Journal of
Nuclear Engineering

MDPI

Article

Artificial Neural Networks as Surrogate Models for Uncertainty Quantification and Data Assimilation in 2-D/3-D Fuel Performance Studies

Carlo Fiorina [1,*], Alessandro Scolaro [1], Daniel Siefman [1], Mathieu Hursin [1,2] and Andreas Pautz [1,2]

[1] Laboratory for Reactor Physics and System Behaviour, EPFL—École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland; alessandro.scolaro@epfl.ch (A.S.); daniel.siefman@epfl.ch (D.S.); mathieu.hursin@epfl.ch (M.H.); andreas.pautz@epfl.ch (A.P.)

[2] Nuclear Energy and Safety Division, Paul Scherrer Institut, 5232 Villigen, Switzerland

* Correspondence: carlo.fiorina@epfl.ch

Received: 22 September 2020; Accepted: 4 November 2020; Published: 10 November 2020

Abstract: This paper preliminarily investigates the use of data-driven surrogates for fuel performance codes. The objective is to develop fast-running models that can be used in the frame of uncertainty quantification and data assimilation studies. In particular, data assimilation techniques based on Monte Carlo sampling often require running several thousand, or tens of thousands of calculations. In these cases, the computational requirements can quickly become prohibitive, notably for 2-D and 3-D codes. The paper analyses the capability of artificial neural networks to model the steady-state thermal-mechanics of the nuclear fuel, assuming given released fission gases, swelling, densification and creep. An optimized and trained neural network is then employed on a data assimilation case based on the end of the first ramp of the IFPE Instrumented Fuel Assemblies 432.

Keywords: 3-D fuel performance; uncertainty quantification; data assimilation; surrogate models; artificial neural networks

1. Introduction

The EPFL and the Paul Scherrer Institute have been developing, in recent years, a multi-dimensional fuel performance tool named OFFBEAT [1–3]. The code is based on OpenFOAM [4] and allows for traditional 1.5-D simulations, as well as for more advanced 2-D and 3-D treatments. In addition, OFFBEAT is being developed for straightforward use in uncertainty quantification (UQ) and data assimilation (DA). This is achieved by exposing to user modification a large number of parameters, and by allowing to directly perturb the global effect of complex phenomena like creep, swelling, fission gas release, etc.

However, UQ and DA techniques for non-linear problems and non-Gaussian distributions are often based on Monte Carlo sampling. The relatively large number of parameters involved in fuel performance studies then requires several thousand or tens of thousands of samples. As a result, computational requirements can quickly become extremely large, notably for 2-D and 3-D simulations. A typical solution to this problem is the development of so-called surrogate models, i.e., minimal fast-running models that can approximate the solution of the full-order model (FOM) in the range of parameters of interest.

Several techniques exist for creating surrogate models, including physical or dimensional simplifications, projection-based reduced order models, and data-driven models. The first two options are intrusive, requiring access to the source code, and time consuming to implement. This paper focuses on data-driven models, which are instead easy to set up, very general, and can benefit from

the quick-paced developments in the field of machine learning. As a drawback, they require data for training, which implies having to run the FOM for a possibly large set of cases. A data-driven model can then be considered as effective if: (1) the number of necessary training samples is much smaller compared to the number of FOM simulations required in the target application (UQ and DA in our case); and (2) the inaccuracies in the model do not significantly affect the results of the target application.

The selection of the surrogate model strongly depends on the FOM and on the phenomenon under investigation. In the case of fuel performance, one can perform UQ and DA analyses based on:

1. The behavior of the fuel during irradiation, which allows to include in the analysis important time-dependent phenomena like densification, swelling, creep and fission gas release; or
2. The behavior of the fuel at a given time, which limits the analysis to beginning-of-life (BOL) situations, or requires assumptions on the above mentioned time-dependent phenomena.

In this preliminary paper, only the second option is investigated. This allows UQ/DA analysis of important BOL parameters like relocation, fabrication tolerances, fuel expansions, etc. (see Ref. [5]). In addition, one can calculate the time dependent parameters, estimate their uncertainty, and verify whether or not the prediction matches experimental data. This can provide information on the accuracy of the models for fission gas release, creep, etc.

A steady-state model takes the model parameters as input and provides a quantity of interest (QoI) as output, acting in essence as an interpolator. For this reason, Artificial Neural Networks (ANN) have been selected as a promising surrogate model for the purpose of this work. ANN are in fact known to perform remarkably well as interpolators. They are very easy to set-up and train (if not too deep), and they can accommodate a large degree of over-fitting. This can come in handy when building surrogate models for different applications (e.g., different UQ studies) without having to finely tune the degree of complexity of the surrogate model. On the other hand, ANNs are not necessarily the best surrogate model for transient scenarios, where one may consider for instance the use of Gaussian Processes.

2. Test Case: The IFPE IFA-432 Assembly

The International Fuel Performance Experiments (IFPE) database collects the data of several experiments, focusing mainly on thermal reactor performance. The rods that have been chosen for the present work are part of the Instrumented Fuel Assemblies 432 (IFA-432) experiment, which was designed to study a BWR-type fuel, up to a burnup of ~40 MWd/kgHM. Each rod was equipped with two thermo-couples, inserted from the top and bottom of the pellet stack. Out of a total of 6 rods, rods 1 and 2 have been chosen for the present study. For these rods, only the measurements of the lower thermo-couples are available. Several experiments were performed at various burnup levels, during which the Fuel Centerline Temperature (FCT) was recorded at the thermo-couple location. As mentioned, the current paper focuses on BOL data. For this purpose, the temperature measurements obtained at the end of the first power ramp (burnup ~0.02 MWd/kgHM) have been employed. An axisymmetric 2-D model has been built using OFFBEAT for each of the two rods, including of course the thermo-couple holes. Figure 1 shows, as an example, the model created for the IFA-432 rod 1, together with a detail of the mesh. More modelling details can be found in Ref. [1].

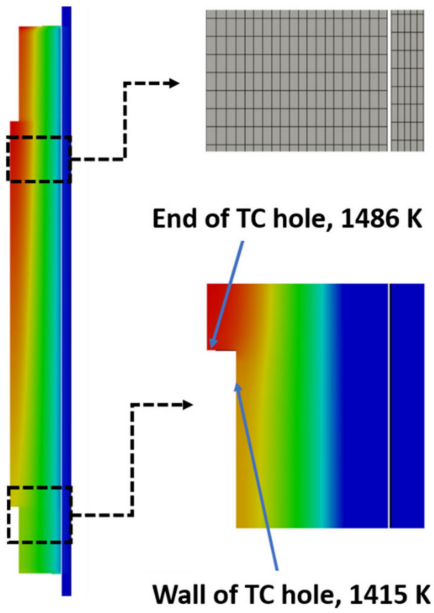

Figure 1. Axisymmetric (2-D r-z) model for the IFA-432 rod 1 [1].

3. Development and Training of the ANN Surrogate

An ANN [6] is normally configured as a sequence of "neurons". Each neuron receives as input the output of a certain number of other neurons (and sometimes from its previous state), it calculates a linear combination of these inputs (using weights), it adds a constant (bias), it applies an activation function, and it outputs the results to a certain number of other neurons. Among the various types of ANNs, multi-layer perceptrons (MLPs), also known as vanilla ANNs, have been selected for this work. They are among the oldest type of ANNs and they are characterized by well-defined layers of neurons, with each neuron receiving as input the outputs of all the neurons in the previous layer, and sending its output to all the neurons in the next layer (Figure 2). Layers that are neither inputs nor outputs are called hidden layers. When multiple hidden layers are employed, the ANN is called "deep". To train an MLP, one needs a set of labelled input data, i.e., a set of input data, each one accompanied by the expected output result. With such a set available, one starts with a random selection of weights and biases, compares the output of the ANN with the expected output, calculate a cost function (typically an L2 norm), back-propagate the error, and perform a gradient descent on weights and biases. In our case, training data will be obtained using OFFBEAT.

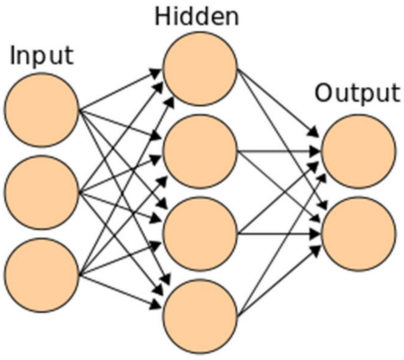

Figure 2. Schematic representation of an ANN.

In this work, the FTC is selected as QoI. Since in this paper the focus is on BOL, this will lead to a single output for the ANN. Sixteen input parameters have been selected, as shown in Table 1. In cases where the parameter is a scalar, its value is directly inputted to the ANN. For cases where the parameter of interest is already the output of a model (viz., creep), OFFBEAT allows applying a perturbation directly to the results of equations and correlations (viz., creep strain). As a matter of fact, all input parameters are normalised to a 0–1 range before inputting them to the ANN. This avoids the neurons of the first layer having to deal with very large differences in the input values. Two different possibilities are reported in Table 1 for standard deviations. One is used for DA and it is based on realistic uncertainties of the data [7]. The second is used for training the ANN and it includes much wider ranges, which in turn allows to train the model for a wider range of applications. Parameters have been chosen in Table 1 to avoid situations with a closed gap, which is non-realistic at BOL.

Table 1. Parameters used to train the ANN and for the DA process.

	Mean Value	Relative Standard Deviation or Half Range [%]		Distribution Type
		ANN Training	DA	
Volume power [MW/m^3] [1]	300	33	-	uniform
Fuel conductivity (perturbation) [-]	1	10	10	normal
Fuel emissivity (perturbation) [-]	1	3	3	normal
Fuel thermal expansion (perturbation) [-]	1	7.5	7.5	normal
Fuel relocation (perturbation) [-]	1	100	5	uniform
Cladding conductivity (perturbation) [-]	1	2.5	2.5	uniform
Cladding emissivity (perturbation) [-]	1	5	5	uniform
Cladding thermal expansion (perturbation) [-]	1	15	15	normal
Fuel roughness [-]	2.2×10^{-6}	12.5	12.5	normal
Cladding roughness [-]	5×10^{-7}	15	15	normal
Diffusive gap conductance (perturbation) [-]	1	10	10	uniform
Contact gap conductance (perturbation) [-]	1	50	50	uniform
Inner fuel radius [mm]	0.875	20	-	uniform
Outer fuel radius [mm]	5.264	20	-	uniform
Gap width [mm]	0.1905	60	-	uniform
Clad width [mm]	0.94	20	-	uniform

[1] Although OFFBEAT allows for non-uniform radial distributions, a flat power profile is considered for simplicity in this investigation. This is not expected to significantly impact the findings of the paper since the profile would be the same at all power levels, thus modifying the proportionality between FCT and power but without introducing significant non-linear effects.

As mentioned, data are necessary to train the ANN. For this purpose, 10,000 combinations of input parameters have been randomly sampled according to the probability density functions reported in Table 1 (ANN training). It is worth noting that running 10,000 simulations with the simple model described in Section 2 requires 20 cores for few days. Calculation time can then become much higher for more refined models or 3-D cases. 70% of the data have been used for training, 15% for testing (i.e., for evaluating the performances of different ANN architectures), and 15% for validation (i.e., for evaluating the performance of the ANN that is selected for application). The output (FCT) distribution of the three data sets is reported in Figure 3. A non-symmetric distribution can be observed. Consequently, the model will in principle tend to be less accurate at relatively high temperatures, where only a few samples are available for training the ANN.

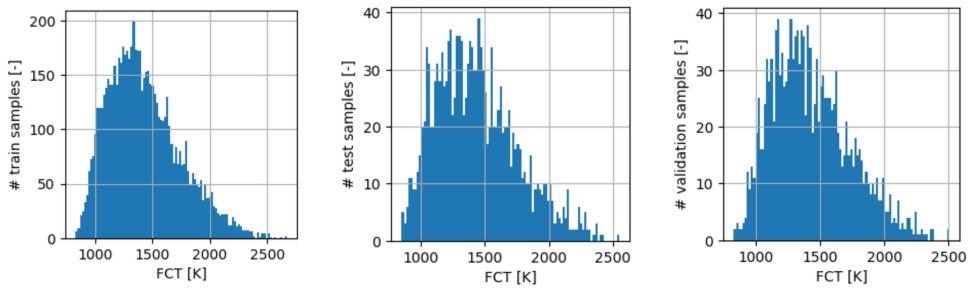

Figure 3. FCT distribution of training, test and validation data.

A few different ANN topologies have been tested and the results are reported in Table 2. As mentioned, all the ANNs feature 16 inputs and one output. Since the ANN is here employed as a regression model, and since only positive values are expected, a so-called ReLU (rectified linear unit) activation function [6] is employed for the output neuron. For the other layers, two possible activation functions are considered, namely: the traditional non-linear sigmoid function; and the exponential linear unit (ELU), which is a quasi-linear function and closely resembles the ReLU (but without discontinuities on the first derivative) [6]. Average absolute error (AAE) was used as cost function. Results are reported in Table 2 in terms of:

1. AAE on the training set (bias), which gives an indication about the capability of the model to fit the data. It should be noted in this sense that a 5% uncertainty on the FCT is common in fuel performance experiments, which corresponds to several tens of degrees.
2. AAE on the test set, which gives an indication about the capability of the model to generalize its results. This represents the main metrics for evaluating the performances of the ANN.
3. Maximum error on the test set, which can complement the average absolute error on the test set with an estimate about the maximum expected error.
4. Degrees of Freedom (DoF), which measures the complexity of the ANN model.

Table 2. Performances of different ANN topologies.

ANN Topology	AAE on Training Set [K]	AAE on Test Set [K]	Max Error on Test Set [K]
No hidden layers—17 DoF	40.8	41.4	241.9
1 linear hidden layer (16 neurons)—289 DoF	2.62	2.69	14.6
1 linear hidden layer (100 neurons)—1801 DoF	2.66	2.67	14.7
1 non-linear hidden layer (16 neurons)—289 DoF	2.42	2.44	14.9
2 linear hidden layers (16 neurons each)—561 DoF	1.91	1.92	10.9
1 non-linear, 1 linear hidden layers (16 neurons)—561 DoF	1.42	1.43	8.22
1 linear, 1 non-linear hidden layers (16 neurons)—561 DoF	1.42	1.47	10.05
2 non-linear hidden layers (16 neurons each)—561 DoF	2.59	2.68	13.95

Please notice that Table 2 reports the results of a single fit. However, a limited variability has been observed between different fits resulting from different initial weights and biases.

The first tested ANN has no hidden layers, which results in a simple linear interpolation of the input parameters. As expected from the complexity and non-linearity of a fuel behavior model, the resulting errors are large on both training and test sets. The use of one hidden layer with 16 neurons significantly improves the results, with a slightly better behavior when using a non-linear activation function. An increased number of neurons in the hidden layer has essentially no impact, showing that combining multiple times the input parameters does not add to the accuracy of the model. A significant improvement in performance is instead obtained when deepening the ANN by one layer, notably when one of the two hidden layers is non-linear and the other linear. The need for a non-linear layer is once again consistent with the non-linearity of the FOM. The performance loss that we observe with two hidden non-linear layers is likely due to the tendency of the sigmoid activation functions to saturate and become inactive [6].

Following the results of Table 2, an ANN with a first linear (ELU) and a second non-linear (sigmoid) hidden layers has been selected for the DA study. Its performance on the validation set show an average error of 1.44 K and a maximum error of 8.0 K, which confirms the good generalization capability of the ANN. However, obtaining 10,000 FOM solutions might be prohibitive in many cases. To better understand the needs of the selected ANN in terms of training data, two subsets of the available labelled data have been created containing 100 and 1000 samples, respectively. The results are reported in Table 3. The degradation of performances with a lower number of data is clearly visible, but a data set with 1000 samples appears to be enough for guaranteeing a reasonable accuracy. It is also worth noting that with 100 and 1000 samples, overfitting becomes possible, with a significant difference between the average absolute error on training and test cases. In other words, the ANN may fit well the data, but it generalizes poorly. Several techniques exist to limit overfitting such as L1/L2 regularization, early stopping, drop-out, etc. Their investigation is however beyond the scope of this work.

Table 3. Performances of the selected ANN for different training data sizes.

Data Size	Average Absolute Error on Training Set [K]	Average Absolute Error on Test Set [K]	Max Error on Test Set [K]
100	7.13	30.3	73.7
1000	2.22	3.14	13.3
10,000	1.42	1.43	8.22

A deeper understanding of the performances of the ANN can be obtained by plotting the distribution of errors. Figure 4 reports the errors on FCTs obtained by training the selected ANN on 1000 data and by applying it to the prediction of the whole 10,000 data set. First, the distribution of errors is roughly symmetrical around zero, which is a desired characteristic to avoid biasing the result of UQ and DA studies. In addition, slightly larger errors (and higher variance) are observed at high temperatures, which is consistent with the lower number of training data in that region. Finally, few large errors are observed at the lower boundary of the temperature range. These errors are associated to four points in the 10,000 data set that are outside of the range included in the 1000 data set (and used for training). This underlines the importance of the training set to cover the full range of inputs/outputs, and for the model to limit overfitting.

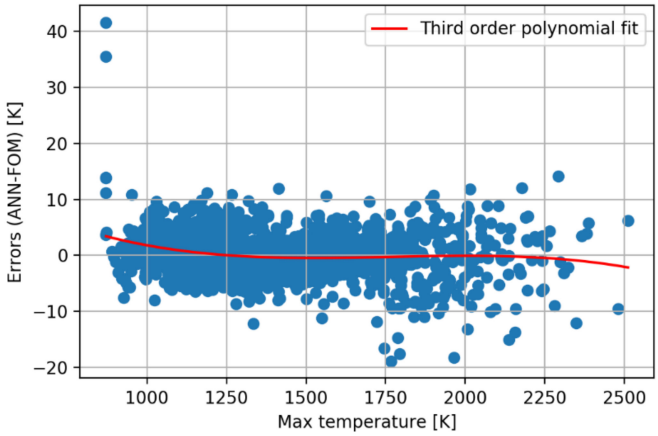

Figure 4. FCTs prediction errors on the 10,000 data set using the selected ANN trained over the 1000 data set.

4. Testing on a DA Study

Rods 1 and 2 of the IFA-432 assembly have been selected for testing the performance of the selected ANN on a practical DA study. The objective is to showcase the effectiveness of the ANN, while the reader should not focus on the specific results of the DA, which might be affected by approximations such as the flat power profile and the lack of uncertainties on the power level. The reader may refer to Ref. [5] for a more detailed DA work on the same test case.

In view of the results of Table 3 and of the computational burden of obtaining 10,000 samples in many practical cases, the ANN trained on 1000 cases has been selected. In particular, the objective is to assimilate the measured FCT values only at the end of the first ramp. This is to avoid assimilating inconsistent data that can be observed along the ramp, where different FCTs are predicted for the same power [5]. The two rods feature the same mean parameters reported in Table 1 the only three differences are in the fuel radius, gap width, and volume power. These values for the two rods are reported in Table 4, together with the corresponding FCTs.

Table 4. Parameters characterizing rods 1 and 2 of the IFA-432 assembly.

	Fuel Radius [mm]	**Gap Width [mm]**	**Volume Power [MW/m^3]**	**Experimental FCT [K]**
Rod 1	5.3405	0.114	358	1388.4
Rod 2	5.264	0.1905	368	1569.6

For the DA, a Bayesian Monte Carlo (BMC) technique has been adopted. The BMC technique consists of randomly sampling the input parameters based on their prior distribution, comparing the corresponding outputs (FCTs in our case) to the experimental ones, and propagating this difference back to the inputs in order to obtain the posterior (adjusted) distributions. More details about this technique can be found in Ref. [5,8,9]. The prior distributions are reported in Table 1. Power and geometrical parameters are assumed as known, while the objective of the DA is to adjust the distribution of physical properties and phenomena. For simplicity, no correlation is considered between the experimental FCTs.

Table 5 shows the result of the DA on the parameters that are mainly affected by the adjustment, and on the corresponding FCTs. The DA has been performed simultaneously on the two pins. Results have been obtained using both OFFBEAT and the trained ANN, showing a very good agreement between the two.

Table 5. Results of the DA process [1].

	Before Assimilation with OFFBEAT		After Assimilation with OFFBEAT		After Assimilation with ANN	
	Mean	**Standard Deviation [%]**	**Mean**	**Standard Deviation [%]**	**Mean**	**Standard Deviation [%]**
Gap diff. conduct.	1.0	10.0	1.00628	5.69	1.00608	5.68
Clad. therm. exp.	1.0	15.0	0.99616	13.17	0.99586	13.17
Fuel conductivity	1.0	10.0	1.05015	5.47	1.05234	5.50
Fuel relocation	1.0	5.0	1.00102	2.87	1.00124	2.88
Fuel therm. exp.	1.0	7.5	1.00531	6.55	1.00503	6.55
FCT Rod 1 [K]	1423.4	5.0	1384.2	2.79	1385.1	2.75
FCT Rod 2 [K]	1644.4	5.0	1599.9	2.76	1601.7	2.74

[1] One may notice some non-physical results, such as a predicted 5% increment in the fuel conductivity. This is the result of few simplifying assumptions in the DA process, and in particular of the lack of uncertainty on fuel power. A detailed DA study is beyond the scope of this work, that mainly aims at testing the effectiveness of ANNs as surrogate models for fuel behavior. The reader may refer to [5] for a more accurate DA work on the same test case.

Figure 5 shows the convergence of the posteriors with respect to the number of Monte Carlo samples employed for the BMC assimilation. One can observe that ~10,000 samples are necessary for an acceptable convergence, corresponding to 20,000 simulations (two rods). Since the time for training and running the ANN is negligible compared to the time necessary for running the FOM, one can conclude that in this case the use of an ANN trained on 1000 samples can lead to a 40 times reduction on the overall computing time necessary for DA. As a matter of fact, one could use the ANN for any rod whose parameters fall within its training range, potentially providing very large saving in terms of computing time.

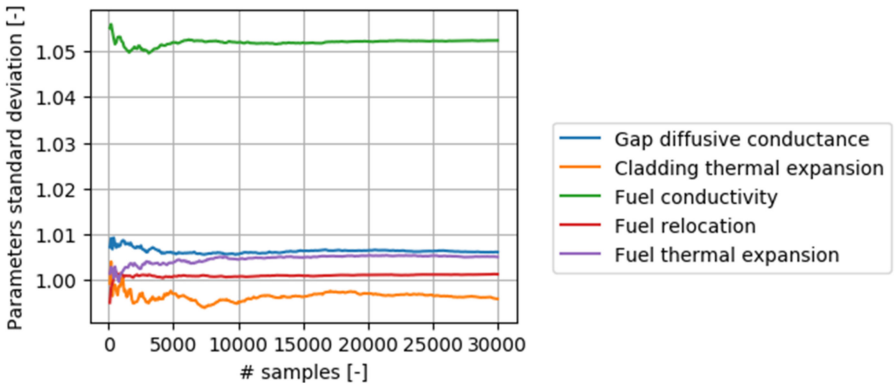

Figure 5. Convergence of posteriors with respect to the number of Monte Carlo samples.

5. Conclusions

In this paper the possibility has been investigated of using ANNs as surrogate models for UQ and DA in fuel performance studies. It has been shown that ANNs have the capability to accurately reproduce the results of the OFFBEAT fuel performance code. To achieve this, one can make use of simple "rectangular" ANNs with one or two hidden layers. Use of a non-linear activation function in at least one of the layers is recommended. As few as 1000 labeled data have been proven to be enough for practical applications, although 10,000 samples provide significantly better results and a very limited risk of overfitting. With 1000 labeled data and a two-layer ANN, one can reduce the computational cost of DA studies by several tens of times, while negligibly affecting the result of the DA process. This reduction of computational costs takes into account the computational time that is needed for training the ANN. This is an important factor, since re-training the ANN will tend to be a

necessary step in fuel behavior studies, where codes and models undergo a continuous development with frequent updates.

The focus of the work was on steady-state cases with a given densification, creep, fission gas, etc. This implies that the results of this work can be applied to DA studies based on: low-burnup experimental data; or, higher-burnup experimental data where history-dependent parameters are available (e.g., by post-irradiation examinations) or can be reasonably guessed/calculated, where reasonably means with an uncertainty that does not jeopardize the whole DA process. In this sense, it is worth mentioning that the performances of the ANNs are expected to deteriorate with an increasing number of input parameters, which will be the case for DA studies on higher-burnup experimental data. This possible deterioration will be the subject of future studies of the authors.

Author Contributions: Conceptualization: C.F., A.S., D.S., M.H.; Methodology: C.F., A.S., D.S., M.H.; Software: C.F., A.S., D.S.; Validation: C.F., A.S., D.S.; Formal analysis: C.F., A.S., D.S., M.H.; Investigation: C.F., D.S.; Data curation: C.F.; Writing—original draft preparation: C.F.; Writing—review and editing: C.F., A.S., D.S., M.H., and A.P.; Visualization: C.F., D.S., M.H.; Supervision: C.F.; Project administration: C.F., M.H., and A.P.; Funding acquisition: A.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Scolaro, A.; Clifford, I.; Fiorina, C.; Pautz, A. The OFFBEAT multi-dimensional fuel behavior solver. *Nucl. Eng. Des.* **2020**, *358*, 110416. [CrossRef]
2. Scolaro, Y.; Robert, I.; Clifford, C.; Fiorina, A.P. Coupling Methodology for the Multidimensional Fuel Performance Code Offbeat and The Monte Carlo Neutron Transport Code Serpent. In Proceedings of the Global/Top Fuel 2019, Seattle, WA, USA, 22–26 September 2019.
3. Scolaro, A.; Clifford, I.; Fiorina, C.; Pautz, A. Multi-dimensional Creep Analysis Using the Novel OFFBEAT. In Proceedings of the IAEA Technical Meeting on Progress on Pellet Cladding Interaction and Stress Corrosion Cracking, Aix-en-Provence, France, 18–11 October 2019.
4. OpenFOAM. Available online: www.openfoam.com (accessed on 1 September 2020).
5. Fiorina, C.; Scolaro, A.; Siefman, D.; Hursin, M.; Pautz, A. Bayesian data assimilation for the 3-D OFFBEAT fuel performance code. In Proceedings of the BEPU 2021 Conference, Lucca, Italy, 30 May–5 June 2021.
6. Géron, A. *Hands-on Machine Learning with Scikit-Learn & TensorFlow*; O'Reilly: Sebastopol, CA, USA, 2019.
7. Che, Y.; Wu, X.; Li, W.; Shirvan, K.; Pastore, G.; Hales, J. Sensitivity and Uncertainty Analysis of Fuel Performance Assessment of Chromia-Doped Fuel during Large-Break Loca. In Proceedings of the Top Fuel Conference, Prague, Czech Republic, 30 September–4 October 2018.
8. Salvatores, M.; Palmiotti, G.; Aliberti, G.; Archier, P.; Jean, C.D.S.; Dupont, E.; Herman, M.; Ishikawa, M.; Ivanova, T.; Ivanov, E.; et al. Methods and Issues for the Combined Use of Integral Experiments and Covariance Data: Results of a NEA International Collaborative Study. *Nucl. Data Sheets* **2014**, *118*, 38–71. [CrossRef]
9. Rochman, D.; Bauge, E.; Vasiliev, A.; Ferroukhi, H.; Pelloni, S.; Koning, A.J.; Sublet, J.C. Monte Carlo nuclear data adjustment via integral information. *Eur. Phys. J. Plus* **2018**, *133*, 537. [CrossRef]