Identification of diversions in spent PWR fuel assemblies by PDET signatures using Artificial Neural Networks (ANNs)

(article starts on next page)

Annals of Nuclear Energy 193 (2023) 110005

Contents lists available at ScienceDirect

Annals of Nuclear Energy

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

Identification of diversions in spent PWR fuel assemblies by PDET signatures using Artificial Neural Networks (ANNs)

Moad Al-dbissi [a,b,*], Riccardo Rossa [b], Alessandro Borella [b], Imre Pázsit [a], 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:
Artificial Neural Networks
Spent nuclear fuel
Nuclear safeguards
Machine learning
Radiation detection

ABSTRACT

Spent nuclear fuel represents the majority of materials placed under nuclear safeguards today and it requires to be inspected and verified regularly to promptly detect any illegal diversion. Research is ongoing both on the development of non-destructive assay instruments and methods for data analysis in order to enhance the verification accuracy and reduce the inspection time. In this paper, two models based on Artificial Neural Networks (ANNs) are studied to process measurements from the Partial Defect Tester (PDET) in spent fuel assemblies of Pressurized Water Reactors (PWRs), and thus to identify at different levels of detail whether nuclear fuel has been replaced with dummy pins or not. The first model provides an estimation of the percentage of replaced fuel pins within the inspected fuel assembly, while the second model determines the exact configuration of the replaced fuel pins. The two models are trained and tested using a dataset of Monte-Carlo simulated PDET responses for intact spent PWR fuel assemblies and a variety of hypothetical diversion scenarios. The first model classifies fuel assemblies according to the percentage of diverted fuel with a high accuracy (96.5%). The second model reconstructs the correct configuration for 57.5% of the fuel assemblies available in the dataset and still retrieves meaningful information of the diversion pattern in many of the misclassified cases.

1. Introduction

One of the most important tasks in nuclear power plants under safeguards is regular inspections to verify that no special nuclear material is missing from the Spent Nuclear Fuel (SNF) assemblies. This task is crucial before the assemblies are transferred to dry cask storage or encapsulated for permanent repository disposal (Vaccaro et al., 2018), where further inspections become more challenging or impossible. In the safeguards community, the search for missing/replaced fuel pins in SNF assemblies is known as detection of partial defects (IAEA, 2022; Lee and Yim, 2020). Spent nuclear fuel is particularly sensitive from a safeguards perspective because of its residual fissile material such as ^{235}U and ^{239}Pu. In the recent years, about 80% of the material placed under safeguards was plutonium contained in SNF (IAEA, 2014).

Several methods of Non-Destructive Assay (NDA) such as the Digital Cherenkov Viewing Device (DCVD) (Branger et al., 2020), the Fork Detector (FD) (Rinard and Bosler, 1988) and the Passive Gamma Emission Tomography (PGET) (Mayorov et al., 2017) among others are used to detect possible diversions in SNF assemblies. These techniques are approved for inspection by the International Atomic Energy Agency (IAEA) and have been extensively applied for many years (IAEA, 2011). The processing and the interpretation of the measurements performed with these techniques mainly relies on the expert judgement of the inspectors. In addition, such investigations are focused on the coarse detection of possible illicit diversion of nuclear material with the exception of the PGET which can provide pin-level resolution.

Recent efforts have been conducted to develop methods that can enhance the processing of the measured data and extract more details of the system configuration. For example, machine learning algorithms were used to quantify the percentage of replaced fuel pins in SNF assemblies (Rossa et al., 2020, 2018; Aldbissi et al., 2022), to predict parameters of SNF assemblies (Mishra et al., 2021), to detect and localize missing radioactive sources within a small grid (Durbin and Lintereur, 2020), to track elemental and isotopic material flows through material balance areas for safeguards (Shoman and Cipiti, 2018). These methods can help to reduce the inspection time and make the identification of diversion patterns more precise, so that the decision process of the inspectors is facilitated.

In this paper, novel insights are discussed about the application of Artificial Neural Networks (ANNs) to quantify and characterize possible partial defects in SNF assemblies of Pressurized Water Reactors

* 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).

https://doi.org/10.1016/j.anucene.2023.110005
Received 6 April 2023; Received in revised form 21 June 2023; Accepted 29 June 2023
Available online 25 July 2023

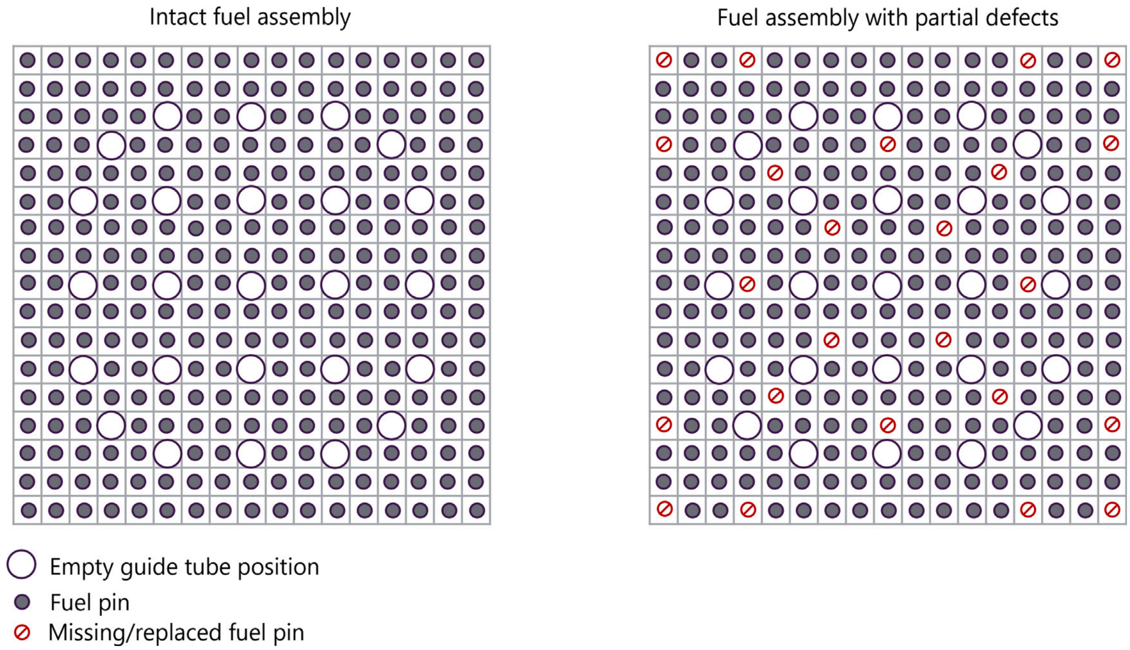

Fig. 1. Intact 17 × 17 PWR spent nuclear fuel assembly (left) and an example of a diversion scenario (right).

(PWRs). For this purpose, a synthetic dataset is used and includes simulated neutron and gamma measurements according to the Partial Defect Tester (PDET), an inspection technique developed by Lawrence Livermore National Laboratory (Sitaraman and Ham, 2007; Ham and Sitaraman, 2011; Ham et al., 2015). The specific objective is to study the performance of two ANN models for the prediction of the total amount of replaced fuel pins within a SNF assembly and for the determination of the exact diversion pattern, respectively.

The paper is structured as follows. The methodology together with the description of the ANN models is introduced in Section 2. The model for the estimation of the percentage of fuel pins replaced in a SNF assembly is examined in Section 3. The performance of the model for the identification of intact and diverted fuel pins in a SNF assembly is discussed in Section 4. Conclusions are drawn in Section 5.

2. Methodology

The general strategy for the verification of the integrity of Spent Nuclear Fuel (SNF) assemblies is to acquire measurements of observable quantities, such as the neutron flux within the assembly, and determine whether the outcome of the measurements is consistent with the declared configuration of the assemblies or not. An underlying assumption is that there is a one-to-one correspondence between the spatial distribution of the observables and the actual composition of the fuel assembly, whether intact or not, which is the basis of the identification of the defects. Furthermore, that the measurable quantities can also be simulated by core calculation methods for any intact or defect fuel assembly.

An approach for safeguards inspections of SNF assemblies is based on the Partial Defect Tester (PDET). The first prototype of the PDET was developed by Lawrence Livermore National Laboratory and consists of a set of neutron fission chambers and gamma-ray ionization chambers that can be inserted in the guide tubes of assemblies of Pressurized Water Reactors (PWRs) (Sitaraman and Ham, 2007; Ham and Sitaraman, 2011; Ham et al., 2015). The measured quantities from the PDET are the passive emission of neutrons and gamma-rays taken at the empty guide tube positions in the SNF assembly. Therefore, these measurements at different locations within the system provide more detailed information for the retrieval of the configuration of the fuel pins.

In current work, a dataset of simulated PDET responses in intact PWR assemblies and diversion scenarios is used to train and test two ANNs models for the identification of possible partial defects at different levels of detail from the PDET data. Artificial neural networks are well-suited for these types of problems (i.e., the task of reconstructing a system configuration from the observables) because they are effective in modelling non-linear relationships between input and output and recognizing complex data patterns.

2.1. PDET dataset

The dataset for the study relies on previous work performed at SCK CEN (Rossa et al., 2020, 2018). It includes the PDET responses simulated via Monte Carlo N-Particle - MCNP code (Werner, 2018) for 17 × 17 PWR spent nuclear fuel assemblies with and without diversion, see Fig. 1.

The dataset contains 196 intact fuel assemblies. Each of them is unique in terms of its Initial Enrichment (IE), Cooling Time (CT) and Burn-Up (BU). There are 107 modelled diversion patterns, both symmetrical and asymmetrical, and have a minimum of 4 up to a maximum of 180 fuel pins replaced by dummy pins made of stainless steel. Each of the 107 diversion scenarios is repeated 9 times, but with different conditions of IE, CT and BU, so the overall number of fuel assemblies with missing fuel pins is 963. The values of IE, CT and BU, which are used for the intact and diverted cases in all the possible combinations, are summarized in Table 1.

Strictly speaking, the one-to-one correspondence between the neutron or gamma-ray flux and the defect configuration exists only between fuel assemblies of the same IE, BU and CT values. It would therefore be a more straightforward procedure if the training set contained only intact and various defect configurations for assemblies of the same IE, BU and CT values as the assembly to be identified. However, in practice, in view of the multitude of assemblies with different IE, BU and CT values, this would require a huge dataset. It is therefore assumed that the relationship between the flux distributions and the defect structure is only mildly affected by the different values of the IE, BU and CT parameters. The study performed actually gives some indications of the significance of these parameters, as it is discussed later.

Table 1
Values of BU, IE and CT included in the PDET dataset.

	Intact fuel assemblies	Diversion scenarios
Burn-Up (MWd/kgU)	5, 10, 15, 20, 30, 40, 60	10, 30, 60
Initial Enrichment (w%)	2, 2.5, 3, 3.5, 4, 4.5, 5	2, 3.5, 5
Cooling Time (years)	1, 5, 10, 50	5

Since the arrangement of a 17×17 PWR fuel assembly has 25 guide tubes (see Fig. 1), the calculated PDET responses for each configuration are 50, i.e., 25 thermal neutron detection rates and 25 gamma-ray detection rates.

2.2. ANN models

Two ANN models are developed to process the PDET responses associated with a PWR spent nuclear fuel assembly and characterize the possible partial defects. The first model identifies a possible diversion in terms of the percentage of replaced fuel pins. The second model identifies if fuel pins are removed and at which positions within the fuel assembly.

Both models are built using the Tensorflow (Abadi et al., 2015) and the Keras (Chollet et al., 2015) open-source software libraries, and they are based on an artificial neural network with 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 (i.e., the backward tuning of the weights and the biases of the neurons through the comparison between the estimated and the correct output) with an efficient convergence rate. The weights and the learning rate of the network are optimized with the Adaptive Moment Estimation (ADAM).

For the estimation of the percentage of replaced fuel pins, the algorithm is designed to process the input associated with a fuel assembly and then assign the fuel assembly to one of the prescribed classes, which correspond to different percentage ranges of diverted material (See Table 3 for the definition of the classes). To obtain an appropriate solution of this so-called multiclass problem, the SoftMax activation function is selected for the output layer and the error (loss) of the algorithm in the optimization process is evaluated from the Categorical Cross-Entropy loss function.

In the case of the identification of the locations where the nuclear fuel pins have been replaced with dummies, the algorithm processes the PDET responses for a fuel assembly and gives as output a probability of being replaced to each of the fuel pins. This is obtained by applying the Sigmoid activation function in the output layer. The Sigmoid function is a typical choice for outputs that are non-mutually exclusive such as in the current application, where each pin is treated independently and can be either present or replaced. For practical reasons, if the probability of a fuel pin to be identified is between 0.5 and 1, the fuel pin is labelled as missing, while, if the probability is less than 0.5, the fuel pin is labelled as present. Then, the problem can be considered a multi-label binary classification, and the performance of the ANN model can be adequately evaluated with the Binary Cross-Entropy loss function.

The input to the ANN models is provided with the PDET responses for the fuel assembly under investigation, which comprises 50 values (see Section 2.1). Therefore, the number of neurons in the input layer of the ANN models is fixed to 50. The set of neutron emissions and the set of gamma-ray emissions obtained from PDET and used as input, are separately normalized so that they are on the same scale and their processing is more consistent.

Table 2
Hyper-parameters optimized for the ANN models using grid search.

	Multi-class model (% of replaced fuel pins)	Multi-label model (individual fuel pins)
Number of Epochs	2000	1000
Batch-size	10	25
Neurons in the hidden layer	50	300

Table 3
Percentage of diversion in fuel assemblies and prescribed class labels.

Percentage of replaced pins (x)	Class label
$x = 0$	0
$x \leq 10\%$	1
$10\% < x \leq 20\%$	2
$20\% < x \leq 30\%$	3
$30\% < x \leq 40\%$	4
$40\% < x \leq 50\%$	5
$x > 50\%$	6

The size of the output layer depends on the model. For the estimation of the percentage of the replaced fuel, the output layer includes a neuron for each of the prescribed classes. For the identification of which fuel pins may be replaced, the output layer has a neuron for each of the fuel pins in the fuel assembly.

The ANN models are trained and tested via an N-fold cross-validation process. Accordingly, the whole dataset is shuffled and divided into N random batches. N-1 of these batches are used to train the ANN, while the remaining one is used for the testing. The procedure is repeated N times so that each of the N batches serves as testing dataset one time. The model accuracy is then taken as the average of the accuracy from the N batches which leads to a less biased result.

A grid search optimization is performed to determine the number of epochs and the batch size in the training process, and the number of neurons for the hidden layer. The values for the two models are summarized in Table 2.

3. Percentage of replaced fuel pins

The first ANN model categorizes a fuel assembly with respect to a set of classes based on the percentage of diverted material. Seven classes are prescribed and are reported in Table 3. The class label 0 is for intact fuel assemblies and the class labels 1 to 6 indicate fuel assemblies with progressively higher numbers of replaced fuel pins.

The ANN model is trained and tested according to a 5-fold cross-validation process, and it reaches an accuracy of 96.5%. More insights are provided by the confusion matrix that summarizes the correct and incorrect predictions for each class, see Fig. 2.

The confusion matrix shows that the misclassified fuel assemblies fall into one or two class higher or lower than their true class. As mentioned in Section 2.2, the model estimates the probability that a fuel assembly belongs to any of the prescribed classes and assigns the fuel assembly to the class with the highest probability. In the current misclassifications, the correct class has always the second highest probability, and the relative differences between the probabilities of the true and predicted labels are below 5% in most cases, see Fig. 3.

The majority of the misclassifications are in classes 0 and 1. This bias might be expected since the number of intact fuel assemblies in the dataset is bigger than the number of fuel assemblies with a specific diversion pattern and since the class with true label 1 (maximum partial defect of 10%) is the closest to the class of intact fuel assemblies with true label 0. On one hand, the cases with partial defect belonging to class 1 but identified as intact cases (false negatives), are of severe concern because diverted material goes undetected. Examples of these misclassified fuel assemblies are shown in Fig. 4 and they share common characteristics, i.e., the removal is symmetric, in a checkered-like

		Predicted Label						
		0	1	2	3	4	5	6
True Label	0	188	8	0	0	0	0	0
	1	18	153	0	0	0	0	0
	2	0	4	208	2	2	0	0
	3	0	2	0	184	2	1	0
	4	0	0	0	1	143	0	0
	5	0	0	0	0	0	144	0
	6	0	0	0	0	0	0	99

Fig. 2. Classification of fuel assemblies based on % of replaced fuel pins; confusion matrix.

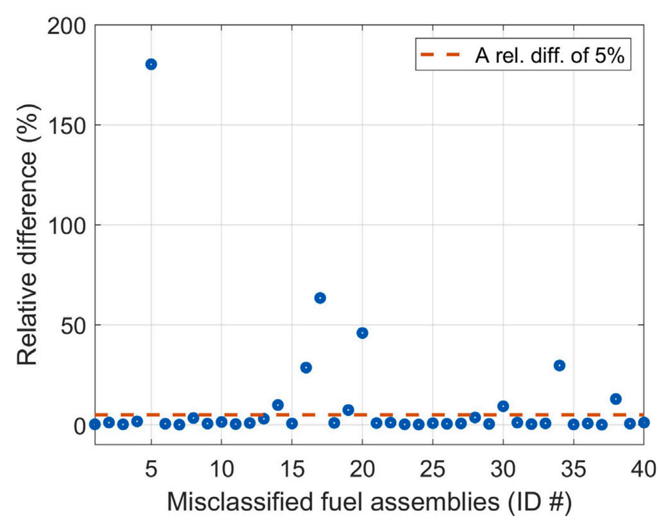

Fig. 3. Classification of fuel assemblies based on % of replaced fuel pins; relative differences between probabilities associated with true and wrongly predicted class label.

pattern, and more focused on the outer edges. On the other hand, intact cases predicted in class 1 (false positives/false alarms) are a less sensitive issue, even though they would require additional resources and time for clarification during an inspection.

The analysis also indicates that higher numbers of misclassified diversions are found at BU of 10 MWd/kgU and IE of 2 w%, BU of 10 MWd/kgU and IE of 5 w%, and BU of 30 MWd/kgU and IE of 2 w%, see Fig. 5. The tendency could be related to the nature of the dataset used to train the model. As described in Section 2.1, the dataset covers more variations in BU, IE and CT for the intact cases than for the diversion scenarios. In addition, the intact fuel assemblies with low BU (between 5 and 30 MWd/kgU) are larger in number than the ones with high BU (between 30 and 60 MWd/kgU), while their values of IE are evenly distributed between 2 and 5 w%. Further investigations are needed to understand the effects of these parameters on the model performance and to identify possible sources of bias.

4. Identification of the replaced fuel pins

The second ANN model introduced in Section 2.2 is applied to determine the exact configuration of the 264 fuel pins within a 17 × 17 PWR fuel assembly. Each individual fuel pin in the assembly is considered and its probability of being replaced is predicted by processing the PDET data. Then, the size of the model output is equal to the number of fuel pins.

Given a 5-fold cross validation process, the model can reconstruct the exact arrangement of the fuel pins in 667 out of the 1159 fuel assemblies available in the PDET dataset, which corresponds to a fraction of 57.5%. The predictions are correct for 97.4% of the intact fuel assemblies (191 out of 196) and 49.4% of the diversion scenarios (476 out of 963). As expected, the performance is better with the cases without partial defects because in the training the intact configurations are more numerous than any of the specific diverted configurations. Yet, the algorithm can detect 94.8% of all the incomplete fuel assemblies as diversion scenarios despite the incorrect number/location of the replaced fuel pins.

As discussed in Section 2.2, the model estimates the probability of a fuel pin to be replaced within the fuel assembly. If the probability is higher than the threshold value of 0.5, the fuel pin is labelled as missing, otherwise as present. The analysis of the distribution of the probabilities for the fuel pins that are correctly and wrongly predicted can provide insights into the behaviour of the model.

The distribution of the probabilities for all the correctly predicted fuel pins is shown in Fig. 6. Two large peaks are respectively found close to the probability values of 0 and 1, and reflect the high confidence in the correct results of the model.

The probability distribution of the wrongly predicted fuel pins is shown in Fig. 7. A bigger portion of the misclassifications (64.5%) has probability less than 0.5 and thus consists of replaced fuel pins predicted as intact, i.e., false negatives. The distribution has two peaks respectively near the probabilities of 0.5 and 0. The misclassifications with probabilities around 0.5 are characterized by a low level of confidence because the difference in probability with the other label (which is the true label) is small. The misclassifications with probabilities close to 0 (thus, far from the threshold) have a higher level of confidence. The tendency to make more misclassifications in favour of false negatives and with higher levels of confidence may depend on the training dataset, where the fraction of the intact fuel pins is large (78.4%). The construction of a more balanced dataset to avoid this type of bias is not

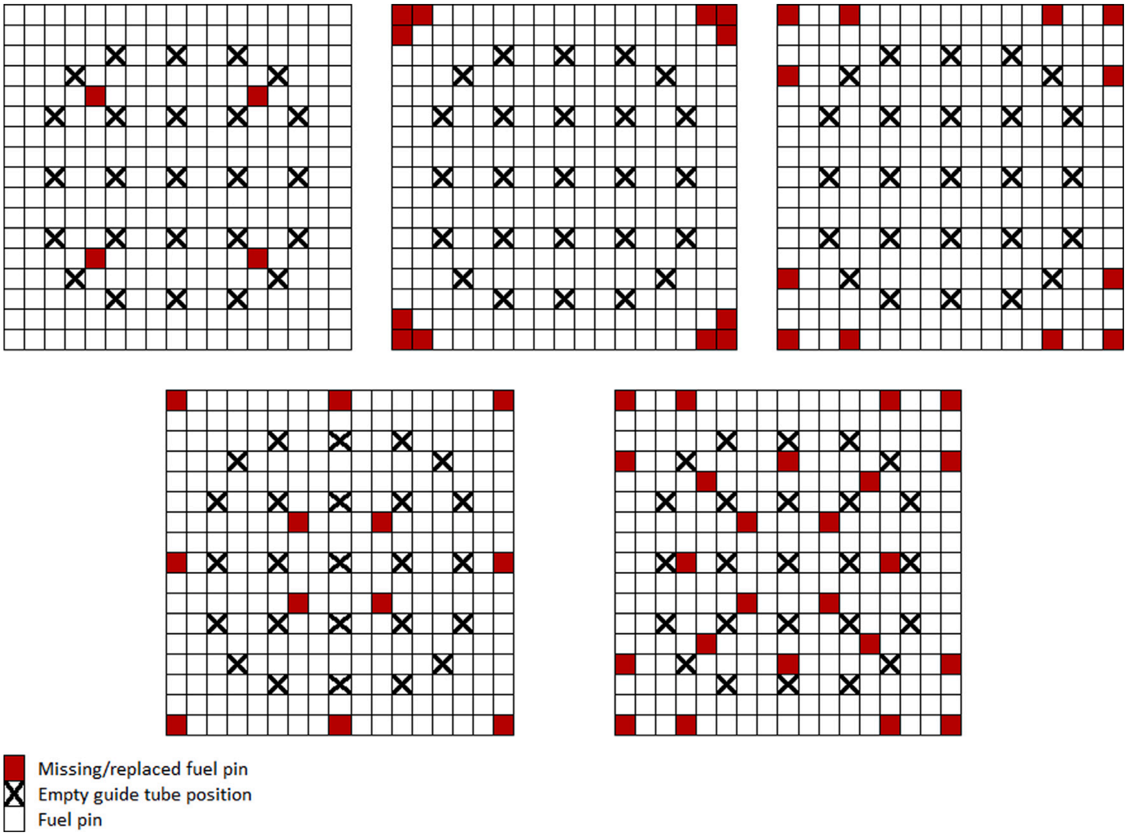

Fig. 4. Classification of fuel assemblies based on % of replaced fuel pins; false negatives.

		Initial enrichment (w%)		
		2	3.5	5
Burn-up (MWd/kgU)	10	7	3	6
	30	6	3	2
	60	3	1	1

Fig. 5. Classification of fuel assemblies based on % of replaced fuel pins; number of misclassifications with respect to BU and IE.

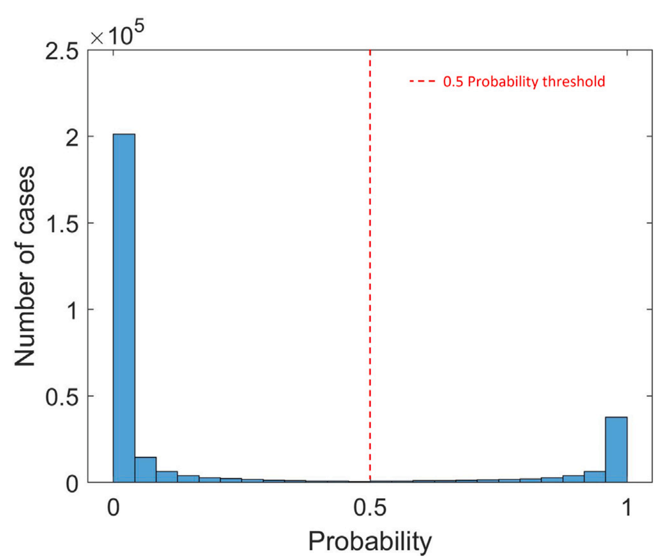

Fig. 6. Classification of the individual fuel pins; probability distribution of the correctly classified fuel pins.

straightforward. Adding more diversion scenarios does not necessarily increase the weight of the replaced pins since fuel assemblies with realistic partial defects still have a significant number of intact fuel pins. Approaches to improve the dataset and to correct possible biases will be studied in the future work.

The misclassified fuel assemblies are 492 in total, and the associated errors may involve one or more fuel pins, see Table 4. The majority of these fuel assemblies have a relatively low number of incorrect fuel pins (between 1 and 20) and are therefore reproduced nearly properly to a significant extent. Fig. 8 shows examples of predicted diversion patterns with different numbers of misclassified fuel pins. Although they are not

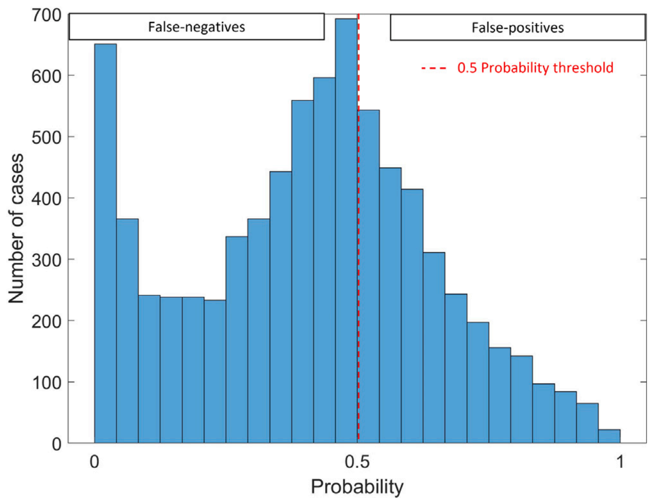

Fig. 7. Classification of the individual fuel pins; probability distribution of the misclassified fuel pins.

Table 4
Number of fuel assemblies with x incorrect fuel pins.

Number of misclassified pins (x)	Number of cases	Fraction of the dataset (%)
$x = 0$	667	57.5
$1 \leq x \leq 10$	238	20.5
$10 < x \leq 20$	160	13.8
$20 < x \leq 30$	46	4.0
$30 < x \leq 40$	19	1.6
$40 < x \leq 50$	8	0.7
$50 < x \leq 60$	6	0.5
$x > 60$	15	1.3

accurate, they can provide a useful indication of the main region of the real diversion, with the exception of the case with 67 wrong fuel pins.

5. Conclusions

Two models based on Artificial Neural Networks (ANNs) are assessed for the task of determining possible partial defects in PWR spent nuclear fuel assemblies. The models are trained and tested using a dataset of simulated responses of the Partial Defect Tester for both intact fuel assemblies and diversion scenarios with different values of burn-up, initial enrichment and cooling time.

The first model assigns spent fuel assemblies to classes defined by percentage ranges of replaced fuel pins. For this purpose, the algorithm processes the PDET information for one fuel assembly, estimates the probability of the fuel assembly to be in each of the prescribed ranges, and assigns the fuel assembly to the range with the highest probability. The results show a classification accuracy of 96.5%. For the misclassifications, the relative differences between the probabilities of the predicted and the true class is below 5% with few exceptions. In addition, the majority of the misclassifications occur between the class of intact fuel assemblies and the class with the smallest partial defects (i.e., less than 10% of replaced fuel pins). Further investigation is required to clarify the effect of the dataset on the model performances when considering different burn-up, initial enrichment and cooling times of the spent fuel.

The second model is developed to identify the presence or absence of each individual fuel pin inside an assembly and thus to retrieve the exact arrangement of the fuel pins from the analysis of the PDET responses. The algorithm estimates the probability of the fuel pins to be replaced and labels the fuel pins with probability higher than the threshold value of 0.5 as missing, otherwise as present. Given the training and testing based on the current PDET dataset, the model fully reconstructs 57.5% of the fuel assemblies and with a high level of confidence. A substantial majority of the misclassified fuel pins consists of false negatives, i.e., replaced fuel is diagnosed as present. The estimated probability of these fuel pins is below 0.5 and the distribution has a large peak close to zero. The aspect of many diverted fuel pins misclassified with probabilities close to zero is important because a high level of confidence is associated with the error and thus it becomes difficult to understand how trustworthy the results are. This bias can be related to the existing imbalance of labels in the datset used for the training process (the intact fuel pins are 78.4% of the total). In addition, a large fraction of the misclassified fuel pins has probabilities around the threshold value of 0.5, so the wrong labelling of the fuel pins is assigned with a lower confidence.

Future work needs to address the issue on the high confidence in the erroneous classifications so that the interpretation of the results can be more robust, and to study methods for the correction of the bias towards false negatives. Nevertheless, in a relevant quantity of incorrect fuel assemblies the predicted diversion patterns differ from the real ones by few fuel pins (e.g., 238 out of the total 1159 fuel assemblies are reproduced with 1 to 10 incorrect fuel pins) and therefore still provide meaningful information on the partial defect and its spatial distribution.

The results presented in the paper are obtained from a 5-fold cross validation procedure where the training and testing batches are randomly selected from the dataset. Since the dataset collects multiple cases of the same assembly configuration (which vary because of burn-up, initial enrichment and cooling time), the ANN models are simultaneously trained and tested using patterns of intact and replaced fuel pins that are often the same. Then, it will be also important to investigate the full capabilities of these models with respect to diversions that are not part of the training phase.

CRediT authorship contribution statement

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

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

Acknowledgements

The project was financially supported by SCK CEN under grant agreement PO4500047684, and the Swedish Radiation Safety Authority under agreements SSM2021-709 and SSM2021-786.

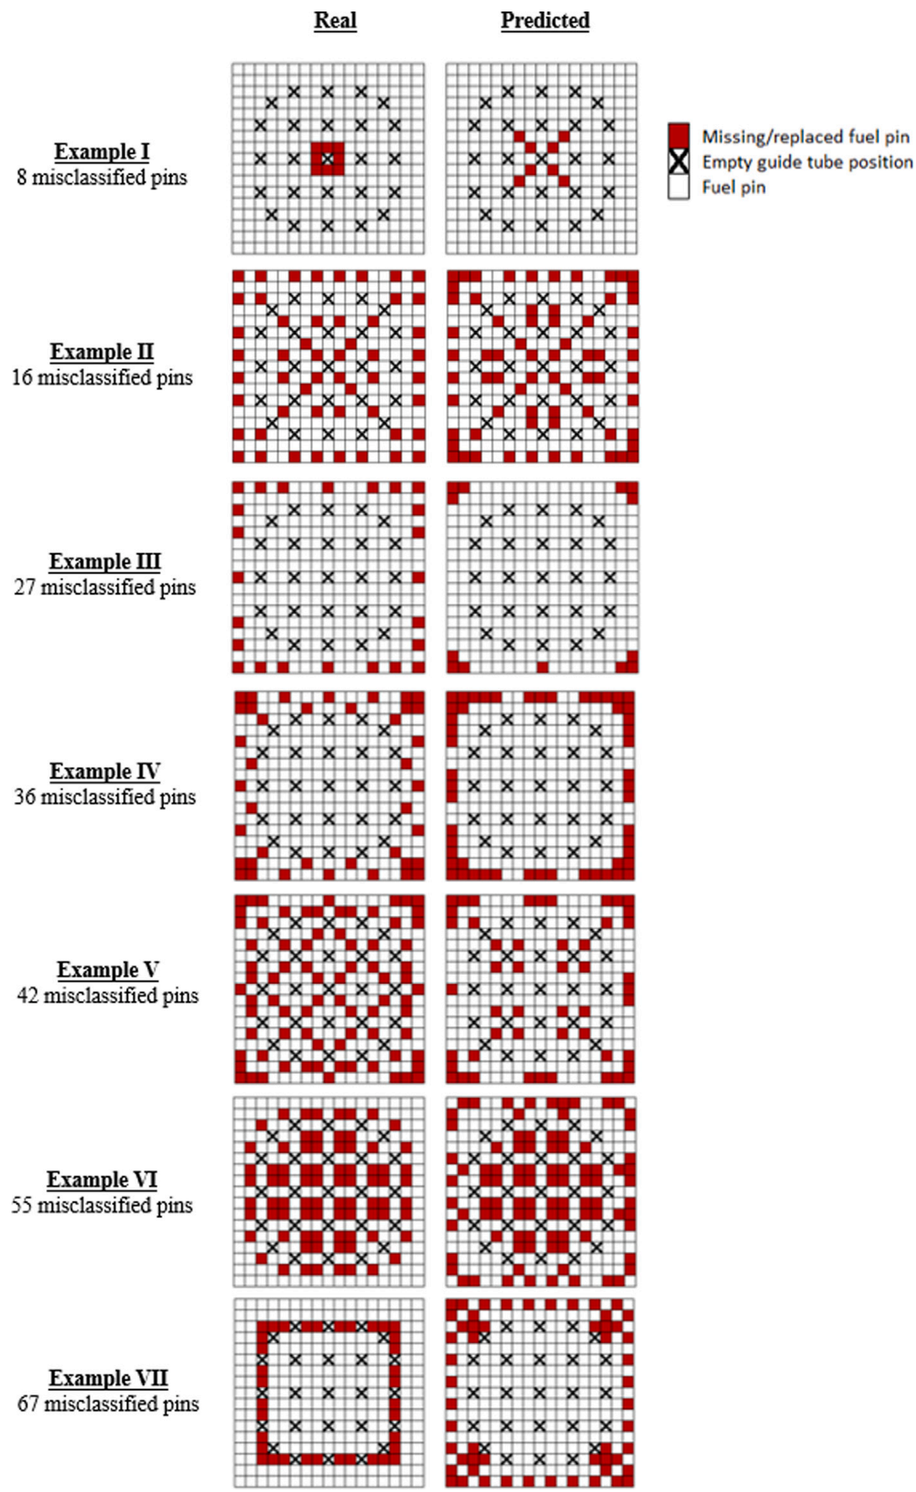

Fig. 8. Examples of misclassified diversion scenarios.

References

Abadi, M., et al., 2015. TensorFlow: Large-scale machine learning on heterogeneous systems.

Aldbissi, M., Borella, A., Pázsit, I., Rossa, R., Vinai, P., 2022. Optimizing neural networks to detect replaced spent fuel pins using the partial defect tester. In: IAEA Symposium on International Safeguards. IAEA, Vienna.

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.

Durbin, M., Lintereur, A., 2020. Implementation of machine learning algorithms for detecting missing radioactive material. J. Radioanal. Nucl. Chem. 324 (3), 1455–1461.

Ham, Y., Kerr, P., Sitaraman, S., Swan, R., Rossa, R., Liljenfeldt, H., 2015. Partial defect verification of spent fuel assemblies by PDET: Principle and field testing in interim spent fuel storage facility (CLAB) in Sweden. In: The 4th International Conference on Advancements in Nuclear Instrumentation Measurement Methods and their Applications. ANIMMA.

Ham, Y., Sitaraman, S., 2011. Partial defect tester: A novel approach to detect partial defects in pressurized water reactor spent fuel. Nucl. Technol. 175 (2), 401–418. http://dx.doi.org/10.13182/NT11-A12312.

IAEA, 2011. Safeguards techniques and equipment. In: International Nuclear Verification Series, No. 1 (Rev. 2), Vienna, Austria.

IAEA, 2014. Annual Report 2013, GC(58)/3, Vienna, Austria.

IAEA, 2022. Safeguards Glossary 2022 Edition.

Lee, H., Yim, M.S., 2020. Investigation of a fast partial defect detection method for safeguarding PWR spent fuel assemblies. Ann. Nucl. Energy 144, 107496. http://dx.doi.org/10.1016/j.anucene.2020.107496.

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.

Mishra, V., Branger, E., Elter, Z., Grape, S., Jansson, P., 2021. Comparison of supervised machine learning algorithms to predict PWR spent fuel parameters. In: INMM/ESARDA Joint Annual Meeting 2021.

Rinard, P.M., Bosler, G.E., 1988. Safeguarding LWR spent fuel with the FORK detector. Los Alamos National Laboratory report LA-11096-MS.

Rossa, R., Borella, A., Giani, N., 2020. Comparison of machine learning models for the detection of partial defects in spent nuclear fuel. Ann. Nucl. Energy 147, 107680. http://dx.doi.org/10.1016/j.anucene.2020.107680.

Rossa, R., Borella, A., van der Meer, K., 2018. Comparison of the SINRD and PDET detectors for the detection of fuel pins diversion in PWR fuel assemblies. In: INMM-59 Annual Meeting.

Shoman, N., Cipiti, B.B., 2018. Unsupervised machine learning for nuclear safeguards. Technical Report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States).

Sitaraman, S., Ham, Y., 2007. Characterization of a safeguards verification methodology to detect pin diversion from pressurized water reactor (PWR) spent fuel assemblies using Monte Carlo techniques. In: INMM 48th Annual Meeting, Tucson, AZ.

Vaccaro, S., Gauld, I., Hu, J., De Baere, P., Peterson, J., Schwalbach, P., Smejkal, A., Tomanin, A., Sjöland, A., Tobin, S., Wiarda, D., 2018. Advancing the Fork detector for quantitative spent nuclear fuel verification. Nucl. Instrum. Methods Phys. Res. A 888, 202–217. http://dx.doi.org/10.1016/j.nima.2018.01.066.

Werner, C.J., 2018. MCNP6.2 release notes.