VARIANCE DECOMPOSITION OF MEDLI2 RECONSTRUCTED HEATING USING NEURAL NETWORKS

H. S. Alpert

AMA at NASA Ames
Entry Systems & Vehicle Engineering
Moffett Field, CA 94305

T. K. West

NASA Langley Research Center
Vehicle Analysis Branch
Hampton, VA 23681

K. A. Thomas

Stanford University
Dept. of Biomedical Data Science
Stanford, CA 94035

ABSTRACT

The Mars Entry, Descent, and Landing Instrumentation (MEDLI2) sensor suite collected data during entry of the Mars 2020 Perseverance rover into Mars' atmosphere. An inverse estimation of the backshell and heatshield surface aeroheating was performed, using the data from the MEDLI2 Instrumented Sensor Plugs, a network of thermocouples embedded within the thermal protection system across the aeroshell. Monte Carlo analysis was conducted to assess the sensitivity of the surface heat rate, temperature, and heat load to uncertainties in thermocouple depth and material properties. In this paper, a variance decomposition method using Sobol indices was employed to understand the relative contributions of each uncertainty parameter. Performing this analysis using results from the inverse analysis tool *FIAT_Opt* was found to require incredibly high computation time, and thus machine learning models were trained and evaluated as a surrogate model for *FIAT_Opt*. This paper demonstrates that machine learning models can be an efficient, accurate alternative to state-of-the-art inverse analysis tools like *FIAT_Opt*, especially for computationally-expensive processes. Using these models, the sensitivity analysis showed that uncertainties in heat capacity and thermal conductivity were the main drivers for the overall uncertainty in peak reconstructed heating and heat load.

Index Terms— Inverse estimation, environment reconstruction, aeroshell, heat flux, neural network

1. INTRODUCTION

The Mars Entry, Descent, and Landing Instrumentation (MEDLI2) sensor suite collected data during entry of the Mars 2020 Perseverance rover into Mars' atmosphere on February 18, 2021. This suite included a network of MEDLI2 Instrumented Sensor Plugs (MISPs) on the heatshield and backshell. Each MISP was comprised of a cylinder made of Thermal Protection System (TPS) material with 1-3 embedded thermocouples (TCs), which was then flush mounted into the heatshield or backshell. The 6 MISPs on the backshell were made of SLA-561V and the 11 on the heatshield were made of PICA, matching the materials in which the MISPs were embedded; their locations are shown in Fig. 1. The backshell MISPs were intended to record the heating in regions of the attached and detached flow, while the distribution of the heatshield MISPs was driven by the desire to capture turbulent flowfield transition. Further detailed information about the MISPs and their embedded thermocouples can be found in [1].

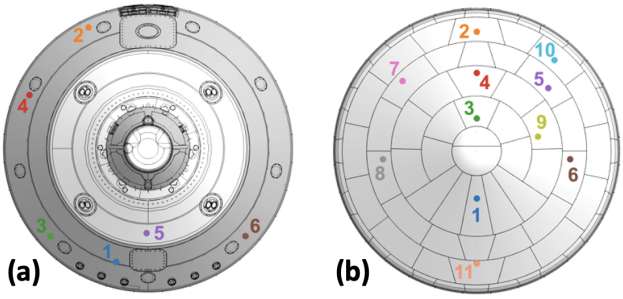

Fig. 1. Locations of MISPs on MEDLI2 (a) backshell and (b) heatshield.

Data from these in-depth TCs were used to reconstruct the aeroheating environment of the vehicle throughout entry. Surface heating was posed as an inverse problem, with the goal of estimating the surface heating by minimizing an objective function of the difference between MISP temperature measurements during flight and the temperature predictions derived from the Fully Implicit Ablation and Thermal response (FIAT) program [2]. FIAT takes the aerothermal environment and aeroshell material stack-up as the inputs and uses equilibrium chemistry models to output the thermal response of the material throughout its depth. To implement the opposite approach, where the in-depth temperature is the input and the aeroheating environment is the output, a NASA-developed program called *FIAT_Opt* is used. *FIAT_Opt* runs through multiple different environments until the output temperature at the TC depth closely matches the flight data. The results from this analysis showed that the backshell experi-

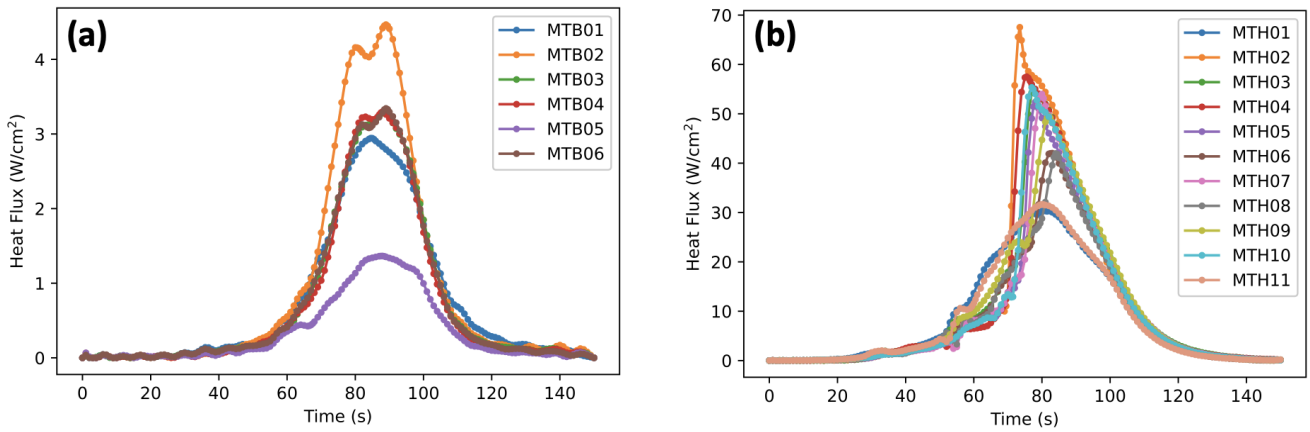

Fig. 2. Reconstructed surface heat flux for MISPs on the (a) backshell and (b) heatshield.

enced a peak combined (radiative and convective) heating of 4.5 W/cm^2 at the backshell location 2, while the heatshield reached 68.4 W/cm^2 at heatshield location 2, as shown in Fig. 2. For the same two locations, the total heat load over the course of the trajectory was 138.3 J/cm^2 and 1672.5 J/cm^2 respectively [3].

A Monte Carlo approach was employed to evaluate the uncertainty associated with the reconstructed surface heating, as there is considerable uncertainty in material properties including virgin and char density (ρ_v and ρ_c), thermal conductivity (k_v and k_c), heat capacity ($C_{p,v}$ and $C_{p,c}$), and emissivity (ϵ_v and ϵ_c), as well as the depth of the TC within the MISP (d). These uncertainties, shown in Table 1, were quantified based on flight-lot material testing [4]. All were assumed to have Gaussian distributions except density, for which a uniform distribution was used.

For the Monte Carlo analysis, 2000 samples were generated with Latin hypercube sampling for each MISP, and the 95% confidence intervals were calculated, as shown in Fig. 3. The 95% confidence intervals on the peak heat flux on the backshell MISPs were $\pm$19% on average, with a range of $\pm$17-22% for the different MISPs. For the heatshield MISPs the average was $\pm$13% with a range of $\pm$7-16% for the various locations.

This paper presents the sensitivity of the reconstructed heat flux to the uncertainty in virgin and char density, thermal conductivity, heat capacity, emissivity, and the TC depth, with the objective of identifying which of these uncertainties are driving the uncertainty in the peak heat flux obtained from the Monte Carlo analysis. Section 2.1 discusses the variance decomposition method used for the sensitivity analysis and the computation time problem it posed for this study. Section 2.2 describes the machine learning approach that was implemented to mitigate the problem. Section 3.1 presents the performance (accuracy and consistency) of the neural networks and Section 3.2 presents the results of the sensitivity analysis for one of the MISPs. Conclusions and ongoing work are discussed in Section 4.

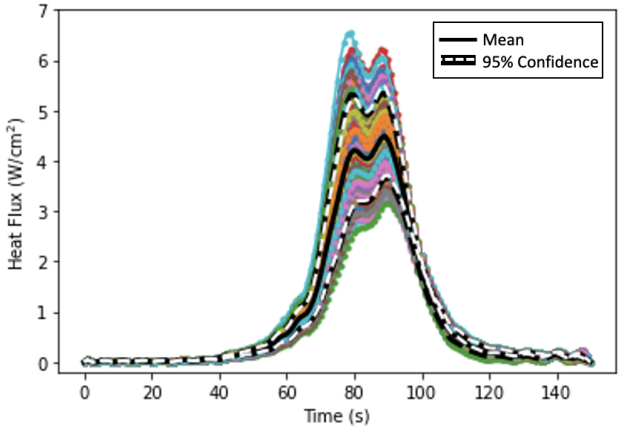

Fig. 3. Reconstructed surface heat fluxes for MTB02. All 2000 cases are plotted in addition to the mean and the the 95% confidence interval.

2. METHODS

2.1. Variance Decomposition

Variance decomposition is an approach used to quantify how much the output of a model can be attributed to uncertainty in each of the model inputs. One such method is to calculate Sobol indices, which quantify the amount of variation in the output that can be attributed to the variation in each of the inputs [5, 6]. This calculation requires two $(N \times k)$ sampling matrices, **A** and **B**, where N is the number of the samples and k is the number of factors, or inputs, to the model. Each of the N rows of the matrix has an associated output $f(A)_j$, where

Table 1. Uncertainties associated with backshell MISPs (SLA-516V) and heatshield MISPs (PICA). A normal distribution is denoted by "N" and uniform by "U".

Backshell						
	$\rho_v = \rho_c$	$C_{p,v} = C_{p,c}$	ϵ_v	ϵ_c	$k_v = k_c$	d
Distribution	U	N	N	N	N	N
$2\sigma/\mu$	$\pm$0.018*	0.15	0.15	0.15	0.13	0.003" **
Heatshield						
	$\rho_v = \rho_c$	$C_{p,v} = C_{p,c}$	ϵ_v	ϵ_c	$k_v = 0.8k_c^\dagger$	d
Distribution	U	N	N	N	N	N
$2\sigma/\mu$	$\pm$0.018*	0.05	0.03	0.03	0.12	0.003" **
*Upper/lower bound for the uniform distribution, **not** $2\sigma/\mu$.						
2σ **not $2\sigma/\mu$.						
†Uncertainty of k_c is found by scaling k_v by 15/12 (i.e., virgin and char are linked).						

j is the row number. The total sensitivity (S_T) for each of the k parameters is calculated as:

$$S_{Ti} = \frac{1}{N} \frac{\sum_{j=1}^{N} f(A)_j (f(A)_j - f(A_B^{(i)})_j)}{V(Y)} \quad (1)$$

where $V(Y)$ is

$$\frac{1}{N} \sum_{j=1}^{N} (f(A)_j - f_0)^2 \quad (2)$$

and i is the column number. The matrix $A_B^{(i)}$ is formed by replacing the i^{th} column in matrix **A** with the i^{th} column from matrix **B**; thus there are k total $A_B^{(i)}$ matrices. Because outputs are required for all N rows of all k $A_B^{(i)}$ matrices, as well as for the **A** matrix itself, this calculation requires $N \times (k + 1)$ *FIAT_Opt* runs. In this study there are 6 parameters with independent uncertainty distributions, hence k=6. An initial test case showed that the Sobol indices do not converge until nearly 30,000 iterations are run, as shown in Fig. 4. With a single *FIAT_Opt* run taking about 40 minutes on a single CPU and access to 20 CPUs at once, calculating the Sobol indices would take 41 days per MISP, and over a year for all MISPs. This demonstrates a clear need for a faster surrogate model that could be used in place of *FIAT_Opt*.

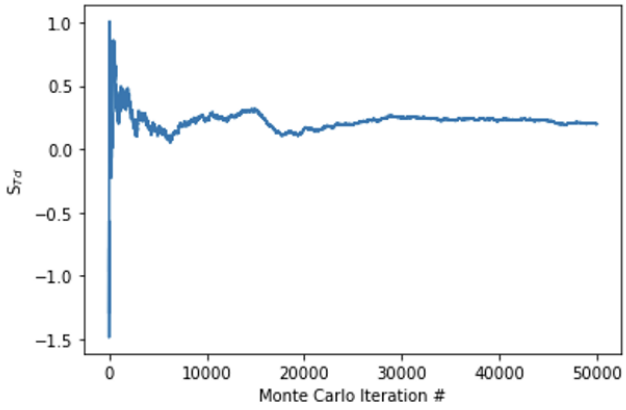

Fig. 4. Evolution of the Sobol index for TC depth over the course of 50,000 samples.

2.2. Machine Learning Approach

In order to significantly reduce the required computation time, machine learning was employed to develop a surrogate model that could be used in place of directly performing computations in *FIAT_Opt*. Each machine learning model was designed to take the same six parameters as inputs (listed in Table 1) and produce an estimate of peak heat flux or heat load as output. A set of 2500 *FIAT_Opt* runs were used to train, validate, and test these models. The values for each material property and the TC depth as well as the corresponding peak heat flux and heat load values were used from each run.

The 2000 runs that were already completed for backshell location 2 as part of the Monte Carlo analysis were randomly split into a training set of 1600 runs and a validation set of 400 runs. An additional 500 *FIAT_Opt* runs were performed to generate data for the test set. Many different model versions were trained with the training data. The model version that performed best on the validation data was kept as the final model. The final model was then evaluated on the test set to obtain the final performance metrics reported in Section 3. Splitting up the data in this way provides a more accurate estimate of how the final models will perform on data that they did not encounter during training.

Three types of machine learning models were trained with the training set: random forests, ridge regression, and deep neural networks (DNN). The neural networks outperformed the random forest and ridge regression models on the vali-

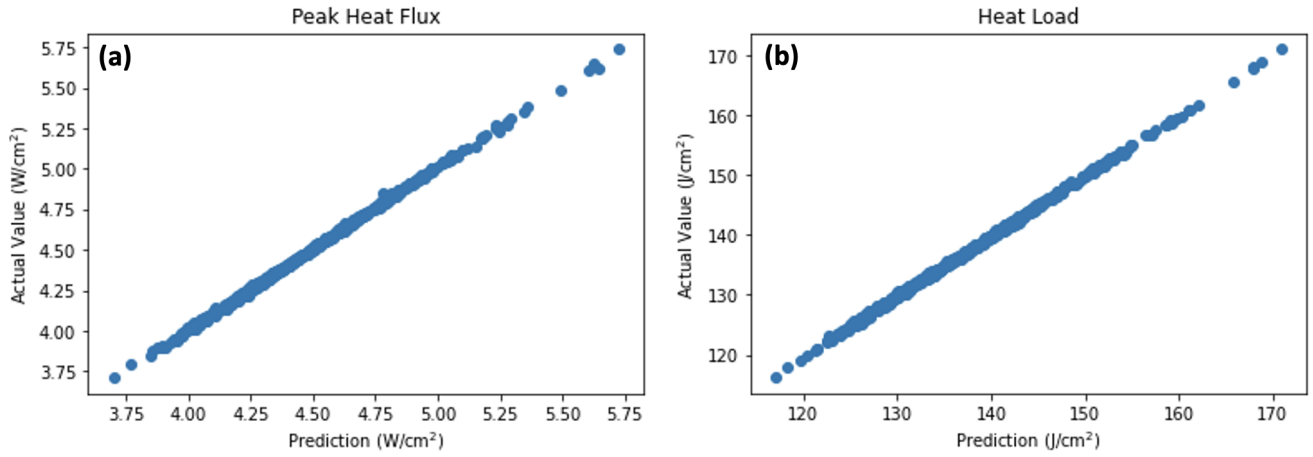

Fig. 5. Plots of actual values (calculated from *FIAT_Opt*) vs. predicted values (from the DNN) for the test set for (a) peak heat flux and (b) heat load.

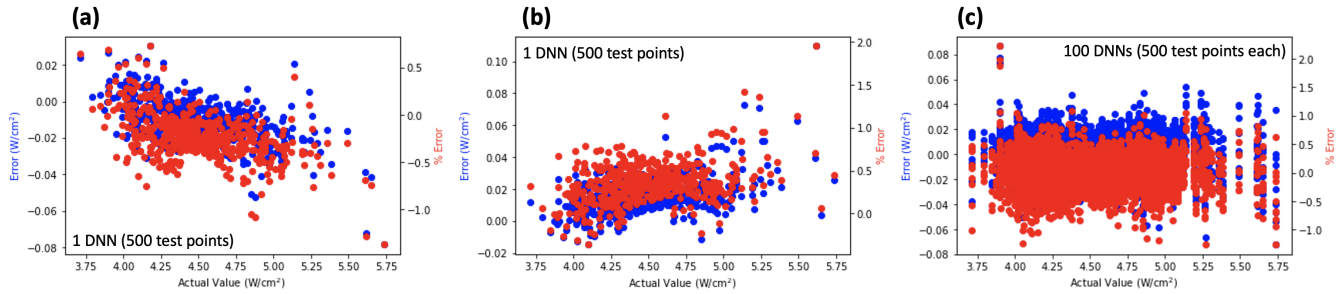

Fig. 6. Bland-Altman plots for (a) a DNN with a negative bias at higher heat fluxes (b) a DNN with positive bias at higher heat fluxes, and (c) all 100 DNNs.

dation set in early testing, so further hyperparameter tuning was only conducted with the neural networks. The final neural network models had a fully-connected architecture with four hidden layers, each with 64 nodes with rectified linear activation functions (ReLU). The output layer of each neural network had a single output. Mean absolute error (MAE) was used as the loss function in order to obtain models that were optimized to estimate *FIAT_Opt* outputs with as little error as possible. Each model was trained for 500 epochs (a single round of training in which each data point is used exactly once) with Adam optimization. This was implemented using Python version 3.8.3, Keras version 2.6.0, and Scikit-learn version 0.23.1.

Training a neural network is a stochastic process. The weights of a neural network are randomly initialized before training is started and the order in which samples from the training set are used to train the model can influence the weights' final values. To evaluate the variation in final model performance caused by these factors, 100 DNNs with the architecture and hyperparameters described above were trained with the training set and evaluated with the test set. Each DNN's MAE and mean average percent error (MAPE) on the test set was calculated. The average and standard deviation of these performance metrics were then calculated across all 100 DNNs.

3. RESULTS

3.1. Neural Network Performance

The DNNs accurately estimate the peak heat flux and heat load calculated by *FIAT_Opt*, as shown in Fig. 5. For the 100 DNNs trained to estimate peak heat flux, the average MAE was 0.0095 W/cm^2 with a standard deviation of 0.0036 W/cm^2, and the average MAPE was 0.2122% with a standard deviation of 0.0796%. The Bland-Altman plots (Fig. 6) show that some of the DNNs are biased in different ways; the DNN in Fig. 6a has a negative bias at higher heat fluxes, while the DNN in Fig. 6b has a positive bias at higher heat fluxes. However, across all 100 DNNs, the bias is fairly uniform (Fig. 6c). For the 100 DNNs trained to estimate heat load, the average

Table 2. Sobol indices for all parameters, with the mean and standard deviation from re-training the neural network 100 times.

	ρ	C_p	k	ϵ_v	ϵ_c	d
Peak Heat Flux						
Mean	0.014	0.550	0.251	0.122	0.058	0.071
Std Dev	0.0009	0.0045	0.0035	0.0025	0.0017	0.0018
Heat Load						
Mean	0.018	0.589	0.212	0.090	0.056	0.075
Std Dev	0.005	0.0082	0.0076	0.0045	0.0036	0.0023

MAE was 0.358 J/cm^2 and the average MAPE was 0.245% with a standard deviation of 0.123%.

Using the DNN as a surrogate model, variance decomposition with 40,000 samples was completed within minutes. There was minor variation in the Sobol indices between each time the DNN was retrained. For the 100 DNNs trained to estimate peak heat flux, the standard deviation of the greatest Sobol index was 0.005, or about 1% of the mean, while the maximum and minimum were 2.2% above and below the mean respectively. For the 100 DNNs trained to estimate heat load, the standard deviation on the greatest Sobol index was 0.008 (1.4% of the mean), the maximum was 3.2% above the mean, and minimum was 3.7% below the mean.

3.2. Sensitivity Analysis

The variance decomposition showed that the uncertainty in heat capacity is the greatest driver in uncertainty of peak reconstructed heat flux, with a Sobol index about twice that of the thermal conductivity (S_{TCp}=0.550 and S_{Tk}=0.251), which is the next greatest. With enough samples the Sobol indices sum to 1; thus the Sobol index for a given parameter is equivalent to the fraction of the uncertainty that can be attributed to that parameter. The results were similar for the heat load, with a slightly higher Sobol index for the heat capacity (S_{TCp}=0.589) and lower Sobol index for the thermal conductivity (S_{Tk}=0.212). The Sobol indices for all of the parameters are shown in Table 2. The results of this analysis were checked by re-running the Monte Carlo analysis, but with different uncertainty distributions. For the nominal case at backshell location 2, with 2σ uncertainties on thermal conductivity and heat capacity of 13% and 15% respectively, the standard deviation for the 2000 cases on the peak heat flux was 0.61 W/cm^2 and on the heat load was 12.3 J/cm^2. When the uncertainty in C_p remained the same but the uncertainty in k was cut in half, the resultant standard deviation on peak heat flux was 0.54 W/cm^2 and on the heat load was 11.4 J/cm^2. In the opposite situation (the uncertainty on k was kept nominal and the uncertainty in C_p was cut in half), the standard deviation on peak heat flux was 0.45 W/cm^2 and on the heat load was 9.6 J/cm^2. This analysis, shown in Fig. 7, points to the same conclusion as drawn from variance decomposition, that the uncertainty in C_p is the main driver of uncertainty in reconstructed peak heat flux. This exercise required days of computation time and does not provide the fraction of uncertainty attributable to each input, and thus was a successful confidence check but has many drawbacks compared to variance decomposition.

4. CONCLUSIONS

The 2000 Monte Carlo iterations conducted for preliminary uncertainty analysis were utilized to train a neural network that was accurate enough to be used in lieu of *FIAT_Opt* for significantly (i.e., 1000x) faster computation time. Three different machine learning models were assessed and the deep neural network was chosen, based on its low MSE when evaluated on the test set. These neural networks were found to be highly accurate in predicting peak heat flux and heat load. The neural network was re-trained 100 times and the Sobol indices were very consistent, showing that the uncertainty in the heat capacity was the main driver in the overall uncertainty of the peak reconstructed heat flux, and the thermal conductivity had the second most impact. Results such as these can be leveraged to provide requirements for material property measurements needed to improve the accuracy of surface heating prediction and ultimately lead to the reduction of design margins in the future. Additionally, the use of neural networks can be implemented more widely to drastically speed up future computations related to inverse heating estimation. Future work includes conducting a similar sensitivity analysis for all of the MISPs on both the backshell and the heatshield to see how the Sobol indices vary across the two materials (SLA-561V vs. PICA) and location.

5. REFERENCES

[1] H. Hwang, D. Bose, T. White, H. Wright, M. Schoenenberger, C. Kuhl, D. Trombetta, J. Santos, T. Oishi, C. Karlgaard, M. Mahzari, and S. Pennington, "Mars 2020 Entry, Descent, and Landing Instrumentation 2 (MEDLI2)," in *AIAA Thermophysics Conference*, 2016.

[2] Y. Chen and F. Milos, "Ablation and Thermal Response Program for Spacecraft Heatshield Analysis," *Journal*

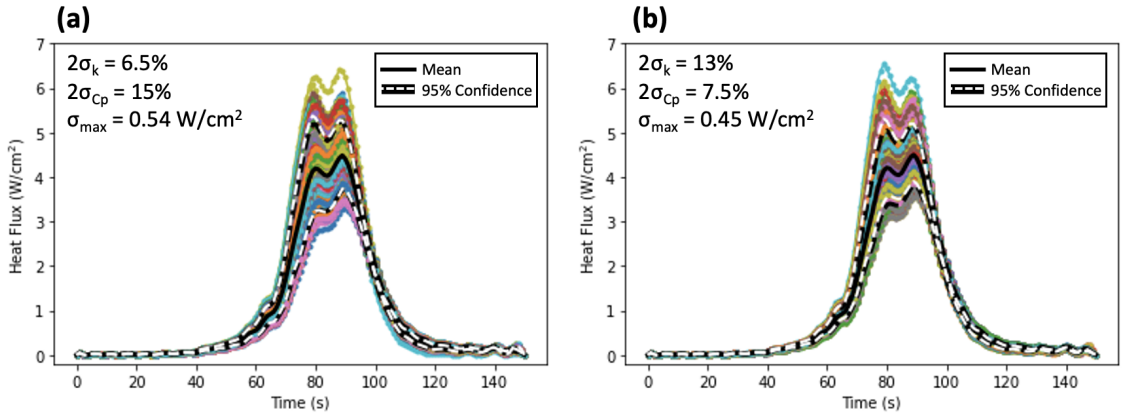

Fig. 7. Plots of all 2000 Monte Carlo runs with the nominal 2σ uncertainties for k and C_p, as well as when the uncertainties are halved for each of those parameters.

of Spacecraft and Rockets, vol. 36, no. 3, pp. 475–483, 1999.

[3] H.S. Alpert, D.A. Saunders, M. Mahzari, J.D. Monk, and T.R. White, "Inverse Estimation of Mars 2020 Entry Aeroheating Environments Using MEDLI2 Flight Data," in *AIAA SciTech Forum*. AIAA, 2022.

[4] J. Monk, J. Feldman, J. Santos, T. White, D. Prabhu, and H. Alpert, "MEDLI2 Material Response Model Development and Validation," in *AIAA SciTech Forum*, 2022.

[5] A. Saltelli, P. Annoni, I. Azzini, F. Campolongo, M. Ratto, and S. Tarantola, "Variance Based Sensitivity Analysis of Model Output. Design and Estimator for the Total Sensitivity Index," *Computer Physics Communications*, vol. 181, pp. 259–170, 2010.

[6] A. Moradi, M. Tootkaboni, and K. Pennell, "A Variance Decomposition Approach to Uncertainty Quantification and Sensitivity Analysis of the Johnson and Ettinger Model," *Journal of the Air and Waste Management Association*, vol. 65, no. 2, pp. 154–164, 2015.