



CHALMERS
UNIVERSITY OF TECHNOLOGY


On the use of neutron flux gradient with ANNs for the detection of diverted spent nuclear fuel

Downloaded from: <https://research.chalmers.se>, 2026-04-05 11:37 UTC

Citation for the original published paper (version of record):

al-Dbissi, M., Pazsit, I., Rossa, R. et al (2024). On the use of neutron flux gradient with ANNs for the detection of diverted spent nuclear fuel. *Annals of Nuclear Energy*, 204.
<http://dx.doi.org/10.1016/j.anucene.2024.110536>

N.B. When citing this work, cite the original published paper.



On the use of neutron flux gradient with ANNs for the detection of diverted spent nuclear fuel

Moad Al-dbissi ^{a,b,*}, Imre Pázsit ^a, Riccardo Rossa ^b, Alessandro Borella ^b, Paolo Vinai ^a

^a Division of Subatomic, High Energy and Plasma Physics Chalmers University of Technology, 412 96, Göteborg, Sweden

^b Belgian Nuclear Research Centre, SCK CEN, 2400 Mol, Belgium

ARTICLE INFO

Keywords:

Neutron flux gradient
Artificial neural networks
Spent nuclear fuel
Nuclear safeguards
Monte Carlo

ABSTRACT

One of the main tasks in nuclear safeguards is regular inspections of Spent Nuclear Fuel (SNF) assemblies to detect possible diversions of special nuclear material such as ²³⁵U and ²³⁹Pu. In these inspections, characteristic signatures of SNF such as emissions of neutrons and gamma rays from the radioactive decay, are measured and their consistency with the declared assemblies is verified to ensure that no fuel pins have been removed. Research in this field is focused on both the development of detection equipment and methods for the analysis of the acquired measurement data. In this paper, the use of the neutron flux gradient, which is not considered in regular SNF verification, is investigated in combination with the scalar neutron flux as input to artificial neural network models for the quantification of fuel pins in SNF assemblies. The training and testing of these ANN models rely on a synthetic dataset that is generated from Monte Carlo simulations of a typical intact pressurized water reactor assembly and with different patterns of fuel pins replaced by dummy pins. The dataset consists of unique scenarios so that the ANN can be assessed over “unknown” cases that are not part of the learning phase. Results show that the neutron flux gradient is advantageous for a more accurate reconstruction of diversions within SNF assemblies.

1. Introduction

During the operation of a nuclear power plant, the fuel assemblies in the core are depleted of their fissile content and accumulate fission products. Therefore, their effectiveness to sustain fission reactions in the core decreases over time until they must be replaced by fresh ones. The Spent Nuclear Fuel (SNF) assemblies discharged from the core are stored in a water pool where their decay heat is dissipated and their radiation emissions are shielded. This is a temporary condition, before the transfer to a final repository or a reprocessing facility. Since the assemblies in the pool have a residual fissile content of ²³⁵U and ²³⁹Pu, regular safeguards inspections are conducted by the International Atomic Energy Agency (IAEA) to detect any illicit removal of SNF and minimize proliferation risks.

Safeguards inspections of SNF use Non-Destructive Assay (NDA) techniques to verify that fuel pins are not diverted from the assemblies. The majority of NDA techniques are based on the passive measurement of characteristic signatures of SNF such as the emissions of neutrons, gamma rays, or Cherenkov light (Branger et al., 2020; Rinard and Bosler, 1988; Mayorov et al., 2017).

Previous work investigated the potential of Artificial Neural Network (ANN) models for the inversion problem of reconstructing the

configuration of the fuel pins in a SNF assembly and characterizing possible diversions from the scalar neutron flux and the gamma rays detected inside the assembly, see Al-dbissi et al. (2023). In the current paper, a novel aspect is explored, i.e., the use of the neutron flux gradient as further input feature to these types of algorithms. The neutron flux gradient, which provides the direction in which the neutron flux increases and the rate of such an increase, has richer information than only the neutron flux, and thus is expected to enhance the predictions. This was already demonstrated, via simulations, for the similar task of localizing neutron sources, see Pázsit (1997), Lindén et al. (1999), Avdic et al. (2001). In addition, progress with miniaturized neutron detectors that consist of neutron scintillators connected to light-guiding fibers, have enabled the possibility of measuring the neutron flux gradient inside the fuel assemblies, see Watanabe et al. (2020), Vitullo et al. (2020), Aldbissi et al. (2022a).

For the purpose of the study, ANN models are developed to process measurements of only the neutron flux or the neutron flux combined with its gradient, and their performances are compared. The work relies on a synthetic dataset generated from Monte-Carlo simulations of a typical intact PWR fuel assembly and assemblies with different patterns of fuel pins replaced by dummy pins.

* Corresponding author at: Division of Subatomic, High Energy and Plasma Physics Chalmers University of Technology, 412 96, Göteborg, Sweden.
E-mail address: moad.al-dbissi@chalmers.se (M. Al-dbissi).




Fig. 1. Monte Carlo model of a PWR fuel assembly; inside each guide tube, the neutron flux is estimated at four locations and used to estimate the gradient.

The paper is structured as follows. The overall methodology and a description of the training dataset are introduced in Section 2. Results and example cases are discussed in Section 3. Conclusions are drawn in Section 4.

2. Methodology

The general strategy for the verification of spent fuel is to determine whether the measurements of the characteristic signatures of an assembly with unknown configuration are consistent with the declared data provided by the power utility. If a discrepancy is detected, an unfolding procedure can be applied to estimate the diversion that might be consistent with the measurement results.

For the unfolding task, ANN models that can process the scalar neutron flux and its gradient measured inside the assembly, are studied. The ANN is trained to “learn” the relationship between sets of measurements and the relative configurations, which might be intact or with patterns of fuel pins replaced by dummy pins. Both the training and the testing of these models make use of synthetic data generated from Monte Carlo simulations.

2.1. The gradient of the neutron flux

In the SNF assembly under inspection, each fuel pin is assumed to be either intact or fully replaced by a dummy pin (which is made of appropriate surrogate materials, e.g., stainless steel), so only the radial position of the pins in the horizontal plane is of interest and the problem is two-dimensional. The neutron flux gradient can be evaluated at some fixed axial elevation, from the neutron flux measured in four points inside each of the available empty guide tubes, as shown in Fig. 1 (Aldbissi et al., 2022b). The two Cartesian components (X and Y) of the gradient can be derived from each pair of points along the same diagonal and can be used to determine eventually the magnitude (absolute value) and the direction of the gradient vector. The feasibility of measuring the neutron flux gradient using a miniaturized detector with four neutron scintillators that can fit into a guide tube of a PWR SNF assembly, was evaluated via Monte-Carlo simulations. The results of this study together with the description of the detector design are discussed in Aldbissi et al. (2022a).

2.2. Training dataset

The dataset for the training and testing of the ANN algorithm relies on Monte Carlo simulations of a standard 17×17 PWR fuel assembly, see Fig. 1. The assembly consists of 264 fuel pins (0.475 cm in radius), each with a helium gap (0.008 cm in thickness) and Zircaloy cladding (0.064 cm in thickness), 24 empty guide tubes (0.602 cm in radius), and a central instrumentation channel (0.602 cm in radius). Although assemblies stored in a water pool may have significant differences in

their history, only one case is considered in the current study where the fuel has an initial enrichment of 3.5 w%, a final burn-up of 40 MWd/kgU, and a cooling time of 5 years. In addition, an ideal case scenario is taken in which all guide tubes and the instrumentation channel are accessible locations where detectors can be placed.

A model of the fuel assembly is developed using the Monte Carlo code Serpent (Leppänen, 2015) and simulations were performed in two-steps. The first step is a burn-up simulation of the declared fresh fuel assembly, assuming a spatially uniform depletion. The irradiation cycle consists of four burn-up steps until a final burn-up value of 40 MWd/kgU is reached. The total length of the irradiation cycle is 1000 days. The irradiation is then followed by a decay cycle equivalent to a cooling time of 5 years in the spent nuclear fuel pool. For the burn-up simulation, the criticality source mode in Serpent is selected, and 500 active generations and 10000 neutron histories per generation are used. The second step is a fixed-source simulation, where the fuel composition obtained from the burn-up simulation is distributed consistently with the diversion patterns of interest and the thermal neutron flux (which typical neutron detectors are mainly sensitive to) and its gradient are estimated in the guide tubes of the assembly. The fixed-source simulation is performed using 5×10^9 neutron histories to balance accuracy and computational time. The statistical error of the Monte Carlo calculation of the thermal neutron flux estimated at the measurement points inside the guide tubes of the assembly is 0.1% on average.

The dataset contains the case of an intact fuel assembly (without defects) and 109 cases with diversion patterns, which are symmetrical or asymmetrical and have a minimum of 4 up to a maximum of 180 fuel pins replaced by stainless steel pins, see examples in Fig. 2. Since the arrangement of a 17×17 PWR fuel assembly has 25 guide tubes, the calculated system responses for each configuration are 75, i.e., 25 values of thermal neutron flux, 25 values of the magnitude of the gradient (absolute value), and 25 values of the angle of the gradient vector (direction). A second dataset is also created, in which the two (X and Y) components of the gradient vector are used instead of the magnitude and the direction.

2.3. Artificial neural networks

Artificial Neural Networks (ANNs) are an advanced machine learning algorithm. They can model non-linear relationships and thus learn and identify complex patterns between inputs and outputs. In the current application the network is tasked with the identification of the locations where the fuel pins have been replaced with dummy pins inside the assembly (if any). Accordingly, the network processes the simulated responses of the neutron flux and/or the gradient for a fuel assembly as input and gives as output a probability of being replaced to each of the fuel pins.

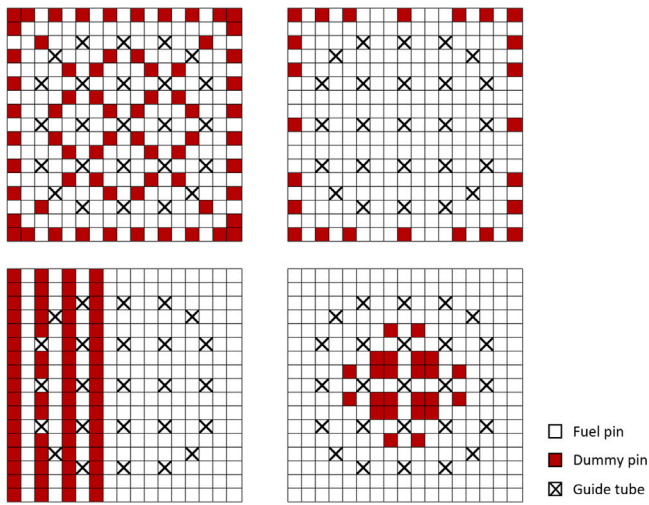


Fig. 2. Examples of SNF diversion patterns included in the dataset.

The ANN model is built using the Tensorflow (Abadi et al., 2015) and the Keras (Chollet et al., 2015) open-source software libraries. Its structure consists of an input layer, a hidden layer and an output layer. The neurons that belong to the input and hidden layers are activated with the Rectified Linear Unit (ReLU) function, which allows for back-propagation with an efficient convergence rate. The neurons that belong to the output layer are activated with the Sigmoid function which is a typical choice for outputs that are non-mutually exclusive, where each pin is treated independently and can be either present or replaced. The model provides as an output, a probability between 0 and 1 for each fuel pin in the assembly. If the probability of a fuel pin is between 0.5 and 1, the fuel pin is labeled as missing (1). If the probability is less than 0.5, the fuel pin is labeled as present (0). The task can then be considered a multi-label binary classification.

The performance of the network is evaluated with the Binary Cross-Entropy loss function. The loss function provides a measure of how closely the distribution of the predictions matches the distribution of the target variables in the training data. In the iterations of the training, the weights and the biases associated with the neurons of the network are optimized by minimizing the value of the loss function via the Adaptive Moment Estimation (ADAM) optimizer.

The number of neurons in the input layer is fixed to the number of input features used for training, and the number of neurons in the output layer is equal to the number of outputs, i.e., one for each of the 264 fuel pins in the assembly. A grid search optimization is performed to determine the number of neurons for the hidden layer and the number of epochs and the batch size in the training process.

The training and testing is performed via a 10-fold cross-validation method. Accordingly, the whole dataset is shuffled and divided into 10 random batches. Nine of these batches are used to train the ANN, while the remaining one is used for the testing. The procedure is repeated 10 times so that each of the 10 batches serves as testing sample one time, and then the results are aggregated. In the current study, the dataset has a small size and the cross-validation method has the advantage of allowing the use of all the cases to test the trained model. In addition, the cross-validation is repeated 5 times with a different initial shuffling of the dataset to reduce the bias in the assessment of the performance.

As mentioned above, the dataset contains 110 fuel assemblies, each of them with a unique configuration. Therefore, the model is always tested over scenarios that are not seen in the training phase and thus its ability to generalize its predictions with respect to unknown data can be investigated. This aspect is relevant because a training dataset with all the possible diversion patterns would require unfeasible computational resources.

Table 1
General form of the confusion matrix.

ANN Model		Real Label	
		Intact (0)	Missing (1)
Predicted Label	Intact (0)	True Neg.	False Neg.
	Missing (1)	False Pos.	True Pos.

2.4. Performance metrics

After the cross-validation process is completed, the model is scored with respect to the number of fuel pins that have been identified correctly in all the fuel assemblies available from the dataset. The predictions of the model fall into 4 categories, see Table 1. The 'True Negatives' are the correctly predicted intact fuel pins, the 'True Positives' are all the correctly predicted missing fuel pins, the 'False Positives' are the intact fuel pins that are wrongly predicted as missing, and the 'False Negatives' are the missing fuel pins wrongly predicted as intact. The sum of True Negatives and False Positives equals the total number of intact fuel pins in the dataset, and the sum of True Positives and False Negatives equals the total number of missing pins.

The performance of the ANN model is then quantified using 4 metrics, i.e., the pin-accuracy, the precision, the recall and the F1 score.

The pin-accuracy A_p is the percentage of the correctly predicted fuel pins (the sum of the true positives and true negatives) out of the total number of fuel pins, considering all the fuel assemblies in the dataset, i.e.:

$$A_p = \frac{TN + TP}{N_{tot}} \quad (1)$$

where TN is the number of true negatives, TP is the number of true positives, and N_{tot} is the total number of fuel pins, both intact and missing, in the dataset.

The precision P is defined as the fraction of correctly predicted missing pins (the true positives) over all the pins predicted as missing (the sum of true and false positives):

$$P = \frac{TP}{TP + FP} \quad (2)$$

where FP is the number of false positives.

The recall R is equal to the fraction of correctly predicted missing pins (true positives) over the total number of missing pins in the dataset (equivalent to the sum of true positives and false negatives):

$$R = \frac{TP}{TP + FN} \quad (3)$$

where FN is the number of false negatives.

The $F1$ score is the harmonic mean of the precision and recall values, which reads as follows:

$$F1 = 2 \times \frac{P \times R}{P + R} \quad (4)$$

3. Results

The ANN was trained with three different options: (a) using only the thermal neutron flux (model N); (b) the neutron flux and the magnitude and direction of the gradient (model $N+G_m+G_d$); and (c) the neutron flux and the two components of the gradient (model $N+C_x+C_y$). A model that processes only the gradient of the neutron flux is not included in the study, but a previous work showed that it has a similar performance to model $N+G_m+G_d$ (Aldbissi et al., 2023). However, since the scalar neutron flux is needed to derive the gradient, it is available to be used as input to the algorithm at no additional cost.

Table 2

Confusion matrix of the model N (a), confusion matrix of the model $N+G_m+G_d$ (b), and the confusion matrix of the model $N+C_x+C_y$ (c).

N		Real Label	
		Intact (0)	Missing (1)
Predicted Label	Intact (0)	20087	4206
	Missing (1)	1572	3175

(a)

$N + G_m + G_d$		Real Label	
		Intact (0)	Missing (1)
Predicted Label	Intact (0)	19268	3322
	Missing (1)	2391	4059

(b)

$N + C_x + C_y$		Real Label	
		Intact (0)	Missing (1)
Predicted Label	Intact (0)	19641	2217
	Missing (1)	2018	5164

(c)

Table 3

Performance metrics for the three ANN models.

Metric	N		$N+G_m+G_d$		$N+C_x+C_y$	
	Mean value	Std. (\pm)	Mean value	Std. (\pm)	Mean value	Std. (\pm)
Pin accuracy	0.80	0.01	0.80	0.01	0.85	0.01
Precision	0.66	0.01	0.63	0.02	0.72	0.02
Recall	0.43	0.01	0.55	0.02	0.70	0.01
F1	0.52	0.01	0.58	0.02	0.70	0.01

As mentioned in Section 2.3 a grid search optimization was performed to determine the hyperparameters associated with each model. The grid for the optimization was defined using the following values: (10, 50, 75, 100, 150, 200, 250, 300, 350, 500) for the number of neurons, (250, 500, 750, 1000, 2000) for the number of epochs, and (1, 2, 3, 5, 10) for the batch size. The three models performed best with 200 neurons in the hidden layer, 1000 epochs, and a batch size of 3.

A comparison between the results obtained from the testing of the three models is carried out to highlight the effects of the gradient of the neutron flux on the identification of replaced fuel pins in SNF assemblies. The performance of the three models is investigated in terms of the overall number of intact and replaced fuel pins that are predicted correctly and the results are averaged over the 5 repetitions of the cross-validation process, see Section 3.1. The ability of the ANN models to reconstruct the full configuration of the fuel assemblies is discussed along with example cases taken from the best repetition of the cross-validation process out of the 5, see Section 3.2.

3.1. Performance at pin level

The results of each ANN model are summarized in the confusion matrices shown in Table 2 and the performances are quantified using the 4 metrics discussed in Section 2.4, i.e., the pin-accuracy, the precision, the recall and the F1 score, see Table 3. The values of the performance metrics are averaged over the five repetitions of the cross-validation process and their relatively low standard deviations reflect that the initial random shuffling of the dataset prior to the cross-validation has a minor effect on the performance of the network.

In terms of pin-accuracy, model N (only the thermal neutron flux) and model $N+G_m+G_d$ (thermal neutron flux together with the magnitude and direction of its gradient) have similar values. The first model is better to predict the intact pins, see Table 2a. The second model is better with missing pins, see Table 2b. Model $N+C_x+C_y$ (thermal neutron flux together with the two x and y components of the gradient)

has a higher pin accuracy since it predict correctly a larger number of missing fuel pins (true positives), see Table 2c.

The precision reflects the general ability of the model to avoid false predictions of both the intact and replaced fuel pins. Again, Model N and model $N+G_m+G_d$ have similar precision values. Model N predicts less missing fuel pins correctly (less true positives), but it gives less errors in terms of intact pins (less false positives). The second model provides a higher number of correct missing pins (more true positives), but it over-predicts pins as missing (more false positives). According to this metric, model $N+C_x+C_y$ performs better since it identifies a higher fraction of true positives (correct missing fuel pins) over false positives (misclassified intact fuel pins).

The recall is a measure of how well the model predicts missing fuel pins. The models that use the gradient (either in magnitude and direction or the two components) have larger recall values and hence can better detect missing fuel pins in comparison to the model based only on the thermal neutron flux. Despite the magnitude and direction of the gradient having a more immediate physical interpretation, the X and Y components are proven to be more beneficial to the ANN for the reconstruction of the diversion patterns (see further details in Section 3.2).

The F1 score is the harmonic mean of the precision and recall values, and confirms that the use of the neutron flux gradient is advantageous and that model $N+C_x+C_y$ performs better than the other considered models.

An under-estimation of missing fuel pins in SNF assemblies (higher number of false negatives) such as the case of the model that relies only on the thermal neutron flux is undesirable from a safeguards perspective since it can lead to diverted nuclear material being undetected.

3.2. Performance at assembly level

The ANN models discussed above, cannot fully reconstruct any of the diversion patterns. This is expected because the size of the dataset

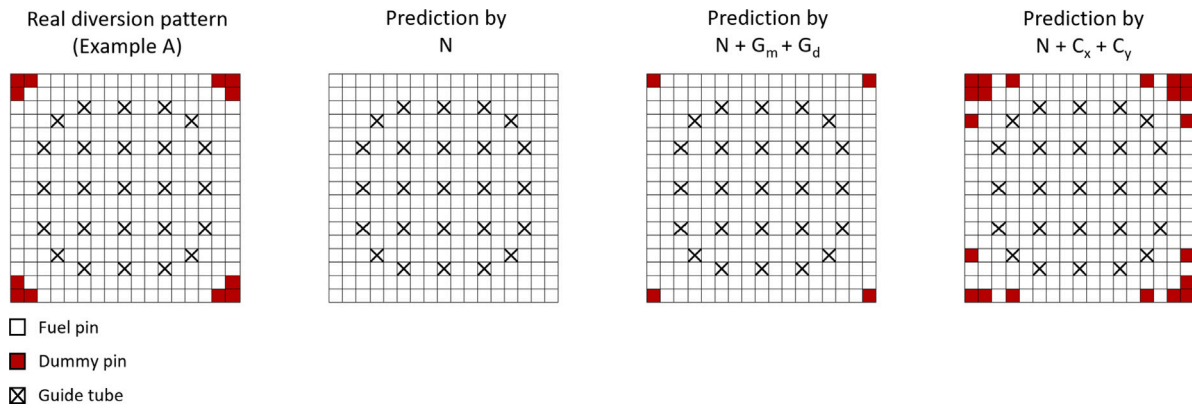


Fig. 3. Example A of a diversion case from the dataset and how it was reconstructed by the three ANN models.

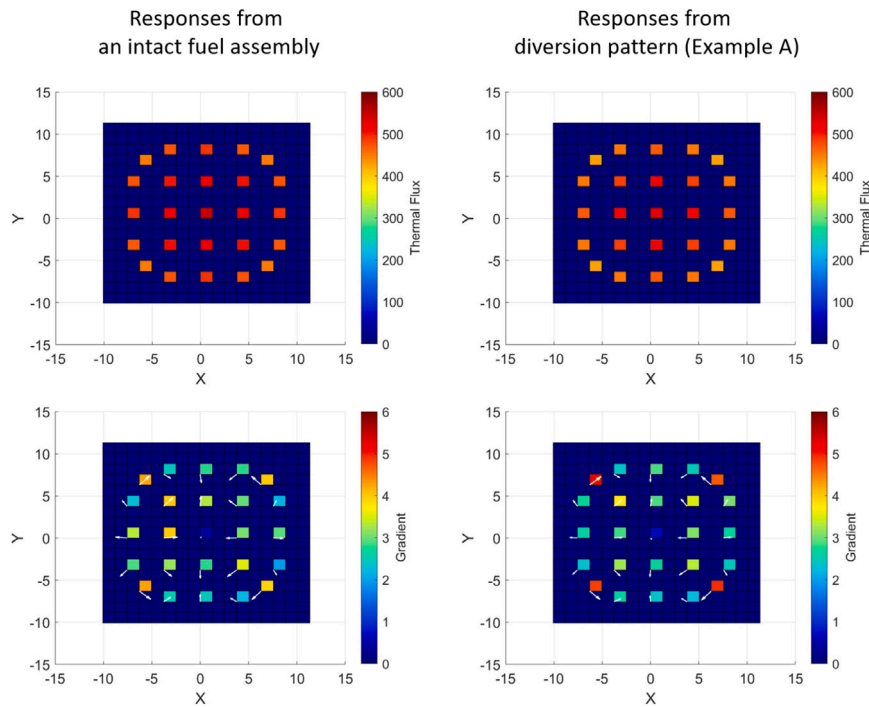


Fig. 4. The simulated responses of the thermal neutron flux and the gradient of the neutron flux from an intact fuel assembly and the diversion pattern from example A.

is relatively small. However, the majority of the predictions, especially from model $N+C_x+C_y$, are close to the real diversion patterns.

An example is shown in Fig. 3. The diversion pattern (referred to as example A) has three fuel pins replaced by dummy pins at each corner of the fuel assembly. Such a scenario is challenging to detect because the amount of replaced nuclear material is small, the affected locations are relatively far from the guide tubes (where detectors could be placed), and the pattern is symmetrical. The use of only the thermal neutron flux fails to detect any of the replaced fuel pins and the assembly is predicted as intact, but the two models trained with the gradient of the neutron flux can determine the lack of fuel pins at the correct positions. A comparison between the simulated signatures of an intact fuel assembly and example A is shown in Fig. 4. The thermal neutron flux in the guide tubes is affected only negligibly by the diversion, while the gradient has significant deviations, which are stronger in the guide tubes closest to the locations of the missing pins. Therefore, the use of the gradient leads to an improved performance of the machine learning algorithm.

In line with the results discussed in Section 3.1, the diversion pattern is better retrieved from the thermal neutron flux and the two

Cartesian components of the gradient than from the thermal neutron flux and the magnitude and direction of the gradient. As reported in Table 4, if the guide tubes close to the diversion (positions 4, 5, 21 and 22 in Fig. 1) are taken, the magnitude of the gradient has a large change in comparison with the case of the intact fuel assembly (between 25.65% and 35.31%) but not the direction (the absolute differences are about 5°), while the x- and y-component are both affected in a significant manner (from 14.05% up to 41.57%). Then, the two Cartesian components of the gradient give more diverse and sensitive information to the neural network. A similar behavior is observed in other symmetrical diversion scenarios included in the dataset.

On the other hand, there exist a few cases in which the predictions do not resemble the real diversion pattern at all. Example B shown in Fig. 5, is representative of these cases. In the assembly, two rows of fuel pins replaced by dummy pins next to the upper edge cause a significant disruption in the distribution of the thermal neutron flux and its gradient within the system, which can be easily seen in the simulated measurements inside the guide tubes. However, model $N+C_x+C_y$ (which processes the thermal neutron flux together with the two components of the gradient and provides the best predictions) fails

Table 4
Gradient in guide tubes 4, 5, 21, and 22, for the intact fuel assembly and Example A.

Guide tube #	Input feature	Intact	Example A	Relative difference (%)
4	Cx	3.02	4.11	36.09
	Cy	2.99	3.41	14.05
	Gm	4.25	5.34	25.65
	Gd	44.71	39.68	11.25
5	Cx	-2.75	-3.85	40
	Cy	2.88	3.36	16.67
	Gm	3.99	5.11	28.07
	Gd	133.68	138.89	3.9
21	Cx	3.31	3.96	19.64
	Cy	-2.7	-3.73	38.15
	Gm	4.27	5.44	27.4
	Gd	320.79	316.71	1.27
22	Cx	-2.67	-3.78	41.57
	Cy	-2.82	-3.64	29.08
	Gm	3.88	5.25	35.31
	Gd	226.56	223.92	1.17




Fig. 7. Example B (left) and its reconstruction via updated model $N+C_x+C_y$ (right).




Fig. 5. Example B of a diversion case (left) and its reconstruction via model $N+C_x+C_y$ (right).




Fig. 6. Additional diversion patterns for the updated dataset.

to reconstruct the pattern correctly. The reason for these anomalies is that example B as well as the cases with the same issue have no common features with the other configurations available in the dataset, therefore the network is not sufficiently trained to characterize them well.

To investigate this issue, the dataset is expanded by adding two new diversion scenarios that resemble example B, see Fig. 6. Then, model $N+C_x+C_y$ is trained with the updated dataset and tested on example B again, see Fig. 7. The predictions are significantly improved since the model has learned from two similar cases.

4. Conclusions

When inspecting SNF assemblies, characteristics signatures such as the neutron and gamma emissions are measured. An artificial neural network model can be used to process the measured signatures and reconstruct the specific system configuration in terms of fuel pins and

identify whether any fuel pins are replaced or not. In this work, the use of neutron flux gradient, which is not considered in standard procedures for SNF verification, is investigated in addition to the neutron flux for the characterization of possible diversions.

For the purpose of the study, a dataset that includes values of the neutron flux and its gradient in the guide tubes of an intact assembly and 109 assemblies with different patterns of replaced fuel pins, is generated using Monte Carlo simulations. Then, three ANN models are trained and tested with the dataset, i.e., a model for the analysis of only neutron flux measurements, a model which combines the neutron flux and the magnitude and direction of its gradient, and a model for the neutron flux together with the Cartesian components of its gradient. The comparison between the results of these model showed that the information about the neutron flux gradient is advantageous for the detection of patterns of replaced fuel pins within the assembly, and that the Cartesian components of the gradient are more effective than the magnitude and direction of the gradient.

The dataset and the cross-validation procedure are designed so that the testing of the ANN models is performed over “unknown” cases, which are not included in the training phase. Although the models above cannot fully identify any of the scenarios, the two models that use the neutron flux gradient do provide results close to most of the real assemblies and thus can generalize to some extent the mapping from the measured signatures to the patterns of replaced fuel pins. Yet, some diversion patterns cannot be reconstructed at all because they have no common features with the rest of the dataset. In order to overcome the issue, additional configurations are needed for the training. A dataset that includes all possible scenarios, is however not feasible because of the limitation of computational resources. Then, future work may explore sampling techniques that allow for better representations of the space of diversion scenarios. In the further development of these models, it will be relevant to consider datasets that also have SNF assemblies with different irradiation histories and with non-uniform burn-up profiles.

CRedit authorship contribution statement

Moad Al-dbissi: Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization, Writing – original draft, Writing – review & editing. **Imre Pázsit:** Writing – review & editing, Supervision, Conceptualization, Formal analysis, Funding acquisition, Methodology. **Riccardo Rossa:** Writing – review & editing, Supervision, Methodology, Formal analysis, Conceptualization. **Alessandro Borella:** Writing – review & editing, Supervision, Methodology, Formal analysis, Conceptualization. **Paolo Vinai:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Formal analysis, Conceptualization.

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

Data will be made available on request.

Acknowledgments

The project was financially supported by SCK CEN under grant agreement PO4500047684, and the Swedish Radiation Safety Authority under agreement SSM2021-786 and SSM2023-4389.

References

- Abadi, M., et al., 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. URL: <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Al-dbissi, M., Rossa, R., Borella, A., Pázsit, I., Vinai, P., 2023. Identification of diversions in spent PWR fuel assemblies by PDET signatures using artificial neural networks (ANNs). *Ann. Nucl. Energy* 193, 110005. <http://dx.doi.org/10.1016/j.anucene.2023.110005>.
- Aldbissi, M., Vinai, P., Borella, A., Rossa, R., Pázsit, I., 2022a. Conceptual design and initial evaluation of a neutron flux gradient detector. *Nucl. Instrum. Methods Phys. Res. A* 1026, 166030. <http://dx.doi.org/10.1016/j.nima.2021.166030>.
- Aldbissi, M., Vinai, P., Borella, A., Rossa, R., Pázsit, I., 2022b. Evaluation of the performance of a neutron gradient detector for partial defect testing in spent nuclear fuel assemblies. In: *Proc. INMM 63rd Annual Meeting*.
- Aldbissi, M., Vinai, P., Rossa, R., Borella, A., Pázsit, I., 2023. Using machine learning for the detection of missing fuel pins in spent nuclear fuel assemblies based on measurements of the gradient of the neutron flux. In: *Proc. INMM/ESARDA Joint Annual Meeting*.
- Avdic, S., Lindén, P., Pázsit, I., 2001. Measurement of the neutron current and its use for the localisation of a neutron source. *Nucl. Instrum. Methods A* 457 (3), 607–616. [http://dx.doi.org/10.1016/S0168-9002\(00\)00773-7](http://dx.doi.org/10.1016/S0168-9002(00)00773-7).
- Branger, E., Grape, S., Jansson, P., 2020. Partial defect detection using the DCVD and a segmented region-of-interest. *J. Instrum.* <http://dx.doi.org/10.1088/1748-0221/15/07/P07009>.
- Chollet, F., et al., 2015. Keras. <https://keras.io>.
- Leppänen, J., 2015. The Serpent Monte Carlo code: Status, development and applications in 2013. *Ann. Nucl. Energy* 82, 142–150.
- Lindén, P., Karlsson, J.K.H., Dahl, B., Pázsit, I., Por, G., 1999. Localisation of a neutron source using measurements and calculation of the neutron flux and its gradient. *Nucl. Instrum. Methods A* 438 (2–3), 345–355. [http://dx.doi.org/10.1016/S0168-9002\(99\)00830-X](http://dx.doi.org/10.1016/S0168-9002(99)00830-X).
- Mayorov, M., White, T., Lebrun, A., Brutscher, J., Keubler, J., Birnbaum, A., Ivanov, V., Honkamaa, T., Peura, P., Dahlberg, J., 2017. Gamma emission tomography for the inspection of spent nuclear fuel. In: 2017 IEEE Nuclear Science Symposium and Medical Imaging Conference. NSS/MIC, pp. 1–2. <http://dx.doi.org/10.1109/NSSMIC.2017.8533017>.
- Pázsit, I., 1997. On the possible use of the neutron current in core monitoring and noise diagnostics. *Ann. Nucl. Energy* 24 (15), 1257–1270. [http://dx.doi.org/10.1016/S0306-4549\(97\)00037-6](http://dx.doi.org/10.1016/S0306-4549(97)00037-6).
- Rinard, P.M., Bosler, G.E., 1988. Safeguarding LWR spent fuel with the FORK detector. Los Alamos National Laboratory report LA-11096-MS.
- Vitullo, F., Lamirand, V., Mosset, J.-B., Frajtag, P., Pakari, O., Perret, G., Pautz, A., 2020. A mm^3 fiber-coupled scintillator for in-core thermal neutron detection in CROCUS. *IEEE Trans. Nucl. Sci.* 67 (4), 625–635. <http://dx.doi.org/10.1109/TNS.2020.2977530>.
- Watanabe, K., Yamanaka, M., Endo, T., Pyeon, C.H., 2020. Real-time subcriticality monitoring system based on a highly sensitive optical fiber detector in an accelerator-driven system at the Kyoto University Critical Assembly. *J. Nucl. Sci. Technol.* 57 (2), 136–144. <http://dx.doi.org/10.1080/00223131.2019.1647895>.