indices is greater than unity.

As the number of sensitivity indices to be computed in Eq. (3) increases exponentially with the number of input parameters, the calculation of high-order indices can become expensive. It is therefore convenient and often sufficient to express sensitivities in terms of first-order and total-order sensitivity indices. The difference between ST_i and S_i provide a measure of how much X_i is involved in interactions with any other input parameter.

For a better understanding, an example is given for a function of the form $Y = f(\mathbf{X})$ with $\mathbf{X} = (A,B,C)$, for which Eq. (3) becomes:

$$S_A + S_B + S_C + S_{AB} + S_{AC} + S_{BC} + S_{ABC} = 1 \tag{4}$$

where S_A, S_B, and S_C are the first-order indices, accounting for the main effects of A, B, and C respectively. The second-order indices S_{AB}, S_{AC} and S_{BC} account for the interaction effects between the pairs of inputs in their subscripts. Accordingly, S_{ABC} is the third-order index which accounts for the interaction effects on the output when A, B and C are varied together. In this example, the total-order indices of inputs A, B and C are given as:

$$ST_A = S_A + S_{AB} + S_{AC} + S_{ABC} \tag{5}$$

$$ST_B = S_B + S_{AB} + S_{BC} + S_{ABC} \tag{6}$$

$$ST_C = S_C + S_{AC} + S_{BC} + S_{ABC} \tag{7}$$

in which the total effect of ST_i is the sum of all the terms in Eq. (3) where the parameter X_i is considered.

In this study, the sensitivities of the simulation setups to the set of 15 input parameters will be evaluated only in terms of first-order (S1) and total-order indices (ST).

2.3.2. Estimating the Sobol indices

For analytical models, the integrals related to the calculation of variances can be solved also analytically. However, this is not the case for the simulation models investigated in this work. The approach used here for estimating the indices assumes that the simulation model is a "black box", and no information on model behaviour is known other than what is perceived through variations in the model's inputs and outputs. The partial variances are calculated by quasi-Monte Carlo estimates and therefore a large number of simulations needs to be conducted, one for each sample of model inputs [20].

Input parameters are assumed to be independent and uniformly distributed within their sampling limits, and samples are generated by employing the Saltelli's sampling scheme [20], which is based on the Sobol sequence. The Sobol sequence is a type of low-discrepancy quasi-random sequence that creates an efficient space filling sampling of the high dimensional parameter space. The SAs, including sampling, estimation of the indices and confidence intervals, are carried out within the SALib Python library (version 1.4.5) [29,30]. The confidence intervals for the sensitivity indices, are calculated using the bootstrap

Table 2
Sampling limits and units of the input parameters considered in the SAs.

	Material	Parameter	Sampling limits	Unit
1	PMMA	Emissivity	[0.799 ; 0.999]	–
2		Absorption coefficient	[6781.9 ; 9175.5]	m^{-1}
3		Refractive index	[2.426 ; 3.281]	–
4		Conductivity at 150 °C	[0.322 ; 0.436]	W m^{-1} K^{-1}
5		Conductivity at 480 °C	[0.021 ; 0.028]	W m^{-1} K^{-1}
6		Conductivity at 800 °C	[3.687 ; 4.988]	W m^{-1} K^{-1}
7		Specific heat at 150 °C	[0.658 ; 0.890]	kJ kg^{-1} K^{-1}
8		Specific heat at 480 °C	[3.237 ; 4.380]	kJ kg^{-1} K^{-1}
9		Specific heat at 800 °C	[6.183 ; 8.366]	kJ kg^{-1} K^{-1}
10	Residue	Emissivity	[0.469 ; 0.635]	–
11		Conductivity	[3.833 ; 5.186]	W m^{-1} K^{-1}
12		Specific heat	[5.009 ; 6.777]	kJ kg^{-1} K^{-1}
13	Backing	Emissivity	[0.375 ; 0.507]	–
14		Conductivity	[2.047 ; 2.769]	W m^{-1} K^{-1}
15		Specific heat	[3.457 ; 4.677]	kJ kg^{-1} K^{-1}

resampling technique. This method involves generating multiple resampled datasets by randomly selecting and replacing data points from the original dataset. Sensitivity indices are then computed for each resampled dataset, creating a distribution of indices. The standard deviation of these resampled sensitivity indices is calculated, and confidence intervals are established based on a confidence level of 95%. These intervals provide an estimate of the uncertainty associated with the sensitivity indices, indicating the range within which the true values are likely to lie.

The sampling limits of the 15 input parameters were defined by taking 15% of variation around the best parameter set determined in the IMP. A restriction is imposed only for the upper limit of emissivity, such that it would not exceed the value of 0.999. The sampling limits and units of each input parameter are shown in Table 2. In order to achieve better uniformity, the sampling scheme requires the number of samples N to be generated as powers of 2, i.e. $N = 2^q$, particularly when sampling high-dimensional parameter spaces. For the Cone Calorimeter, $q = 17$, and for the flame spread simulation, $q = 15$. Fewer samples are considered for the flame spread simulation due to the significantly enlarged computing time in comparison to the Cone Calorimeter simulation.

The influence of the 15 input parameters on two types of outputs are evaluated, namely "multiple-value" or "single-value" outputs. The simulated HRR is a time-series, therefore it is a multiple-value output. In this case, the sensitivity indices are calculated at every point in time, and, consequently, are presented also as time-series. This is convenient to evaluate how the influence of a given input parameter varies over the course of the simulation.

With respect to the single-value outputs, two indirect quantities are calculated from the simulated HRRs: the root mean square error (RMSE) for the Cone Calorimeter; and the ROS for the flame spread simulation, see Section 2.2.1. It is important to evaluate the effect of the input parameters on the RMSE because it is typically taken as the cost function in the optimisation [31]. Here, the RMSE is calculated as:

$$\mathrm{RMSE} = \sqrt{\frac{1}{\mathrm{N}} \sum_{i=1}^{\mathrm{N}} \left(\mathrm{simulated\_HRR}_i - \mathrm{measured\_HRR}_i\right)^2} \tag{8}$$

where N is equal to the number of points of the HRR curve, and the measured_HRR is the experimental HRR curve shown in Fig. 1(c), which was used as target in the IMP in earlier work [2]. Table 3 shows for each setup the number of simulations, the output of interest taken in the SAs, their type, and how the Sobol indices are presented.

The 131,072 Cone Calorimeter simulations consumed about $1.6 \cdot 10^4$ core-hours, whereas the 32,768 flame spread simulations demanded substantial $1.2 \cdot 10^7$ core-hours in total. These numbers were estimated considering that each individual simulation took as much time as the elapsed wall clock time of their respective reference simulation case:

Table 3
Summary of the different simulation outputs considered in the SAs based on the Sobol indices.

Simulation setup	Number of simulations	Output of interest	Type of output	Sobol indices
Cone Calorimeter	131,072	HRR	multiple-value	time-series
		RMSE	single-value	bar plots
Flame spread	32,768	HRR	multiple-value	time-series
		ROS	single-value	bar plots

7.2 minutes for the Cone Calorimeter simulation; and 14.1 h for the flame spread simulation. All simulations were executed on a high performance computing cluster comprising 268 worker nodes with a total of 17,152 cores. Each worker node features 2 AMD EPYC 7452 32-Core processors, each operating at a base clock speed of 2.350 GHz, and is equipped with 256 GB of memory, corresponding to 4 GB per core.

The Sobol indices lead to a straightforward and concise way of ranking the input parameters according to their importance, by providing a quantitative measure of sensitivities. This is advantageous when dealing with multi-dimensional parameter spaces. Yet, the indices do not provide the type of relation (i.e. linear, non-linear) between model output and the individual inputs, which is also a meaningful aspect of the analysis. This gap is filled by a complementary qualitative analysis with scatterplots. The scatterplots aim at showing the relation between the output and the two most influential parameters that have been previously identified by the indices. Only single-value outputs (RMSE and ROS) are contemplated in the analysis with scatterplots.

3. Results and discussion

3.1. Effects on the HRRs

The effects of the 15 input parameters on the HRRs of the Cone Calorimeter and the flame spread simulations are discussed in terms of the time-series of the ST and S1 indices, expressing respectively the total-order and the main effects of each parameter. The ST and S1 indices calculated over the simulated HRR of the Cone Calorimeter simulation are shown in Figs. 7(a) and 8(a). Figs. 7(b) and 8(b) show the ST and S1 indices with respect to the HRR of the flame spread simulation up to the initial 300 seconds, and Figs. 7(c) and 8(c) show the indices up to 2000 seconds.

In Figs. 7 and 8, confidence intervals are expressed by the shaded areas around the curves, indicating the uncertainty in the estimation of the indices. The uncertainty in the confidence intervals arises from the inherent variability in the sampled data and the statistical methods employed, reflecting the precision of the Sobol index estimates. Wider intervals indicate greater uncertainty due to limited sample size, while narrower intervals, as observed in the Cone Calorimeter case, result from a larger number of samples, enhancing the reliability of the estimates through the quasi-Monte Carlo approach.

Regarding the Cone Calorimeter simulation setup, the SA was conducted for the C2, C3 and C5 cases of fluid cell sizes shown in Fig. 3. As expected, the obtained sensitivity indices were very similar, showing that the HRR response to the input parameters is maintained across the evaluated fluid grid resolutions. The C7 case was left out due to the enlarged computing time that would have been necessary to run the simulations. Nonetheless, due to the similarity in the HRR shapes presented Fig. 3, it is assumed that the sensitivities to input parameters would be also similar for the C7 case. In light of the similarity and for the sake of brevity, only the SA results with respect to the reference case (C3) of the Cone Calorimeter will be presented here. The remaining results can be found in the supplementary material data available online [22].

As can be seen in Fig. 7(a), within the initial 5 seconds of the Cone Calorimeter simulation, only three parameters have non-zero

ST indices: PMMA emissivity, conductivity at 150 °C and specific heat at 150 °C. This means that these are effectively the only parameters affecting the HRR up to this point in time, whereas all the other parameters remain unimportant. At about 10 seconds, the importance of the specific heat at 150 °C momentarily drops, while the influence of the same property at 480 °C increases. This is explained by the dependency of the specific heat on temperature, established by the piecewise linear function. As the sample heats up, higher temperatures are reached, causing the value of the property to change, and consequently its influence over the HRR. Soon after that, at about 25 seconds, the importance of the specific heat at 150 °C increases again, and it becomes, together with the specific heat at 480 °C, the two most important parameters to affect the HRR of the Cone Calorimeter. At the same time, the initial total-order effects of PMMA emissivity and conductivity decrease to a practically negligible value for the rest of the simulation time.

It is important to distinguish, however, two stages of influence of the values of specific heat on the HRR, which become evident when Figs. 7(a) and 8(a) are compared. The stages are defined by the difference between their respective ST (Fig. 7(a)) and S1 (Fig. 8(a)) indices, indicating the degree of interaction effects between the two parameters. In the first stage, ranging from about 25 to 120 seconds, the difference between ST and S1 is small, indicating therefore that the interaction effects are also small. This implies that the HRR is affected predominantly by main effects of the two parameters, expressing that the effects of changing them individually is dominant to the HRR at this stage. In the second stage, from 120 to 150 seconds, an approximately synchronised and significant increase of the ST indices is observed. At the same time, their corresponding S1 indices decline in a similar trend, as presented in Fig. 8(a). The large discrepancy between ST and S1, together with the fact that all the other parameters have ST indices close to zero, indicate that strong interaction effects between the values of specific heat at 150 °C and at 480 °C dominate the HRR from this point forward in the Cone Calorimeter simulation. Still, at around 125 seconds, a timid increase in the ST indices of the PMMA and the residue emissivities are almost solely related to interaction effects, given that their correspondent S1 indices are very close to zero.

By comparing Figs. 7(b) and 8(b) to Figs. 7(a) and 8(a), a very similar ranking of effects can be identified between the initial 100 seconds of the flame spread simulation, and the initial 125 seconds of the Cone Calorimeter simulation. In the flame spread setup, the initial 100 seconds corresponds to the heating of the dark brown part of the sample by the prescribed external heat flux to start ignition (see Section 2.2). Given that this approach is the same as the one used in the Cone Calorimeter for heating up the whole sample, the similarity in modelling for using the `EXTERNAL_FLUX` function in FDS is clearly captured by the sensitivity indices.

The vertical blue dashed line at 100 seconds in Figs. 7(b), 7(c), 8(b) and 8(c) marks the end of the external heat flux and thus the transition from the ignition phase to the phase where the spread is self-sustained. The transition from one phase to another is highlighted by an abrupt change in the importance of some parameters. The ignition patch receives a large heat flux, consisting of the prescribed external flux of 65 kW/m^2 and the heat feedback of the flame. This heat flux is larger than that during the steady-state throughout the self-sustained spread. Once the external flux is shut off, the initial flame shrinks rapidly. The flame takes some time to recover and grow to the size of the steady-state. This is reflected in the increased effect of the specific heat of PMMA at 150 °C between 100 and 270 s. However, after approximately 270 s of simulation, their roles reverse, and the specific heat at 480 °C becomes more influential than that at 150 °C for the remaining simulation time. As more of the material heats up and its temperature increases, the specific heat at 480 °C prevails, as presented in Fig. 7(c). In addition, it can be seen that despite oscillations, after the initial 300 seconds, the specific heat values along with emissivity, conductivity at 150 °C of PMMA, and the specific heat of the insulation material remain as the most influential parameters affecting the HRR. Another important observation concerns the least important parameters, whose effects on the HRR, although smaller, are not insignificant. The only exception is the refractive index, whose indices turned out to be zero in both simulation setups.

Figs. 7(b), 7(c), 8(b) and 8(c) also show that the specific heat of the insulation material transitions from having no relevance during the ignition phase to becoming one of the most influential parameters affecting the spread. This discrepancy is potentially related to the amount of heat flux heating up the material in the two different stages. Throughout ignition, the sample receives substantial 65 kW/m^2 plus the heat feedback from the flame. Under the influence of a higher heat flux, thermal equilibrium is reached more quickly. The rapid heating minimises the time period during which the insulation layer can impact the temperature distribution within the sample, which, in turn, impacts pyrolysis and the HRR. In contrast, during the spread phase, the total heat flux is considerably lower, and the overall system is less dominated by the intensity of the heat input. Therefore, changes in the specific heat of the insulation material, which directly affect its ability to store thermal energy, can have a comparatively larger impact on the temperature distribution of the sample.

The existence of a transition in parameter importance, defined by the end of the ignition phase and the start of a self-sustained spread, highlights the differences between the Cone Calorimeter and the flame spread heating conditions. That is because both in the ignition phase of the flame spread simulation, and in the Cone Calorimeter setup, the sample is subject to a considerably higher heat flux than that of the self-sustained spread. As previously discussed, the elevated heat flux tends to dominate the thermal behaviour of the system, effectively masking the effects of changes in parameters that are crucial to the spread, such as the specific heat of the insulation material. Consequently, given the Cone Calorimeter insensitivity to changes in this parameter, attempting its estimation through an IMP based on the Cone Calorimeter would involve a high level of uncertainty. This implies that not all parameters that are important to the flame spread can be well estimated with IMPs based solely on the Cone Calorimeter setup.

3.2. Effects on the RMSE and ROS

The effects of the 15 input parameters on the RMSE and on the ROS are presented respectively in Figs. 9(a) and 9(b). Since both are single-value outputs, sensitivities are described by a single set of indices for each case. Confidence intervals are presented by error bars. For indices very close to one, such as the ST of specific heat at 480 °C in Fig. 9(a), confidence intervals can indicate that possible values for the associated sensitivity index exceed unity. In general, increased sample sizes ensure narrower confidence intervals and indices that fall within the expected range of 0 to 1. Nonetheless, the presented results were considered satisfactory for the objectives of this work. This consideration was supported by a preliminary analysis of the effect of the sample size on the indices estimation for the Cone Calorimeter case, encompassing 8192, 16384, 32768, 65536, and 131,072 samples. The results can be found on the supplementary material data [22]. Moreover, some S1 indices present negative values, which are due to numerical artefacts in the estimates. This issue has been reported before as a common characteristic of Saltelli's method that is often associated with the value of the index being close to zero.

As can be seen from Fig. 9(a), the specific heat at 150 °C and at 480 °C are effectively the only two input parameters that significantly affect the RMSE (calculated between simulated HRR curves of the Cone Calorimeter and experiment, see Section 2.3.2). Moreover, the influence is characterised by strong interaction effects, due to the negligible S1 and dominant ST values. This means that the effect on the RMSE highly depends on how the values of these two parameters are combined. Indeed, the combination between the specific heat at 150 °C and at 480 °C defines the slope of the linear curve that relates

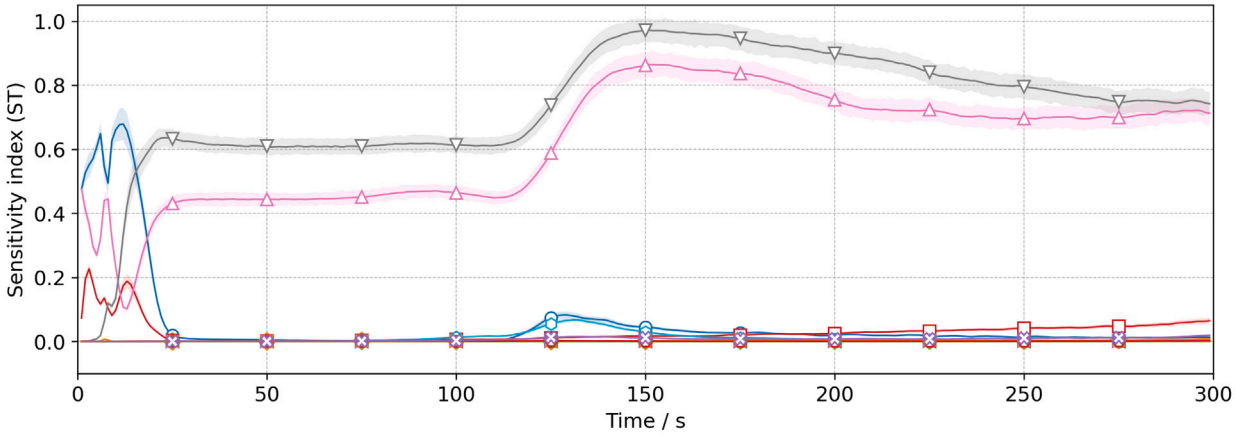

(a) Total-order effects on the HRR of the Cone Calorimeter simulation.

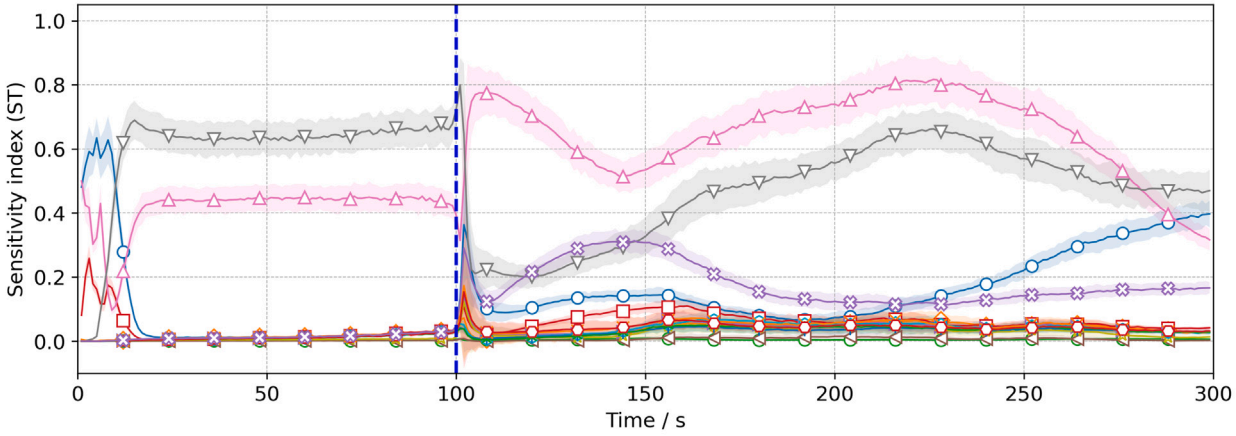

(b) Total-order effects on the HRR of the flame spread simulation (zoom up to 300 seconds).

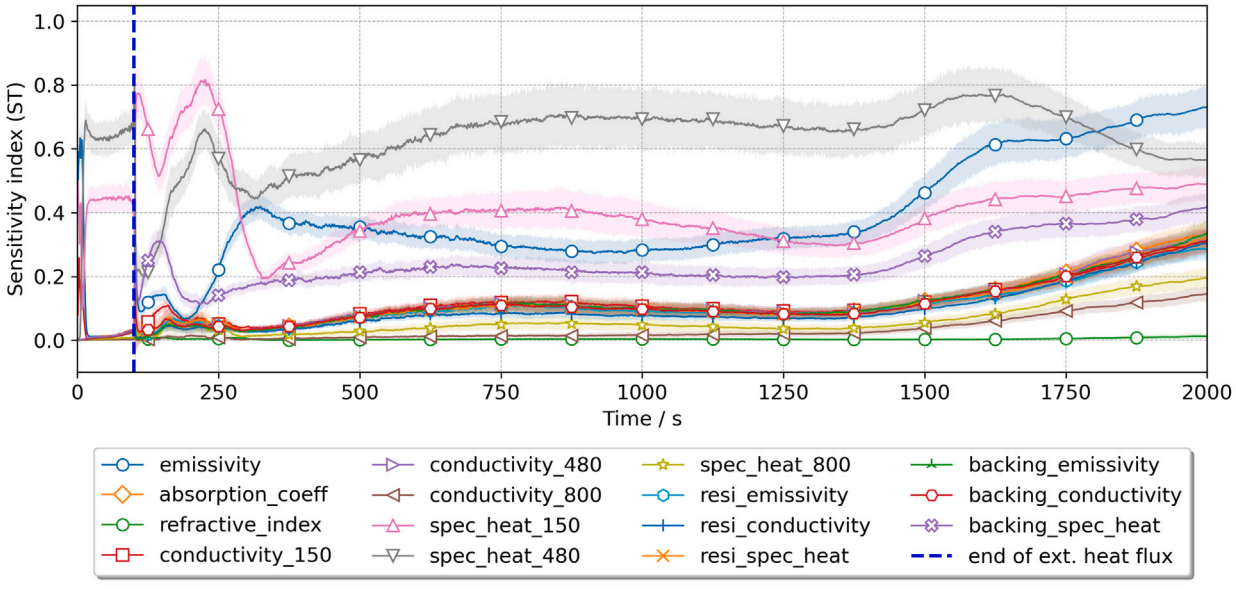

(c) Total-order effects on the HRR of the flame spread simulation.

Fig. 7. Time-series of ST indices, indicating total-order effects on the HRRs. The vertical blue dashed line in plots (b) and (c) highlight the end of the external heat flux at 100 seconds in the flame spread simulation.

the two values in the piecewise linear function, see Fig. 2(b). The slope is in turn related to how fast the change in the specific heat will occur as a function of temperature. For example, if the specific heat changes from a low value to a high value abruptly, suddenly more energy is required to cause the temperature of the material to change. In this case, this would reduce the local rates of pyrolysis, slowing down

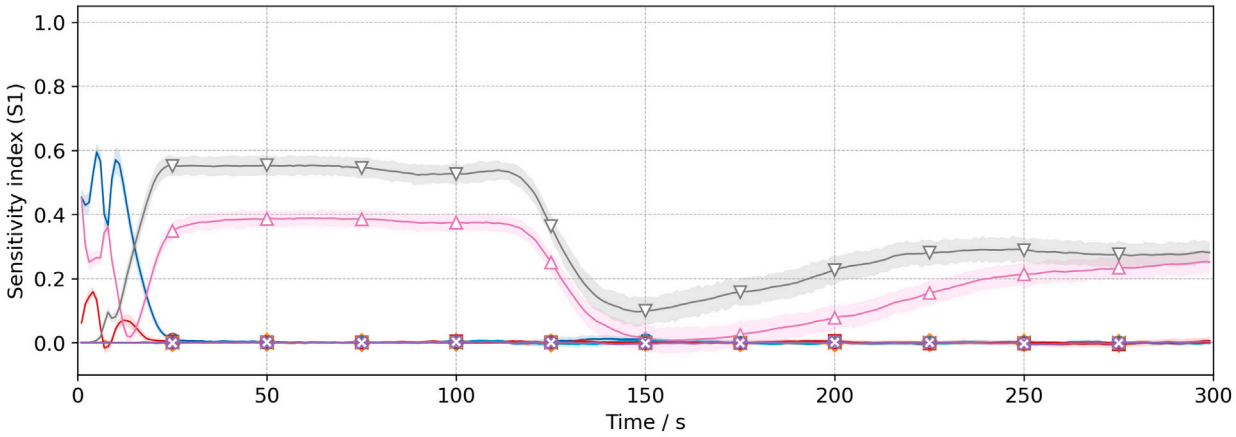

(a) Main effects on the HRR of the Cone Calorimeter simulation.

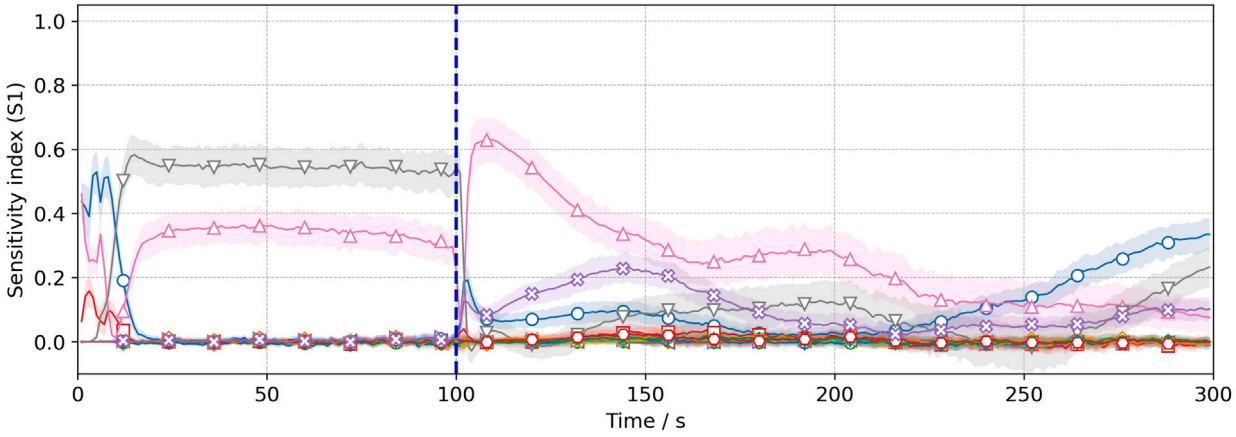

(b) Main effects on the HRR of the flame spread simulation (zoom up to 300 seconds).

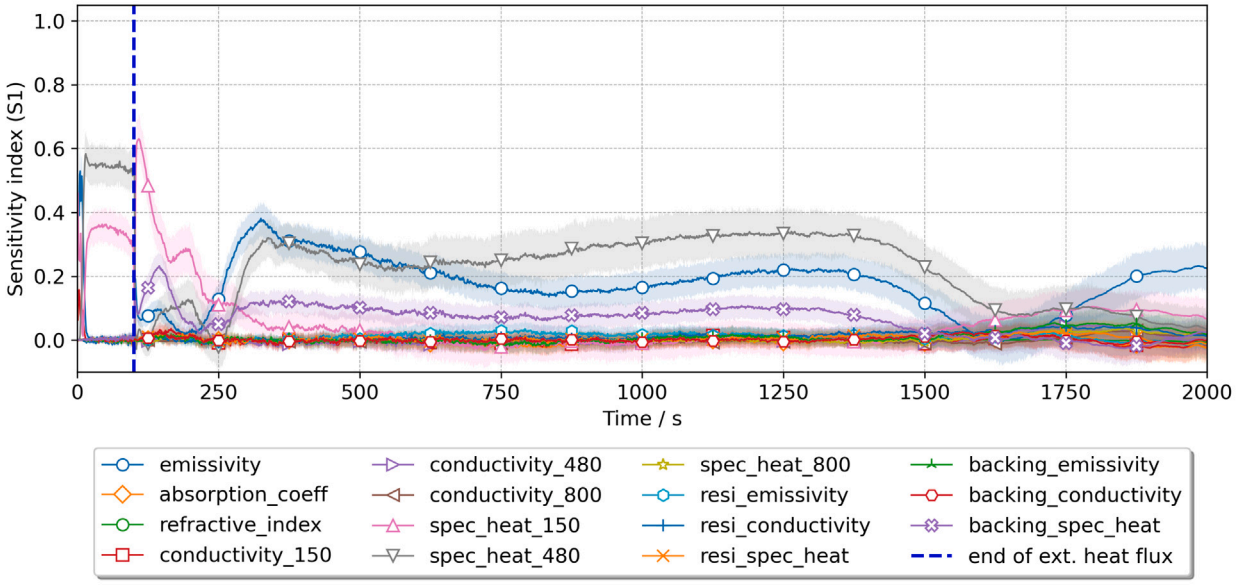

(c) Main effects on the HRR of the flame spread simulation.

Fig. 8. Time-series of S1 indices, indicating main effects on the HRRs. The vertical blue dashed line in plots (b) and (c) highlight the end of the external heat flux at 100 seconds in the flame spread simulation.

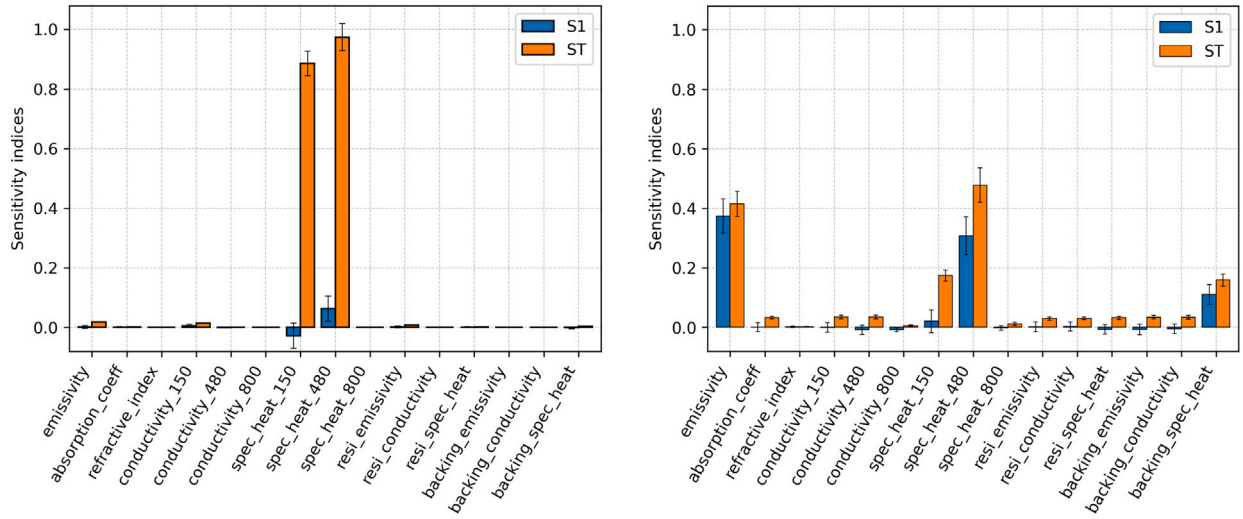

(a) Effects on the RMSE - Cone Calorimeter simulation.

(b) Effects on the ROS - horizontal flame spread simulation.

Fig. 9. Sobol sensitivity indices indicating the effects of the 15 input parameters on the single-value outputs of the two simulations setups.

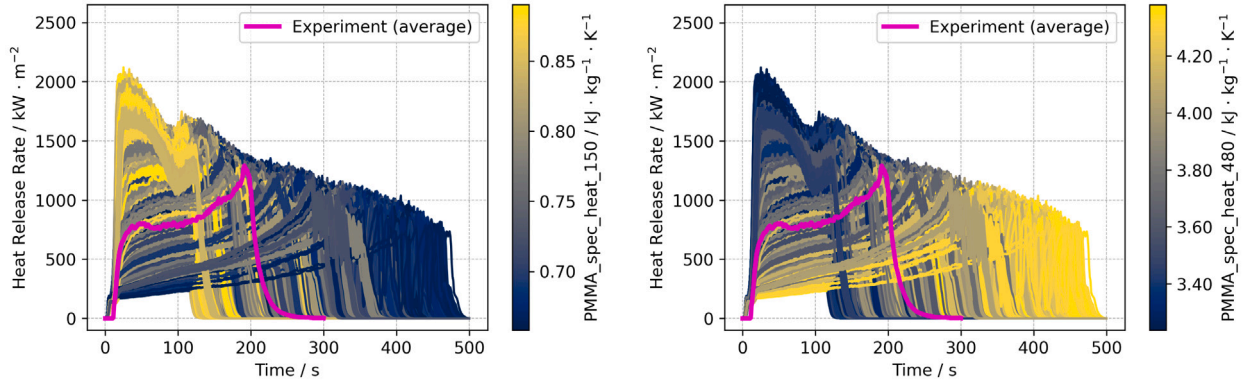

(a) Colour map: samples of specific heat at 150 °C.

(b) Colour map: samples of specific heat at 480 °C.

Fig. 10. Visualisation of the interaction effects between the time-dependent values of specific heat on the HRR curves of the Cone Calorimeter simulation.

the production of combustible gases which then burn, releasing heat. Ultimately, this translates into flattened HRR curves. This change in the HRR shape determines how much each simulated curve deviates from the experimental HRR curve. This effect is then captured by the RMSE.

A complementary approach to visualise these strong interaction effects is presented in Fig. 10. The plots show all 131,072 simulated HRR curves from the Cone Calorimeter, with colours representing their respective sample values of specific heat at 150 °C in Fig. 10(a), and specific heat at 480 °C in Fig. 10(b). The samples of these parameters are used to create colour maps, which are then assigned to the HRR curves. This provides a visual representation of the specific heat values used in each simulation. The colour map illustrates how the shape of the HRR is influenced by combinations of values of these two crucial parameters. For example, it becomes evident that flattened HRR curves, which also burn for longer times, result from a combination of low values of specific heat at 150 °C (dark blue in Fig. 10(a)), and high values of specific heat at 480 °C (yellow in Fig. 10(b)). The magenta curve is the HRR from Cone Calorimeter experiments, taken as reference for determining the RMSE.

In comparison to what is observed for the RMSE, a different scenario of sensitivities is identified for the ROS, as depicted in Fig. 9(b). The ST and S1 indices show that PMMA emissivity and specific heat at 480 °C are, in order, the two most important parameters to affect the ROS amongst the ones investigated here, and no meaningful interaction effects are observed. The higher importance of PMMA emissivity to the ROS is reasonable, given that in the model, the flame is considered optically thick, and as such, radiation is assumed to control the steady-state propagation process by preheating the unburnt sample ahead of the flame. The optically thick assumption for horizontally spreading flames over PMMA samples of this size is supported by the heat flux measurements reported by Jiang et al. [14]. In their work, for a PMMA sample of the same thickness and similar width (10 cm), radiation was the dominant mode of heat transfer from the flame to the solid. It should be noted however, that the lower part of a real flame, where the leading edge is located, the flame is usually weakly radiative and dominated by convection, therefore being considered optically thin. For PMMA flames, this characteristic is perceived by a transparent faint blue colour at the flame base, see e.g. Morrisset et al. [32]. Since this effect is not accounted for in the model, dependence on radiation might be overpredicted, which would be reflected in an increased sensitivity of the ROS to changes in the material emissivity. Moreover, uncertainties related to the specified radiative fraction, which controls the amount of radiated heat from the flame, might contribute as well to an artificially enlarged sensitivity to emissivity. Amongst the input parameters with lower importance, the specific heat of the insulation material and the specific heat of PMMA at 150 °C are the two most important ones.

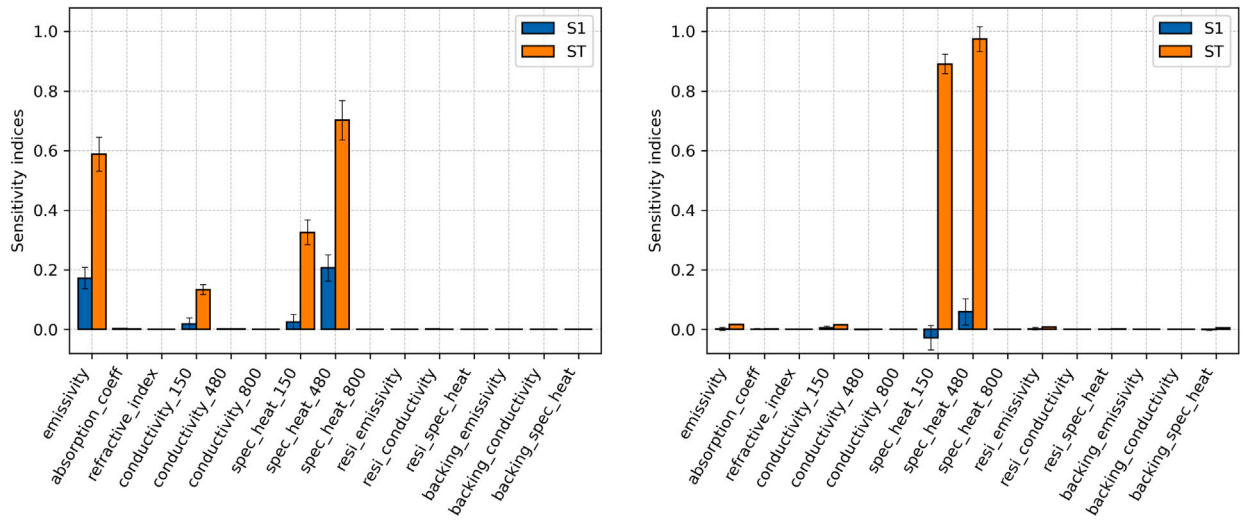

(a) Effects on the RMSE over 0 - 20 seconds.

(b) Effects on the RMSE over 20 - 300 seconds.

Fig. 11. Sobol sensitivity indices indicating the effect of input parameters on the RMSE calculated over different stages of the Cone Calorimeter simulation.

At this point, it is important to emphasise what is the main implication to the IMP revealed by the sensitivity indices shown in Fig. 9. The RMSE is a common approach to measure deviations between two sets of data, and therefore it is commonly used as cost function during the inverse modelling [31]. This means that whatever influence an input parameter has on the model output (in this case, the HRR of the Cone Calorimeter), it should be reflected in the RMSE for an effective estimation. However, comparison of Figs. 7(a) and 9(a) suggests that the initial importance of PMMA emissivity and specific heat at 150 °C to the HRR (Fig. 7(a)) is not manifested to the same level in the RMSE. This implies that not only is it necessary that the direct model output is sufficiently sensitive to the inputs that are important to the flame spread, but also that the cost function is. In this regard, neither of these two requirements were met, since there are several other input parameters affecting the flame spread (reflected both in the HRR and ROS) that have little or no importance to the HRR in the Cone Calorimeter and/ or to the RMSE.

Motivated by the results shown in Fig. 7(a), where differences exist between the initial 20 seconds and the rest of the simulation, additional SAs were conducted taking two different extracts of the RMSE as output of interest. In the first extract, the RMSE is calculated up to 20 seconds of simulation time (RMSE-0-20) and deviations to the experimental data are calculated accordingly up to the 20th second. Similarly, in the second extract, the RMSE is calculated from 21 to 300 seconds (RMSE-20-300). Sensitivity indices are presented for these two approaches in Figs. 11(a) and 11(b), respectively. As expected, the results for each extract present a clear correspondence to the different sensitivity profiles shown in Fig. 7(a). The rank of parameter importance up to 20 seconds to the HRR is equivalent to the rank shown in Fig. 11(a) to the RMSE-0-20. The same is true for the second extract. The dominance of the specific heat values after the 20th second is reflected in the RMSE-20-300. This analysis helps to understand why the default RMSE seems not to be significantly influenced by any other parameter than the specific heat at 150 °C and at 480 °C. Since the importance of these two values is higher for the most part of the simulation, the brief importance of the other parameters (emissivity, conductivity at 150 °C) gets diluted when the whole HRR time-series is condensed in a single RMSE value.

These observations highlight the importance of making use of cost functions that are as sensitive as the model outputs that they intend to represent. This way, a new design of cost functions can be defined. Instead of optimising for the global RMSE, which may be dominated by only a subset of sensitive parameters, a combination of RMSE at different phases of the experiment may cover the full set of sensitive parameters. In the case of the Cone Calorimeter investigated here, this would involve considering partial RMSEs calculated up to the initial 20 s and another one calculated from 21 to 300 s of the experiment, such that the estimations of emissivity and conductivity at 150 °C are improved. Still, this strategy would not be sufficient to estimate well all parameters that are important to the spread, such as the specific heat of the insulation material, as the Cone Calorimeter is insensitive to it. Alternatively, optimisation targets measured from bench-scale flame spread experiments could be employed rather than those from the Cone Calorimeter. Specifically, the ROS could be measured and incorporated as target in the IMP, enabling a comparison with the simulated ROS calculated for example as introduced in Section 2.2. Yet, this strategy may be computationally expensive and further research on this idea is needed, particularly because it does not exist to date a typical bench-scale flame spread experiment as the Cone Calorimeter that provides standardised guidelines for measuring the HRR and the MLR along with the ROS. A good starting point towards such an experiment could be the standard reaction-to-fire test methods described in ISO 5658-2, ISO 9239-1, and ISO 12468-1, which could be adapted to allow for the measurement of such quantities as well as analysis of samples at the Cone Calorimeter scale.

All in all, in addition to revealing differences between the Cone Calorimeter and the horizontal flame spread setups, the SA on the Cone Calorimeter disclosed that only 4 out of 15 input parameters have non-negligible influences. In terms of thermophysical properties, these four parameters are in fact three: emissivity, conductivity, and specific heat of the PMMA sample. This observation is similar to the results found in the work of Fleurotte et al. [16], who conducted a SA based on the Morris method to determine which parameters are more important to the HRR of a Cone Calorimeter model. In their work, PMMA emissivity and specific heat capacity are among the most influential parameters, along with activation energy and density. Yet, it is important to bare in mind that other parameters which were not included in the SAs can play a significant role in the model, and their effects on the output must be assessed in case they are included in the optimisation strategy. Thus, ranking the parameters according to their importance allows for model simplification, by excluding the non-influential inputs and/ or fixing them in the optimisation. This approach can potentially reduce the computing time by orders of magnitude, depending on the characteristic of the optimisation method used [2].

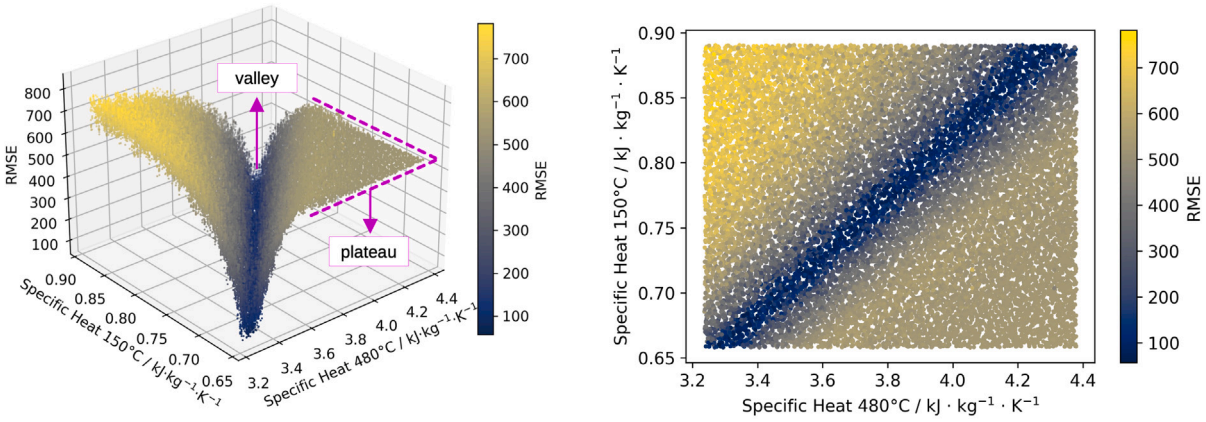

(a) 3-D surface. (b) 2-D projection.

Fig. 12. Values of RMSE plotted against its two most influential parameters.

3.3. Scatterplots

The calculated values of RMSE in the Cone Calorimeter setup are plotted against the sampled values of the two most influential parameters: the specific heat at 150 °C and at 480 °C. The result is a 3-D plot, where a well-defined surface allows the graphical interpretation of the interaction effects between the two inputs on the RMSE, see Fig. 12(a). The 3-D surface reveals a dark-blue valley for which the RMSE values are minimised when certain combinations of the two parameters are taken. A 2-D projection over the axes of input parameters shown in Fig. 12(b) clarifies that such combinations belong to a linear shaped subset of samples in their parameter space. Another interesting region in the 3-D surface is the plateau formed by nearly constant values of RMSE. The existence of a plateau reveals a significant portion of the input space that leads to no meaningful change in the RMSE. This is particularly important for the optimisation, because it can decrease its efficiency and lead to convergence to local minima.

A very similar relation to the one shown in Fig. 12, is presented in the work of Batiot et al. [17], where the effects of two interacting parameters on the quadratic error is discussed also in the context of their consequences to the IMP and the optimisation. In their work, the two interacting parameters are the pre-exponential factor A and the activation energy E of the Arrhenius equation, and the quadratic error is calculated over the material MLR. Batiot et al. [17] used the Sobol indices to discuss the well-known compensation effect between A and E in terms of the interaction effects captured by the second-order index. Given the similarity between the applied methodologies and the produced outcomes, it could be said also here that the linear relation between values of specific heat at 150 °C and at 480 °C translate into compensation effects and trade-offs during the IMP.

The 2-D scatterplots of the ROS versus emissivity and specific heat at 480 °C are shown respectively in Figs. 13(a) and 13(b). The parameters are the two most influential ones to impact the ROS, according to what is shown in Fig. 9(b). It can be seen that the relationship between the ROS and each individual parameter can be approximated by a linear function. However, whereas increasing values of emissivity act to increase the ROS, an opposite effect on the ROS is observed when the specific heat at 480 °C is increased, similarly to what was previously discussed in Section 3.2. This behaviour is coherent to the modelling of heat transfer mechanisms. The higher the value of emissivity is, the more heat by radiation is absorbed by the material. More absorbed heat causes local temperatures in the material to rise faster, enhancing the pyrolysis rates. Increased pyrolysis rates accelerates the production of fuel gases which subsequently burn, releasing more heat which then feeds the positive feedback loop that sustains the spread. On the other hand, higher specific heats make it more difficult to heat up the material and thus they lower the pyrolysis rate, thereby reducing the ROS, as can be seen in Fig. 13(b).

4. Conclusions

The SAs conducted in this study provided meaningful information on the differences in parameter importance to a Cone Calorimeter and a flame spread simulation conducted with FDS. Sobol indices of the input parameters suggested that the Cone Calorimeter simulation is not sufficiently sensitive to all of the parameters that are important to the flame spread. This is an issue because unimportant parameters to the Cone Calorimeter simulation will be estimated with higher degree of uncertainty during the IMP, which is then carried over to the flame spread simulation. In addition, it was revealed that the brief importance of some parameters in the Cone Calorimeter is diminished when the temporal development of the HRR is summarised in a single-value output, as the global RMSE value. Only the values of specific heat at 150 °C and at 480 °C seemed to influence the RMSE through strong interaction effects, whereas the importance of the remaining parameters is negligible. Moreover, the relation between the RMSE and its two most important parameters presented by scatterplots helped to visually identify subsets of the input space that could lead to minimised values of RMSE or convergence to local minima during an IMP. A possible solution for a more effective parameter estimation seems to rely on a combination of RMSE calculated at different phases of the experiment, such that the full set of sensitive parameters is covered. Still, many of the important parameters to the flame spread are practically unimportant throughout the Cone Calorimeter simulation. This limitation could possibly be overcome by IMPs based on bench-scale flame spread experiments, in which the ROS is taken as target in the optimisation, ensuring that the parameters relevant to the spread are more accurately estimated.

CRediT authorship contribution statement

Tássia L.S. Quaresma: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Tristan Hehnen:** Writing – review & editing, Validation, Conceptualization. **Lukas Arnold:** Writing – review & editing, Validation, Supervision, Software, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

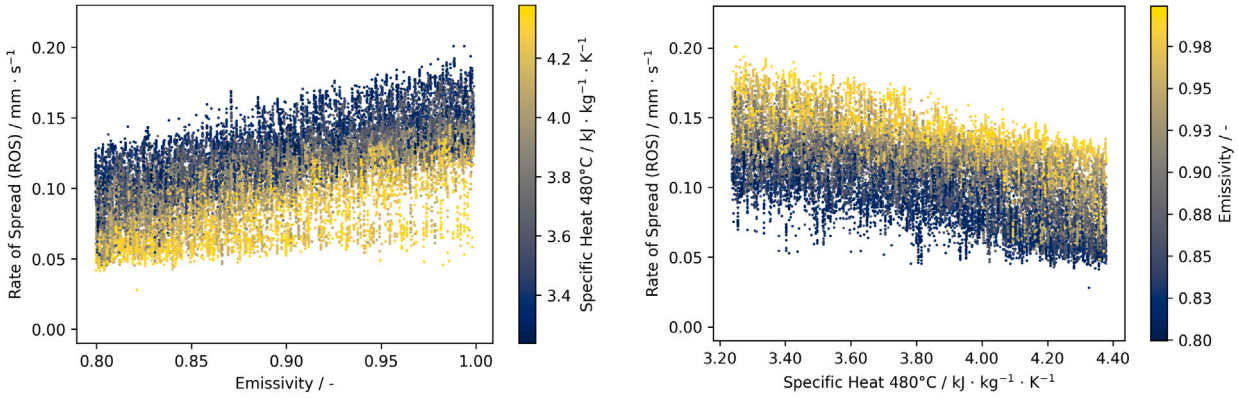

(a) (b)

Fig. 13. Scatterplots of (a) ROS plotted against emissivity and coloured by the specific heat at 480 °C; and (b) ROS plotted against specific heat at 480 °C and coloured by emissivity.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data is made publicly available here: https://doi.org/10.5281/zenodo.7618897.

Acknowledgements

We gratefully acknowledge the computing time granted through the project on the CoBra-system, funded by the German Federal Ministry of Education and Research with the grant number 13N15497. This research was partially funded by the German Federal Ministry of Education and Research with the grant number 13N15497.

References

[1] K. McGrattan, S. Hostikka, J. Floyd, R. McDermott, M. Vanella, Fire Dynamics Simulator (version FDS6.7.6-810-ge59f90f-HEAD) – User's Guide, US Department of Commerce, Technology Administration, NIST, 2005, http://dx.doi.org/10.6028/NIST.SP.1019.

[2] T. Hehnen, L. Arnold, PMMA pyrolysis simulation – from micro- to real-scale, Fire Saf. J. 141 (2023) 103926, http://dx.doi.org/10.1016/j.firesaf.2023.103926.

[3] A. Viitanen, S. Hostikka, J. Vaari, CFD simulations of fire propagation in horizontal cable trays using a pyrolysis model with stochastically determined geometry, Fire Technol. 58 (5) (2022) 3039–3065, http://dx.doi.org/10.1007/s10694-022-01291-6.

[4] G. Fiola, D. Chaudhari, S. Stoliarov, Comparison of pyrolysis properties of extruded and cast poly (methyl methacrylate), Fire Saf. J. 120 (2021) 103083, http://dx.doi.org/10.1016/j.firesaf.2020.103083.

[5] T. Hehnen, L. Arnold, S. La Mendola, Numerical fire spread simulation based on material pyrolysis—An application to the CHRISTIFIRE phase 1 horizontal cable tray tests, Fire 3 (3) (2020) 33, http://dx.doi.org/10.3390/fire3030033.

[6] F. Yang, C. Rippe, J. Hodges, B. Lattimer, Methodology for material property determination, Fire Mater. 43 (6) (2019) 694–706, http://dx.doi.org/10.1002/fam.2721.

[7] T. Nyazika, M. Jimenez, F. Samyn, S. Bourbigot, Pyrolysis modeling, sensitivity analysis, and optimization techniques for combustible materials: A review, J. Fire Sci. 37 (4–6) (2019) 377–433, http://dx.doi.org/10.1177/0734904119852740.

[8] T. Beji, B. Merci, Numerical simulations of a full-scale cable tray fire using small-scale test data, Fire Mater. 43 (5) (2019) 486–496, http://dx.doi.org/10.1002/fam.2687.

[9] L. Arnold, T. Hehnen, P. Lauer, C. Trettin, A. Vinayak, PROPTI – A generalised inverse modelling framework, J. Phys. Conf. Ser. 1107 (3) (2018) 032016, http://dx.doi.org/10.1088/1742-6596/1107/3/032016.

[10] L. Chad, S. Stoliarov, J. Lord, I. Leventon, A methodology for predicting and comparing the full-scale fire performance of similar materials based on small-scale testing, Fire Mater. 42 (7) (2018) 710–724.

[11] S. Hostikka, A. Matala, Pyrolysis model for predicting the heat release rate of birch wood, Combust. Sci. Technol. 189 (8) (2017) 1373–1393, http://dx.doi.org/10.1080/00102202.2017.1295959.

[12] E. Kim, N. Dembsey, Parameter estimation for comprehensive pyrolysis modeling: guidance and critical observations, Fire Technol. 51 (2) (2015) 443–477, http://dx.doi.org/10.1007/s10694-014-0399-0.

[13] O. Korobeinichev, M. Gonchikzhapov, A. Tereshchenko, I. Gerasimov, A. Shmakov, A. Paletsky, A. Karpov, An experimental study of horizontal flame spread over PMMA surface in still air, Combust. Flame 188 (2018) 388–398, http://dx.doi.org/10.1016/j.combustflame.2017.10.008.

[14] L. Jiang, C.H. Miller, M.J. Gollner, J. Sun, Sample width and thickness effects on horizontal flame spread over a thin PMMA surface, Proc. Combust. Inst. 36 (2) (2017) 2987–2994, http://dx.doi.org/10.1016/j.proci.2016.06.157.

[15] A. Karpov, O. Korobeinichev, A. Shaklein, A. Bolkisev, A. Kumar, A. Shmakov, Numerical study of horizontal flame spread over PMMA surface in still air, Appl. Therm. Eng. 144 (2018) 937–944, http://dx.doi.org/10.1016/j.applthermaleng.2018.08.106.

[16] M. Fleurotte, O. Authier, S. Bourbigot, G. Fontaine, G. Debenest, A. Amokrane, Sensitivity analysis applied to the pyrolysis of charring and non-charring materials under a cone calorimeter, J. Fire Sci. 40 (6) (2022) 405–443, http://dx.doi.org/10.1177/07349041221134486.

[17] B. Batiot, T. Rogaume, A. Collin, F. Richard, J. Luche, Sensitivity and uncertainty analysis of Arrhenius parameters in order to describe the kinetic of solid thermal degradation during fire phenomena, Fire Saf. J. 82 (2016) 76–90, http://dx.doi.org/10.1016/j.firesaf.2016.03.007.

[18] Y. Ding, I. Leventon, S. Stoliarov, An analysis of the sensitivity of the rate of buoyancy-driven flame spread on a solid material to uncertainties in the pyrolysis and combustion properties. Is accurate prediction possible?, Polym. Degrad. Stab. 214 (2023) 110405, http://dx.doi.org/10.1016/j.polymdegradstab.2023.110405.

[19] I. Sobol, Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, Math. Comput. Simul. 55 (1–3) (2001) 271–280, http://dx.doi.org/10.1016/S0378-4754(00)00270-6.

[20] A. Saltelli, Making best use of model evaluations to compute sensitivity indices, Comput. Phys. Commun. 145 (2) (2002) 280–297, http://dx.doi.org/10.1016/S0010-4655(02)00280-1.

[21] K. Ujjwal, J. Aryal, S. Garg, J. Hilton, Global sensitivity analysis for uncertainty quantification in fire spread models, Environ. Model. Softw. 143 (2021) 105110, http://dx.doi.org/10.1016/j.envsoft.2021.105110.

[22] T. Quaresma, T. Hehnen, L. Arnold, Sensitivity Analysis for an Effective Transfer of Material Properties from Cone Calorimeter to Flame Spread Simulations - Data Set, Forschungszentrum Jülich and Bergische Universität Wuppertal, Germany, 2023, http://dx.doi.org/10.5281/zenodo.7618897, Zenodo.

[23] I. Leventon, B. Batiot, M. Bruns, S. Hostikka, Y. Nakamura, P. Reszka, T. Rogaume, S. Stoliarov, Measurement and computation of fire phenomena (MaCFP) condensed phase material database, 2022, http://dx.doi.org/10.18434/mds2-2586, github.com/MaCFP/macfp-db. Commit: 7f89fd8.

[24] C.G. McCoy, J.L. Tilles, S.I. Stoliarov, Empirical Model of flame heat feedback for simulation of cone calorimetry, Fire Saf. J. 103 (2019) 38–48, http://dx.doi.org/10.1016/j.firesaf.2018.11.006, URL https://www.sciencedirect.com/science/article/pii/S0379711218303035.

[25] J. Quintiere, Principles of Fire Behaviour. Cengage Learning, CRC Press, Boca Raton, 1998.

[26] J. Vogelsang, L. Arnold, fdsreader - Fast and easy-to-use Python reader for FDS data, 2013, GitHub repository. https://github.com/FireDynamics/fdsreader.

[27] A. Saltelli, K. Aleksankina, W. Becker, P. Fennell, F. Ferretti, N. Holst, S. Li, Q. Wu, Why so many published sensitivity analyses are false: A systematic review of sensitivity analysis practices, Environ. Model. Softw. 114 (2019) 29–39, http://dx.doi.org/10.1016/j.envsoft.2019.01.012.

[28] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, D. Cariboni, M. Saisana, S. Tarantola, Global Sensitivity Analysis: the Primer, John Wiley & Sons, 2008.

[29] T. Iwanaga, W. Usher, J. Herman, Toward SALib 2.0: Advancing the accessibility and interpretability of global sensitivity analyses, Socio-Environ. Syst. Model. 4 (2022) 18155, https://sesmo.org/article/view/18155.

[30] J. Herman, W. Usher, SALib: An open-source Python library for sensitivity analysis, J. Open Source Softw. 2 (9) (2017) http://dx.doi.org/10.21105/joss.00097.

[31] P. Lauer, T. Hehnen, C. Trettin, F. Brännström, L. Arnold, Role of the cost function for material parameter estimation, in: Fire and Evacuation Modelling Technical Conference, 2020, http://dx.doi.org/10.5281/zenodo.4007595.

[32] D. Morrisset, R.M. Hadden, A. Law, Quantifying the controlling mechanisms of opposed flow flame spread: Influence of orientation, material, and external heating, Fire Saf. J. 142 (2024) 104048, http://dx.doi.org/10.1016/j.firesaf.2023.104048.