Reconstructing patterns of missing fuel pins in PWR spent nuclear fuel assemblies via CNNs

Annals of Nuclear Energy 237 (2026) 112465

Contents lists available at ScienceDirect

Annals of Nuclear Energy

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

Reconstructing patterns of missing fuel pins in PWR spent nuclear fuel assemblies via CNNs

Moad Al-dbissi *, Imre Pázsit, Paolo Vinai

Department of Physics and Astronomy Chalmers University of Technology, SE-412 96, Gothenburg, Sweden

ARTICLE INFO

Keywords:
Machine learning
Spent nuclear fuel
Convolutional Neural Networks
Neutron flux
Monte Carlo
Nuclear safeguards

ABSTRACT

Convolutional Neural Networks (CNNs) are investigated for identifying and localizing missing fuel pins in 17 × 17 PWR spent nuclear fuel assemblies under partial defect scenarios, where selected fuel pins are replaced with dummy pins. A CNN based on the Inception architecture is developed and trained using Monte Carlo–simulated data. Thermal neutron flux, fast neutron flux, and their spatial gradient are tallied at the 24 guide tube positions and used as input features. In return, the network provides a binary classification for each fuel pin, (0) for intact pins and (1) for replaced pins. Results show that the network performs better at identifying intact fuel pins than replaced ones. Incorporating the x and y components of the flux gradient alongside the scalar flux significantly improves the classification of both intact and replaced pins. A more detailed evaluation reveals that, despite strong global F1 scores, substantial variability remains in the detection of replaced fuel pins across different scenarios. However, a proximity-based classification criterion shows that the CNN reliably identifies the correct region of diversion for 90% of the cases. The impact of fuel burnup is also examined, with results indicating that models trained using fast neutron flux features and their gradients exhibit more stable performance and reduced sensitivity to burnup variations.

1. Introduction

One of the main objectives of nuclear safeguards is material accountancy and control (IAEA, 2022); for that reason, Spent Nuclear Fuel (SNF) assemblies discharged from nuclear power plants are usually subjected to regular inspections to verify their content. The general strategy of these inspections relies on acquiring measurements of observable quantities, such as neutron flux, gamma dose rates, or Cherenkov light (Branger et al., 2020; Rinard and Bosler, 1988; Mayorov et al., 2017), and determining whether the outcome of the measurements is consistent with the declared configuration of the assemblies or not.

Recent research efforts have shown the potential of machine learning to enhance the processing of measured data in nuclear systems and extract more details of their configuration. For example, machine learning algorithms were used to quantify the percentage of replaced fuel pins in SNF assemblies (Rossa et al., 2020, 2018), to predict characteristic parameters of SNF such as Burn-Up (BU), Initial Enrichment (IE) and Cooling Time (CT) (Grape et al., 2020), to detect and localize missing radioactive sources within a small grid (Durbin and Lintereur, 2020), to detect anomalies in the actinide inventories for a SNF reprocessing facility (Shoman et al., 2020), and to identify and localize perturbations in nuclear reactor cores from neutron flux measurements (Pázsit et al., 1996; Durrant et al., 2021; Tasakos et al., October,2021).

In previous work (Al-dbissi et al., 2023, 2024), Machine Learning (ML) models based on a feed-forward fully-connected Artificial Neural Network (ANN) were investigated for processing simulated measurements of the thermal neutron flux from 17x17 PWR SNF assemblies to identify partial defects and reconstruct the specific patterns of missing fuel pins. ML models have the potential to reduce the amount of expert judgment and improve the accuracy of the results, thereby facilitating the decision making process of the safeguards inspectors.

Building upon the previous work, the present study explores three aspects that provide novel insights. First, a different network architecture, a Convolutional Neural Network (CNN), is explored for improving the identification of missing fuel pin patterns within SNF assemblies. CNNs are tailored for problems in which the input data have a grid-like structure (Ketkar and Moolayil, 2021), such as a 17 × 17 PWR fuel assembly, and are designed to preserve and exploit spatial relationships in the data. Unlike feed-forward fully-connected neural networks, CNNs maintain the spatial structure of the input through convolutional layers that apply filters over local regions.

Second, given the practical limitation that training datasets cannot include all possible partial defect configurations, exact pin-level

* Corresponding author.
E-mail address: moad.al-dbissi@chalmers.se (M. Al-dbissi).

https://doi.org/10.1016/j.anucene.2026.112465
Received 23 February 2026; Received in revised form 8 May 2026; Accepted 13 May 2026
Available online 20 May 2026

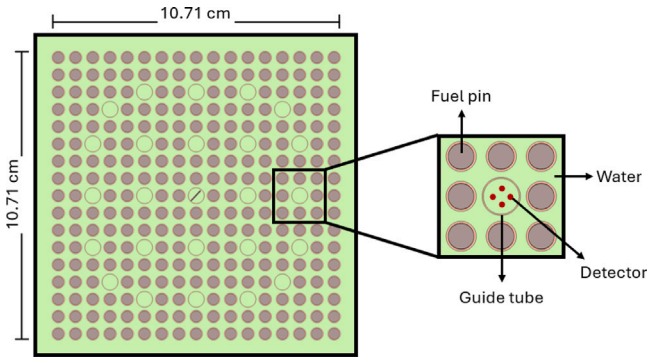

Fig. 1. A 2D model of a 17x17 PWR fuel assembly.

localization may not always be achievable. In such cases, predicted missing pins may still appear in the vicinity of their true locations. To address this, a proximity-based classification criterion is investigated, allowing for a more flexible interpretation of the output and enabling identification of the regions within the assembly where partial defects are present.

Third, since SNF assemblies encountered in safeguards applications span a wide range of burnup conditions, training a separate model for every possible burnup is impractical. This work therefore examines the robustness of the CNN when trained on a limited set of burnup values and tested under different burnup conditions, providing insight into the sensitivity of the network to variations in fuel irradiation history.

The paper is structured as follows. Section 2 introduces the problem under investigation. Section 3 describes the CNN model and the training dataset. The performance of the CNN is evaluated and compared with an upgraded version of the feed-forward fully-connected ANN in Section 4. A proximity-based classification criterion is also introduced in Section 4. The impact of fuel burnup on the performance of the CNN is studied in Section 5. Finally, the conclusions are drawn in Section 6.

2. Problem setup

A standard 17 × 17 PWR fuel assembly model is considered in this study. The assembly consists of 264 fuel pins enclosed in Zircaloy cladding (0.475 cm in radius), 24 empty guide tubes (0.63 cm in radius), and a central instrumentation channel, see Fig. 1.

It is assumed that both thermal and fast neutron fluxes, as well as their gradients, can be measured at the empty guide tube locations. The gradient of the neutron flux provides additional information on the spatial variation and directional behavior of the neutron flux. Since only the radial distribution of fuel pins in the horizontal x-y plane is relevant for the present application, the problem is treated as two-dimensional. At a fixed axial elevation z, the neutron flux and its gradient can be estimated by measuring the flux at four positions within each empty guide tube. The scalar flux at a given location is taken as the average of the four measurements, while the two Cartesian x and y components of the gradient are derived from differences between detector pairs aligned along the corresponding axes. Such measurements could be achieved by inserting a bundle of four small detectors, e.g., fiber-coupled scintillators as described in Watanabe et al. (2020), Vitullo et al. (2020), into each guide tube. In line with our previous work (Al-dbissi et al., 2022), a Monte-Carlo model of the gradient detector arrangement, shown in the inset of Fig. 1, is used in this work.

In earlier work (Al-dbissi et al., 2023, 2024), flux and gradient measurements were considered at 25 locations by including the central instrumentation channel. In practice, however, this channel is inaccessible due to the presence of the top grid plate. Consequently, only the 24 empty guide tube positions are considered in the present study.

Using the measured neutron fluxes and their gradients as input, a Convolutional Neural Network (CNN) is then employed to reconstruct the fuel pin configuration of the assembly and to identify the patterns associated with missing or replaced fuel pins under partial defect scenarios.

3. CNN architecture and training data

3.1. CNN architecture

A CNN inspired by the Inception architecture (Szegedy et al., 2015) is built using the Tensorflow (Abadi et al., 2015) and Keras (Chollet et al., 2015) open source software libraries. It consists of an input layer, an initial hidden convolutional layer, two Inception modules, a final hidden convolutional layer, and an output layer, see Fig. 2. In each Inception module, features of the problem are extracted in four parallel branches, three of them apply convolutional layers that scan the assembly according to different filter sizes (namely, 4 × 4, 5 × 5, and 6 × 6) while the fourth applies a MaxPooling layer that reduces the dimensions of the input followed by a 1 × 1 convolutional layer. The outputs from the four branches are then concatenated. The use of multiple filters in each module diversifies the spatial scales at which features are captured, thereby improving the extraction process.

All convolutional layers in the network use padding of the type 'same' and a ReLU activation function, except for the output layer, which uses a Softmax activation function. The network is compiled with an Adam optimizer and a weighted categorical cross-entropy loss function and fitted with 250 epochs and a batch size of 3. The number of epochs and the batch size are chosen based on a grid search process.

3.2. Dataset for training and testing

The CNN is trained and evaluated using a dataset of Monte-Carlo simulated 17 × 17 PWR spent nuclear fuel assemblies with an initial enrichment of 3.5 w%, a final burnup of 60 MWd/kgU, and a cooling time of 5 years. The dataset is generated with the neutron transport code SERPENT (Leppänen, 2015) with simulations performed using 5×10^9 neutron histories to balance computational cost and statistical accuracy, resulting in an average relative uncertainty of approximately 0.1% in the tallied neutron flux. The dataset consists of one intact configuration and 109 diverted cases, in which selected fuel pins are replaced with stainless-steel dummy pins. Stainless-steel is commonly used to represent non-fissile dummy fuel pins in safeguards applications due to its moderate neutron interaction properties, its similarity to structural materials already present in the assembly, and for being economic and easy to manufacture. Each diverted case in the dataset corresponds to a unique diversion pattern, with the number of replaced fuel pins ranging from 4 to 180.

For each assembly, the thermal and fast neutron fluxes, as well as the x and y components of their respective gradients, are tallied at the 24 empty guide tube locations and used as input. In return, the network provides a 17 × 17 output grid representing the fuel assembly, where each of the 264 fuel pin positions is assigned a binary label: 0 for an intact pin and 1 for a replaced pin. The central instrumentation channel and guide tube positions in the output grid are assigned a label of 2 and are masked so that they are excluded from the prediction process.

The dataset exhibits a significant class imbalance, with intact fuel pins being approximately three times more common than the replaced ones. Such imbalance can lead to a bias in the predictions toward the majority class, making it more difficult for the CNN to correctly identify replaced pins. To mitigate this, a weighted categorical cross-entropy loss function is used. By assigning higher weights to the minority class, the loss function encourages the network to place greater emphasis on learning features associated with replaced pins, thereby reducing the impact of the imbalance during training.

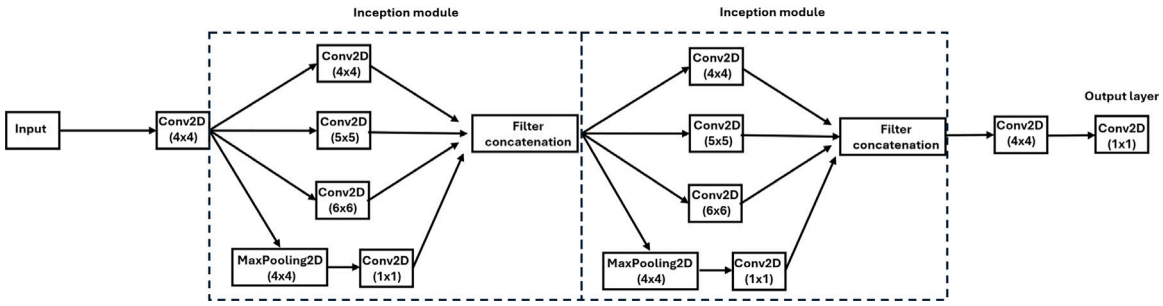

Fig. 2. A schematic of the architecture of the inception-based CNN.

Table 1
Confusion matrix.

		True label	
		Intact (0)	Replaced (1)
Predicted label	Intact (0)	True negative	False negative
	Replaced (1)	False positive	True positive

Given the limited size of the dataset, a Leave One Out Cross-Validation (LOOCV) scheme is adopted. In this approach, the CNN is trained and tested 110 times, with each case used once as the test sample while the remaining 109 cases form the training set. This strategy maximizes the use of the available data and avoids potential variability associated with a random train–test split.

4. Results

4.1. Overall performance

At the end of the LOOCV process, the predictions from the 110 iterations are concatenated and categorized according to a confusion matrix, see Table 1. The performance is assessed by the global precision, recall, and F1 score calculated across all the predictions.

The value of the recall represents the accuracy of the network in predicting a certain label. For label 1 (replaced pins), the recall is defined as the number of correctly predicted replaced pins (true positives) divided by the total number of replaced pins in the dataset (true positives + false negatives). While for label 0 (intact pins), the recall is defined as the number of correctly predicted intact pins (true negatives) divided by the total number of intact pins (true negatives + false positives).

The value of the precision represents the quality of the predictions for a certain label. For label 1, the precision is defined as the number of correctly predicted replaced pins (true positives) divided by the total number of pins predicted as replaced (true positives + false positives). While for label 0, the precision is defined as the number of correctly predicted intact pins (true negatives) divided by the total number of pins predicted as intact (true negatives + false negatives).

To ensure an accurate reconstruction of the fuel pin configuration in a SNF assembly, both recall and precision are critical; therefore, the F1 score, defined as the harmonic mean of recall and precision for a given label, is evaluated and used as the main performance metric.

Values of the F1 score resulting from the training of the CNN using different combinations of the available input features are listed in Table 2. The results show that the network performs generally better in identifying the intact fuel pins (label 0) compared to the replaced ones (label 1). A clear improvement is noticed in the network's ability to classify both intact pins and replaced pins as additional input features are used in the training. Using only the thermal or fast neutron flux results in a similar performance, but combining both increases the F1 score for both labels. Noticeable gains are also observed when using the scalar flux alongside the x and y components of its gradient, indicating that local spatial variations of the flux can provide important information for distinguishing pin configurations and identifying replaced fuel pins. Using all available features (both fluxes along with their associated gradients) leads to the best overall performance.

Table 2
Values of the F1 score obtained from training the CNN using different combinations of input features.

		Input features					
		T	F	T, F	T, $G_{Tx,y}$	F, $G_{Fx,y}$	T, $G_{Tx,y}$, F, $G_{Fx,y}$
F1 score	Label 0	0.82	0.84	0.88	0.89	0.90	0.91
	Label 1	0.60	0.60	0.67	0.70	0.70	0.73

T: Thermal neutron flux
F: Fast neutron flux
$G_{Tx,y}$: The x and y components of the gradient of the thermal neutron flux
$G_{Fx,y}$: The x and y components of the gradient of the fast neutron flux.

Table 3
Values of the F1 score obtained from training the feed-forward, fully-connected ANN using different combinations of input features.

		Input features					
		T	F	T, F	T, $G_{Tx,y}$	F, $G_{Fx,y}$	T, $G_{Tx,y}$, F, $G_{Fx,y}$
F1 score	Label 0	0.75	0.80	0.84	0.82	0.86	0.89
	Label 1	0.56	0.59	0.64	0.63	0.66	0.72

T: Thermal neutron flux
F: Fast neutron flux
$G_{Tx,y}$: The x and y components of the gradient of the thermal neutron flux
$G_{Fx,y}$: The x and y components of the gradient of the fast neutron flux

4.2. Comparison with feed-forward fully-connected ANN

The performance of the CNN is compared with that of the feed-forward, fully connected ANN that was previously studied in Al-dbissi et al. (2024). To ensure a fair comparison, the ANN was slightly modified and trained under the same conditions as the CNN. Specifically, the standard binary cross-entropy loss was replaced with a weighted loss function consistent with that used for the CNN, and a Leave One Out Cross-Validation (LOOCV) scheme was adopted in place of the original 10-fold cross-validation. The ANN was trained using the same dataset and the same combinations of input features, as summarized in Table 3.

The results indicate that, similar to the CNN, the ANN performs better at identifying intact fuel pins (Label 0) than replaced pins (Label 1). Incorporating information related to the fast neutron flux and its gradient consistently improves the ANN's performance, with the best overall results obtained when all available features are used. A direct comparison with the results in Table 2 shows that the CNN achieves slightly better performance than the ANN across all feature combinations considered.

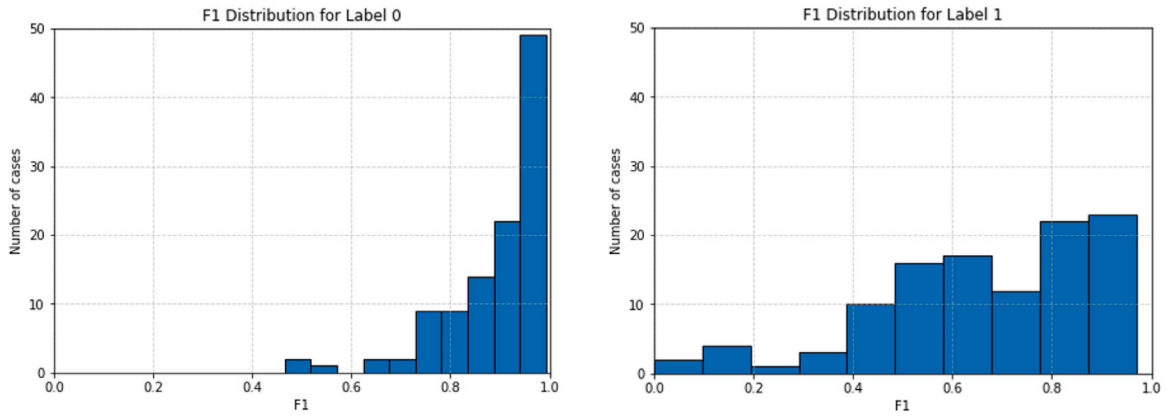

Fig. 3. F1 score distribution for label 0 (left), and the F1 score distribution for label 1 (right).

4.3. Performance over individual samples

The global F1 score values reported in Table 2 provide a useful overview of the CNN's performance, however, they can overlook variability between individual cases. Considering the CNN model that shows the best global results, i.e., the one trained with both thermal and fast fluxes along with their associated gradients, the F1 score is evaluated separately for each individual scenario (per-sample) in the dataset.

The distributions of the F1 score for labels 0 and 1 across all 110 scenarios are illustrated in Fig. 3. For Label 0 (intact pins), a majority of 92 cases have a F1 score higher than 0.80 with only one case falling below 0.50, indicating that the CNN correctly identifies intact pins with high consistency across most scenarios. For Label 1 (replaced pins), the F1 distribution shows that the model has a more widely spread performance across the dataset. While 41 cases have a F1 score higher than 0.80 there are also 21 cases with values below 0.50. This suggests that although the model performs well for many configurations, the detection of replaced pins remains more challenging.

The six worst-performing cases, for which the F1 score for Label 1 is below 0.20, are presented in Fig. 4. Case A corresponds to the only intact configuration in the dataset. Since the training set includes one intact case and 109 diverted cases, the model exhibits a tendency to predict the presence of a diversion pattern even when the assembly is fully intact. This suggests that, for the specific objective of reconstructing diversion patterns, testing only on diverted assemblies may be more appropriate. A two-step classification process could also be introduced. In the first step, a binary decision would be made to determine whether an assembly is intact or not. For this, a separate model should be developed and trained on a balanced dataset containing both intact and diverted configurations. Assemblies identified as defective could then be analyzed by the method developed in this study in order to reconstruct the corresponding diversion pattern.

The remaining cases in Fig. 4 represent diversion scenarios for which the CNN performs poorly, either failing to detect missing pins or placing them in incorrect locations. Given the small size of the training dataset and the limited variety of patterns it contains, reduced performance on certain cases is expected.

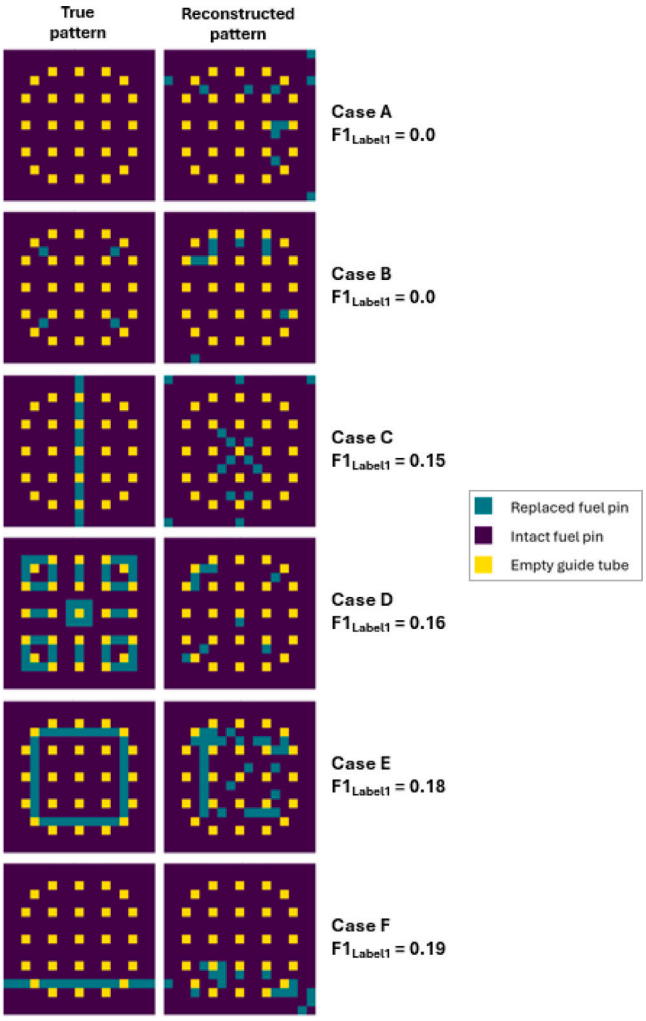

Fig. 4. Cases with the lowest F1 score for label 1 (<0.2).

4.4. Proximity-based criterion

For the scenarios shown in Fig. 4, with the exception of case A, the reconstructed patterns provide indications of the regions where the fuel pins are replaced. In other words, many false positives occur nearby false negatives, suggesting that the model identifies the correct area within the assembly. Based on this observation, a proximity-based classification criterion is introduced so that an actual missing fuel pin is considered correctly identified if it falls within a 3x3 region defined by the predicted missing fuel pin and its first and second neighbors, where the region is constrained by the assembly boundaries.

The distribution of the F1 score when applying the proximity criterion for label 1, is shown in Fig. 5. The values are substantially improved across the 110 cases. With 99 cases having F1 scores above 0.80, out of which 23 cases having a score of 1.0. Only two cases have a F1 score below 0.50, one of which is the intact configuration. These

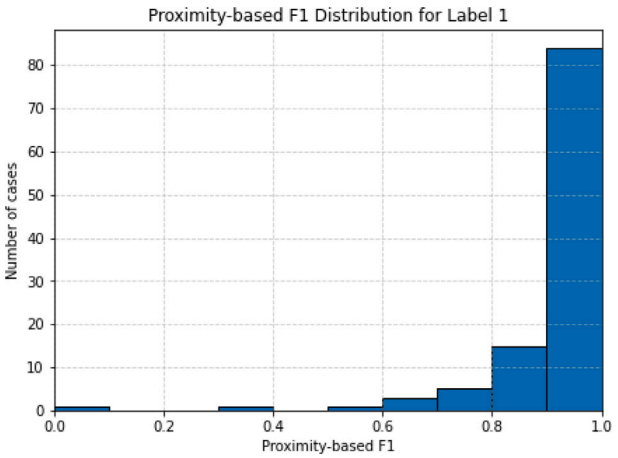

Fig. 5. The distribution of the proximity-based F1 score for label 1.

results demonstrate that the CNN model consistently recognizes the correct region where the pattern of missing fuel pins occurs.

5. Impact of fuel burnup

SNF assemblies can span a wide range of burnup levels, reflecting differences in operating history and irradiation conditions. The burnup level influences the fuel isotopic composition and, consequently, alters both the magnitude and spatial distribution of the neutron flux within the assembly. Since in the current application, the CNN relies on neutron signatures to reconstruct the fuel pin configuration, variations in burnup may affect the network's ability to identify replaced fuel pins correctly.

During operation, fuel assemblies can be reshuffled in the core to promote more uniform fuel utilization. Hence, in our study we assume that the burnup over an assembly is uniform, although it might not always be the case.

5.1. Evaluation on unseen burnup levels

To assess the CNN's ability to generalize across different burnup levels, two additional versions of the dataset described in Section 3 were generated with final burnup levels of 20 MWd/kgU and 40 MWd/kgU, alongside the original dataset at 60 MWd/kgU. All three datasets contain the same diversion scenarios and share the same initial enrichment and cooling time. A three-fold cross-validation approach is then applied, where in each fold the CNN is trained on two of the available datasets and tested on the third. This means that the network is evaluated on diversion scenarios it has already seen during training, but under different burnup conditions. The resulting F1 scores for label 1 using different combinations of input features are shown in Fig. 6.

The results show that for the following models: the one trained using the thermal neutron flux (T), the one trained using the thermal and fast neutron flux (T, F), the one trained using the thermal neutron flux and its gradient (T, $G_{Tx,y}$), and the one trained using all available features (T, $G_{Tx,y}$, F, $G_{Fx,y}$), the network performs best when tested on the 40 MWd/kgU dataset and worst when tested on the 20 MWd/kgU dataset. This suggests that the network's performance is strongly influenced by burnup when thermal neutron flux and its gradient are included as input features. In contrast, for the models trained using only features related to the fast neutron flux (F) and (F, $G_{Fx,y}$), the performance remains relatively consistent across all burnup levels. The model trained using the fast neutron flux and its gradient (F, $G_{Fx,y}$) provides the best overall results with F1 scores ranging between 0.97 and 0.99. This indicates that the localization of missing fuel pins using fast neutron flux and its gradient is less sensitive to burnup variations.

Table 4
Comparison Between Standard and Proximity-Based F1 Scores for the (T, $G_{Tx,y}$, F, $G_{Fx,y}$) model.

Test case	No. of replaced pins	Burnup (MWd/kgU)					
		20		40		60	
		Original	Proximity	Original	Proximity	Original	Proximity
1	12	0.31	0.67	0.32	0.47	0.33	0.50
2	24	0.49	0.68	0.69	1.00	0.81	0.98
3	48	0.86	1.00	0.99	1.00	0.98	0.98
4	63	0.94	1.00	0.93	1.00	0.99	1.00
5	80	0.64	0.90	0.45	0.94	0.49	0.94
6	128	0.77	0.99	0.88	0.99	0.93	1.00

5.2. Evaluation on unseen scenarios and burnup levels

The CNN is further evaluated under conditions where both the diversion scenarios and the burnup levels are unseen during training. For the purpose, each of the three burnup datasets is divided into a training set of 104 cases and a test set of 6 cases. The six test scenarios are selected to include configurations with both small and large numbers of replaced fuel pins, as well as symmetric and non-symmetric diversion patterns. A three-fold cross-validation is then applied, in which the CNN is trained using the training sets from two burnup levels and tested on the test set from the third. The resulting F1 scores for Label 1 are shown in Fig. 7.

Models that rely mainly on thermal neutron flux and its gradient, namely (T) and (T, $G_{Tx,y}$) in the figure, show similar F1 scores to the previous analysis, achieving best performance when tested at a burnup of 40 MWd/kgU. Meanwhile, the remaining models trained using the fast neutron flux and its gradients, i.e., (F) and (F, $G_{Fx,y}$), or a mix of fast and thermal features, (T, F) and (T, $G_{Tx,y}$, F, $G_{Fx,y}$), exhibit lower F1 scores compared to the previous analysis, however, their performance remains more consistent across different burnup levels. The (T, $G_{Tx,y}$, F, $G_{Fx,y}$) and (F, $G_{Fx,y}$) models have the best overall results, with F1 scores ranging between 0.72 and 0.80.

The proximity-based criterion introduced in Section 4.4 is applied to the test cases of the (T, $G_{Tx,y}$, F, $G_{Fx,y}$) model, and its impact on the individual F1 scores across the different burnup levels is summarized in Table 4. The results show a consistent improvement in F1 score for all cases and across all three burnup levels. This suggests that the network reliably identifies the correct region of the replaced fuel pins even across different levels of burnup.

6. Conclusions

Convolutional Neural Networks (CNNs) were investigated for the task of identifying and localizing replaced fuel pins in 17 × 17 PWR spent nuclear fuel (SNF) assemblies under partial defect scenarios. A CNN based on the Inception architecture was developed and trained using a Monte-Carlo simulated dataset consisting of one intact configuration and 109 diversion scenarios in which selected fuel pins are replaced with dummy pins. Thermal neutron flux, fast neutron flux, and the spatial gradients of both fluxes were tallied at the 24 empty guide tube positions and used as input features. The CNN produces a 17 × 17 output grid representing the fuel assembly, where each fuel pin is classified as intact (0) or replaced (1).

Different combinations of input features were evaluated during training. The results show that all models perform better at identifying intact fuel pins compared to the replaced ones. A noticeable improvement in F1 score was gained when the gradient components were included together with the scalar flux. The best performance was achieved when all available input features were used to train the network.

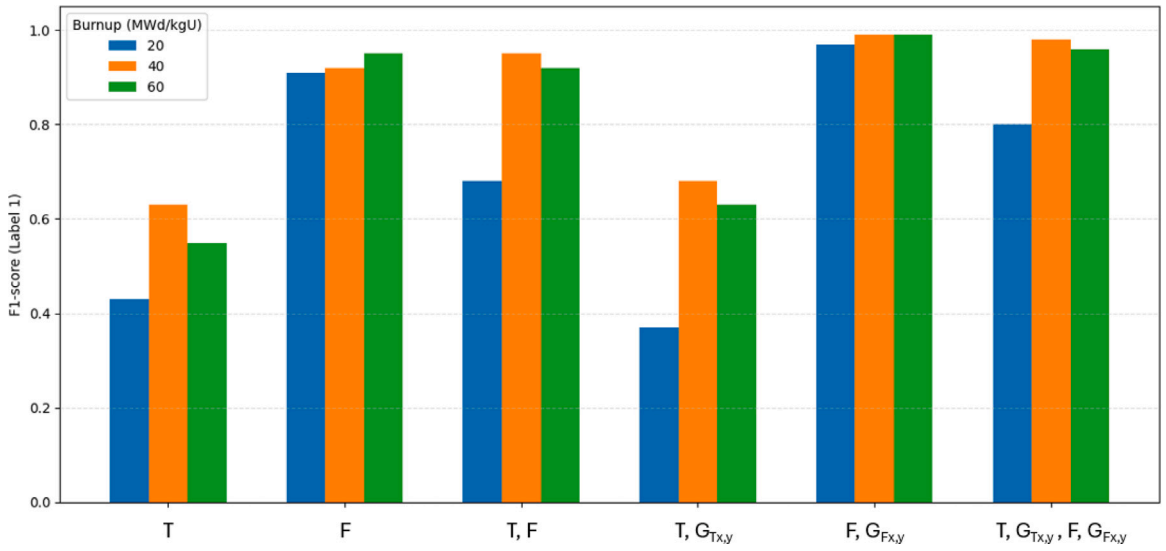

Fig. 6. F1 score for label 1 at three different burnup levels using a 3-fold cross-validation.

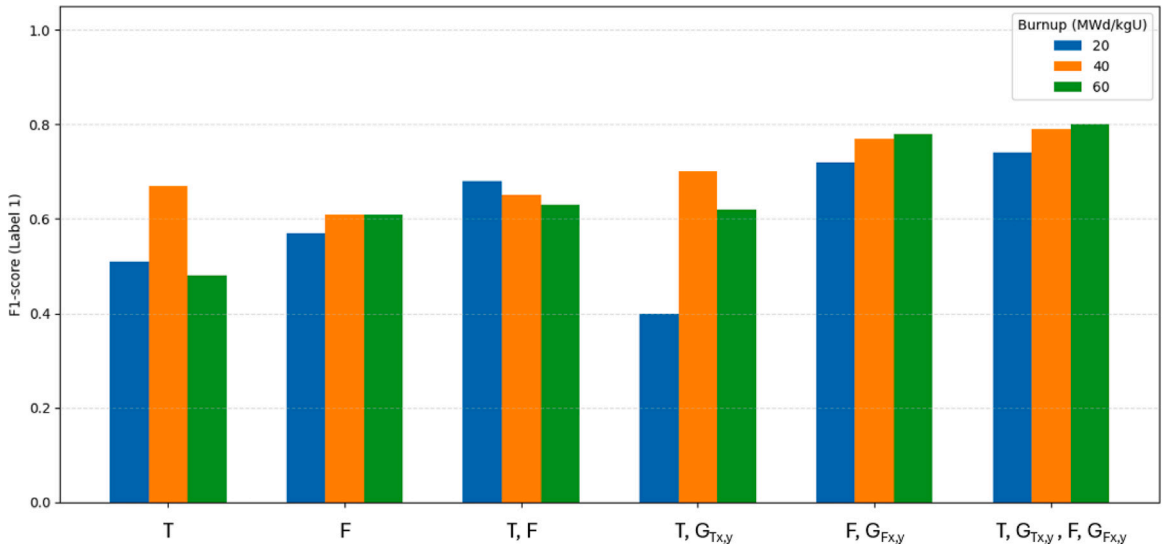

Fig. 7. F1 score for label 1 at three different burnup levels using an unseen test-set.

The CNN's performance was also compared with a feed-forward, fully connected ANN. For all feature combinations considered, the CNN achieved better results than the ANN.

Analysis of the F1 score distribution across individual test cases revealed a noticeable variability in the network's ability to localize replaced fuel pins. This behavior can be attributed to the small size and limited diversity of the training dataset. However, when a proximity-based classification criterion was applied, the network successfully identified the region of diversion within the assembly even when the exact pin configurations were not recovered.

The performance of the CNN was also examined using three datasets generated at burnup levels of 20, 40, and 60 MWd/kgU. The results show that models relying on thermal neutron flux and its gradient were more sensitive to burnup variations, whereas models trained with fast neutron flux and its gradient exhibited better and more consistent performance. The best results were obtained when using either all available input features or the fast flux together with its gradient. For those specific models, the analysis suggests that exposure to a diverse range of partial-defect patterns during training appears to have a greater influence on model performance than exposure to fuel assemblies uniformly burnt at different levels.

Future work will focus on expanding the training dataset to include a wider variety of diversion scenarios in order to allow the CNN to learn more detailed spatial features and further improve its ability to accurately localize replaced fuel pins.

This study is based on the assumption that the removed fuel pins are replaced with stainless steel dummy pins. Although the use of stainless steel is common in this context, future work could also investigate the adaptability of the CNN to scenarios where different materials are used for the replaced fuel pins.

CRediT authorship contribution statement

Moad Al-dbissi: conceptualization, Methodology, Software, Formal analysis, Writing – original draft, Writing – review & editing, visualization. **Imre Pázsit:** conceptualization, Methodology, Writing – review & editing, supervision. **Paolo Vinai:** conceptualization, Methodology, Formal analysis, Writing- review & editing, Supervision.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT (OpenAI) to assist with language refinement, including sentence rephrasing and grammatical corrections. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

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.

Acknowledgment

The project was financially supported by the Swedish Radiation Safety Authority, Sweden under agreement SSM2023-4389.

Data availability

Data will be made available on request.

References

Abadi, M., et al., 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org. URL https://www.tensorflow.org/.

Al-dbissi, M., Pázsit, I., Rossa, R., Borella, A., Vinai, P., 2024. On the use of neutron flux gradient with ANNs for the detection of diverted spent nuclear fuel. Ann. Nucl. Energy 204, 110536. http://dx.doi.org/10.1016/j.anucene.2024.110536.

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.

Al-dbissi, M., Vinai, P., Borella, A., Rossa, R., Pázsit, I., 2022. Conceptual design and initial evaluation of a neutron flux gradient detector. Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip. 1026, 166030. http://dx.doi.org/10.1016/j.nima.2021.166030.

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.

Durrant, A.M., Leontidis, G., Kollias, S., Torres, A., Montalvo, C., Mylonakis, A., Demaziere, C., Vinai, P., 2021. Detection and localisation of multiple in-core perturbations with neutron noise-based self-supervised domain adaptation. In: Proc. Int. Conf. MC 2021. American Nuclear Society.

Grape, S., Branger, E., Elter, Z., Pöder Balkeståhl, L., 2020. Determination of spent nuclear fuel parameters using modelled signatures from non-destructive assay and random forest regression. Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip. 969, 163979. http://dx.doi.org/10.1016/j.nima.2020.163979.

IAEA, 2022. International atomic energy agency safeguards glossary. URL https://www.iaea.org/publications/15176/iaea-safeguards-glossary.

Ketkar, N., Moolayil, J., 2021. Convolutional neural networks. In: Deep Learning with Python: Learn Best Practices of Deep Learning Models with PyTorch. A Press, Berkeley, CA, pp. 197–242. http://dx.doi.org/10.1007/978-1-4842-5364-9_6.

Leppänen, J., 2015. The Serpent Monte Carlo code: Status, development and applications in 2013. Ann. Nucl. Energy 82, 142–150.

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., Garis, N.S., Glöckler, O., 1996. On the Neutron Noise Diagnostics of Pressurized Water Reactor Control Rod Vibrations —IV: Application of Neural Networks. Nucl. Sci. Eng. 124 (1), 167–177. http://dx.doi.org/10.13182/NSE96-A24232.

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, URL https://www.sciencedirect.com/science/article/pii/S0306454920303789.

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. Proc. INMM-59 Annual Meeting.

Shoman, N., Cipiti, B., Grimes, T., Wilson, B., 2020. Advances in machine learning for safeguarding a PUREX reprocessing facility. In: INMM 61st Annual Meeting.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition. CVPR, http://dx.doi.org/10.1109/CVPR.2015.7298594.

Tasakos, T., Ioannou, G., Verma, V., Alexandridis, G., Dokhane, A., Stafylopatis, A., October,2021. Deep learning-based anomaly detection in nuclear reactor cores. In: Proceedings of the International Conference on Mathematics and Computational Methods Applied To Nuclear Science and Engineering (M&C 2021).

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.