Detection of partial defects in spent fuel assemblies with the help of machine learning

Riccardo Rossa, Alessandro Borella

Belgian nuclear research centre SCK CEN. Boeretang 200, B-2400 Mol, Belgium

Abstract

The detection of partial defects, i.e. missing fuel pins, in spent fuel assemblies is a challenging tasks during a safeguards inspection. Spent fuel assemblies are composed of a large number of fuel pins (e.g. 264 pins in a PWR 17x17 assembly geometry) and a proliferator can develop a practically infinite number of diversion scenarios to replace fuel pins with dummy pins.

Safeguards inspections on spent fuel rely on non-destructive assays (NDA) to verify the presence of spent fuel and to detect fuel pins diversions. One of the NDA that has been developed specifically for the detection of fuel pins diversions is the Partial Defect Tester (PDET). The PDET instrument combines several detectors to measure the neutron and gamma-ray fluxes in several positions across the fuel assembly cross-section.

Given the almost infinite number of possible diversion scenarios in a fuel assembly and the multivariate information resulting from a PDET measurement of a fuel assembly, machine learning (ML) has been chosen to tackle the data analysis. First, a large dataset has been developed with the results of Monte Carlo simulations representing measurement of complete fuel assemblies as well as assemblies with dummy pins. Then, a class label has been assigned to each dataset observation according to the percentage of replaced fuel pins. The problem of partial defect in a fuel assembly has been treated as a classification problem, using as input features the detector responses from the PDET instrument and as output values the class label of the observation. The k-nearest neighbor (kNN) approach has been applied to develop the ML models.

The results show that the use of multiple detectors in different locations within and outside the fuel assembly is helpful in the classification problem. By tuning the hyper-parameters the kNN models were able to reach a >99% classification accuracy. A detailed discussion highlights the observations that are misclassified, with particular attention to the missed detections where assemblies with missing fuel pins are classified as complete assemblies by the ML models.

1 Introduction

Nuclear safeguards are a set of technical measures to ensure that nuclear materials and technologies are used for peaceful activities and not for the development of nuclear weapons (IAEA, 1972). Safeguards inspections are carried out by the International Atomic Energy Agency (IAEA) according to international treaties such as the Non-Proliferation Treaty and the Additional Protocol (IAEA, 1970), (IAEA, 1997).

Among the materials placed under safeguards spent fuel accounts for the vast majority of the inventory due to the presence of irradiated ^{235}U and Pu. According to the IAEA Safeguards Implementation Report (IAEA, 2020) at the end of 2019 more than 77% of significant quantities under safeguards were in form of irradiated plutonium. Moreover, due to the discharge of spent fuel from operating reactors, the number of significant quantities of irradiated Pu increased by roughly 25% compared to 2009. A significant quantity is defined by the IAEA as “the approximate amount of nuclear material for which the possibility of manufacturing a nuclear explosive device cannot be excluded” (IAEA, 2002).

Spent fuel assemblies are composed of a large number of fuel pins (e.g. 264 pins in a PWR 17x17 assembly geometry) and a proliferator can develop a practically infinite number of diversion scenarios to replace fuel pins with dummy

pins. Several non-destructive assay (NDA) instruments are already used routinely in the safeguards inspections and they rely on neutron, gamma-ray, or Cherenkov emissions from spent fuel (IAEA, 2011). Given the almost infinite number of possible diversion scenarios in a fuel assembly and the multivariate information resulting from NDA measurement of a fuel assembly, machine learning is being increasingly chosen to tackle the data analysis (Rossa, 2019), (Rossa, 2020-b), (Grape, 2020), (Bachmann, 2021).

This paper describes in Section 2 the methodology applied for this study giving information on the available dataset, NDA technique, and machine learning approach. The results of the analysis are presented in Section 3 in terms of classification accuracy and discussing misclassification cases. The conclusions from the analysis are summarized in Section 4 together with the outlook for the next steps in the research.

2 Methodology

2.1 Dataset of spent fuel assemblies

A dataset containing the simulated detector responses from the considered NDA technique was used for the analysis. This section offers an overview of the dataset, whereas the approach for the determination of the detector responses is described in Section 2.2.

The first part of the dataset contains so-called complete fuel assemblies, where all fuel pins are present and contribute to the neutron and gamma-ray emission of the spent fuel. This section totaled 196 simulations and was developed to account for the changes in radiation emission due to the different irradiation history (i.e. initial enrichment, discharge burnup, and cooling time) of the fuel in the reactor. For this reason the fuel in this part of the dataset had the following parameters:

- Initial enrichment: 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0%;
- Discharge burnup: 5, 10, 15, 20, 30, 40, 60 GWd/tHM;
- Cooling time: 1, 5, 10, 50 years.

The second part of the dataset contains so-called defect fuel assemblies, where some of the fuel pins have been replaced by dummies made of stainless steel. A total of 107 diversion scenarios (i.e. fuel replacement by dummies) were considered, and for each diversion scenario a total of 9 fuel compositions were considered, all with cooling time of 5 years, and combining the following parameters:

- Initial enrichment: 2.0, 3.5, 5.0%;
- Discharge burnup: 10, 30, 60 GWd/tHM.

Both symmetric and asymmetric patterns were taken into account in the dataset, and some of the diversion scenarios are shown in Figure 1.

2.2 Detector responses from the Partial Defect Tester

The dataset presented in Section 2.1 contains the detector responses from the Partial Defect Tester (PDET) (Sitaraman, 2007) , (Ham, 2009), (Ham, 2015). The PDET combines a set of neutron and gamma-ray detectors placed simultaneously in different areas of the fuel assembly to obtain the spatial distribution of the radiation emitted.

The first step in the computation of the PDET detector responses was the fuel depletion calculation. This was carried out with the ORIGEN-ARP code (Gauld, 2011) to simulate the fuel irradiation in a nuclear reactor. The geometry of the PWR 17x17 fuel assembly was taken as reference and the results were included in the SCK CEN spent fuel library (Rossa, 2020-a).

As second step the fuel isotopic composition and radiation emission were included in a Monte Carlo model of the PDET detector. The MCNPX code was used for this step and Figure 2 shows the example of the geometry for a defect assembly. The fuel pins are depicted in black whereas the dummy pins are shown in grey. The detectors positions

included both the central guide tube locations (highlighted in red) as well as locations outside the fuel assembly (highlighted in green).

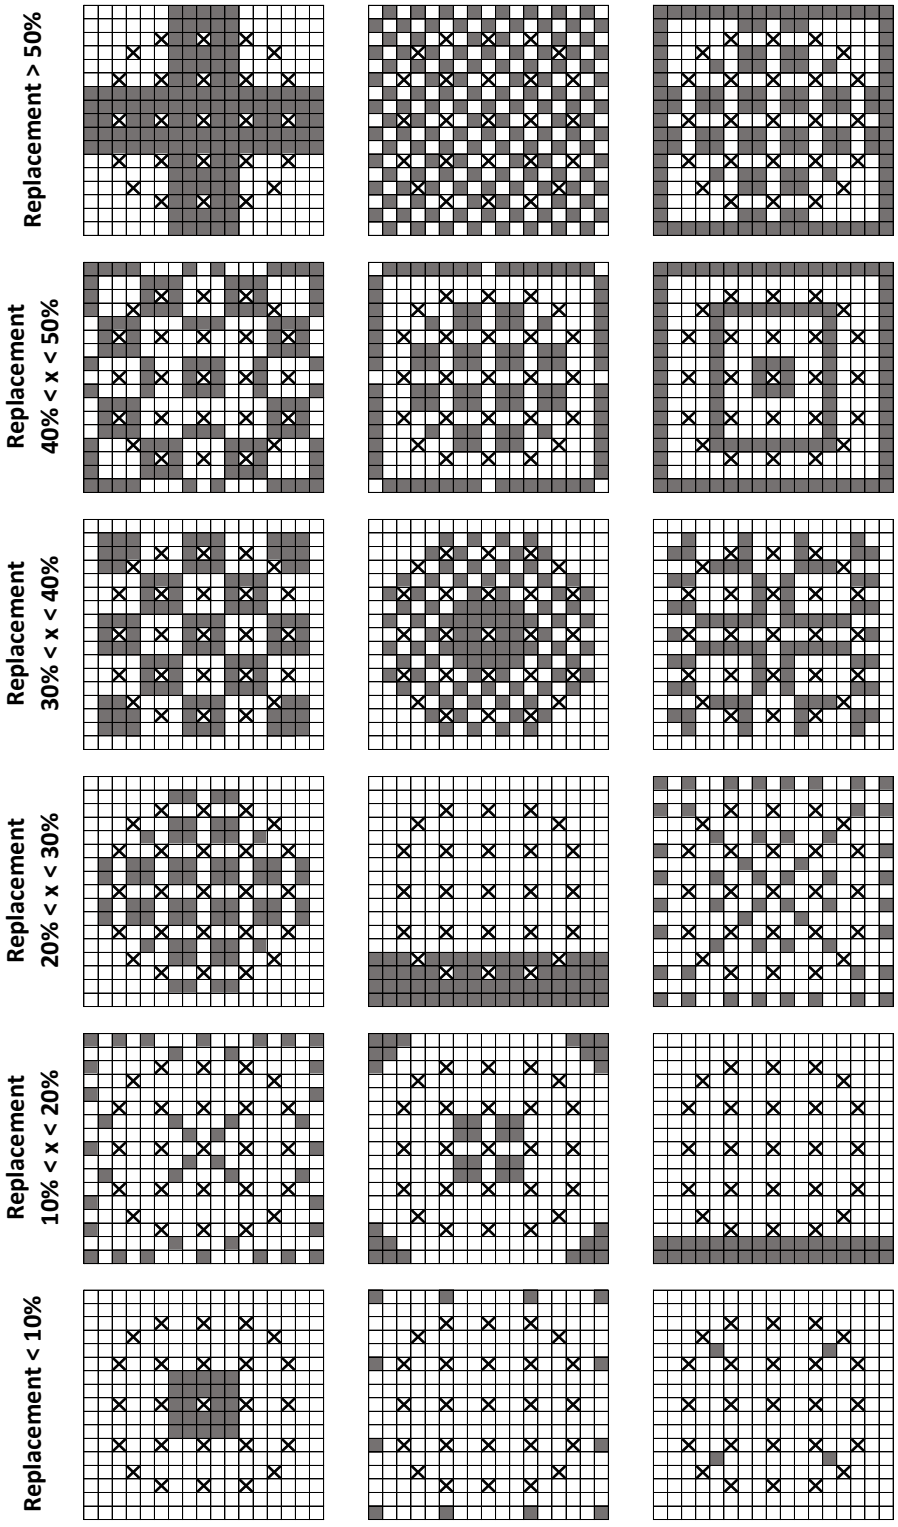

Figure 1: examples of diversion scenarios. The fuel pins are depicted in white, the dummy pins are depicted in grey, and the guide tube positions are marked with a cross. The percentage of replaced fuel pins is indicated for each row in the figure.

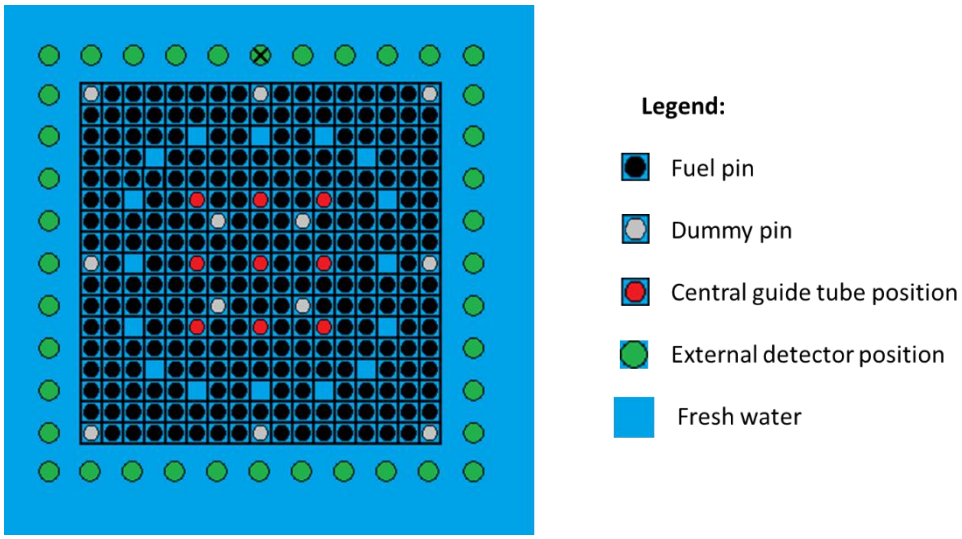

Figure 2: Monte Carlo model of the PDET instrument in case of a simulation with defect fuel.

Based on the results of previous work (Rossa, 2020-b) two detector responses were considered for PDET:

- Bare ^{238}U fission chamber (*FA*): detector sensitive mainly to fast neutrons;
- Ionization chamber (*P*): detector sensitive to gamma-rays.

The detector responses from the Monte Carlo model are so-called tallies that are by default normalized per starting source particle. Therefore, the tallies need to be multiplied by the source intensity to obtain the absolute estimate of the detector response (i.e. count rate for the fission chamber and electric current for the ionization chamber).

The neutron count rates are obtained by calculating the neutron fluence at given positions and multiplying by the neutron induced fission cross section of ^{238}U and the amount of ^{238}U in the fission chamber (Borella, 2017). The ionization chamber current is obtained from the deposited energy in the active volume of the detector with the procedure described in Rossa (2016), (Borella, 2018).

The detector responses (i.e. neutron count rate or electric current) obtained for each simulation were normalized by the value of the response calculated for the detector marked with a cross in Figure 2. Finally, the average values of the detector responses from the nine central guide tubes locations (*Cen*, red positions in Figure 2) and from the forty external locations (*Ext*, green positions in Figure 2) were then used as input features to the machine learning models. Table 1 offers the list of input features included in each machine learning model.

2.3 Machine learning approach

2.3.1 Definition of the classification problem

Machine learning can be defined as the study of computer algorithms that improve automatically through experience (Mitchell, 1997). Applications of machine learning can be found nowadays in a broad range of applications such as pattern recognition, financial fraud detection, and cancer prognosis (Bishop, 2006), (Ngai, 2011), (Kourou, 2015). Narrowing to the nuclear field the increase of machine learning uses is also observed with examples on gamma-ray spectroscopy and fuel reprocessing (Hellesen, 2017), (Orton, 2011).

Machine learning models can be divided into the two broad categories of supervised and unsupervised learning (Murphy, 2012). The former indicates problems where the input data (also called features) has an associated output while the for latter the input data does not have an associated output. In the case of supervised learning the machine learning model is first developed during a training phase, where observations with known inputs and outputs are fed to the model. In the following phase, called prediction phase, new observations with unknown outputs are fed to the

model in order to predict the corresponding output. In the field of supervised learning, distinction can be made between regression and classification problems depending on the nature of the associated output. Regression concerns problems where the output can assume a continuous value (e.g. temperature reading) whereas in classification problems the output can assume only discrete labels (e.g. yes/no).

The detection of partial defects considered in this paper is an example of classification problem. The input data are the detector responses from PDET whereas the output data is the class label of the fuel assembly. Seven class labels were defined in this problem depending on the percentage of diverted fuel pins in a fuel assembly. The input features selected for each model are shown in Table 1, whereas the number of observations for each class label is included in Table 2.

2.3.2 The k-nearest neighbors machine learning approach

The k-nearest neighbors (kNN) approach was selected for the development of the machine learning models based on the results of previous research (Rossa, 2020-b).

The kNN approach is used to classify an observation with unknown class label (also called query observation) based on the distance from k neighboring observations in the training dataset (Iggane, 2002). The kNN machine learning models were developed using the Scikit-learn package from Python (Pedregosa, 2011). The following parameters were chosen to develop the kNN models:

- Uniform weights for all observations
- Brute algorithm to calculate the distance between observations
- Minkowski metric to calculate the distance between observations

Once the parameters above have been defined the user has to select the features to be used as input to the model. The model then calculates the distance between the query observation and all observations in the training dataset. The k training observations with smaller distance to the query observations are retained, and the class label of the query observation is taken as majority of the class labels of the retained observations of the training dataset. The number of k neighboring observations was the hyper-parameter optimized during the analysis and varied between 6 and 15. The boundaries for the variation of the hyper-parameter were selected from previous research to reduce the risk of overfitting while keeping a large classification accuracy (Rossa, 2020-b).

2.3.3 Estimation of models performance

The overall accuracy of the developed models was evaluated using the 5-fold cross-validation approach (Mitchell, 1997). The observations of the complete dataset were first split randomly into 5 groups (also called folds) of equal size. Four of these folds were used as training dataset during the training phase, whereas the last fold was used as validation dataset during the prediction phase. Following an iterative process a machine learning model was developed using four of the folds as training dataset and the remaining fold as validation dataset. The iteration continued until all folds have been used as validation dataset. For each iteration the predicted class label was retained and the model accuracy was calculated as the ratio between the number of correct classifications and the total number of observations in the dataset.

The model performance was further evaluated focusing on the misclassified observations where defect fuel assemblies were identified as complete fuel assemblies. These misclassifications were judged the most critical from a safeguards point of view because they lead to missed detection of diversion. Three models with largest classification accuracy were selected, one using only features from central detector positions, one using only features from external detector positions, and one using features from both central and external detector positions. A confusion matrix was constructed for each of these models, where the observations are placed in the matrix according to their true class label and predicted class label. In general, observations that fall along the main diagonal of the confusion matrix are correct classifications, whereas observations outside the main diagonal indicate misclassifications. The further away from the diagonal, the larger the degree of misclassification.

Table 1: Overview of features used for each machine learning model.

	Machine learning model identification								
	1	2	3	4	5	6	7	8	9
Feature 1	*CenFA*	*CenP*	*CenFA*	*ExtFA*	*ExtP*	*ExtFA*	*CenFA*	*CenP*	*CenFA*
Feature 2			*CenP*			*ExtP*	*ExtFA*	*ExtP*	*CenP*
Feature 3									*ExtFA*
Feature 4									*ExtP*

Table 2: overview of observations included in each class of the dataset. The percentage of fuel diversion, number of diversion scenarios, and number of observations are reported.

Class label	Percentage of fuel diversion	Number of diversion scenarios	Number of observations
0	0%	N.A.	196
1	< 10%	19	171
2	10% <x< 20%	24	216
3	20% <x< 30%	21	189
4	30% <x< 40%	16	144
5	40% <x< 50%	16	144
6	> 50%	11	99

3 Results

3.1 Classification accuracy

The classification accuracy of the models are shown in Table 3 as a function of the number of neighbors set for the model. In general a decrease of the classification accuracy can be seen by increasing the number of neighbors set for the model, as was also shown in previous results (Rossa, 2020-b). By increasing the number of neighbors the observations are increasingly distant from the query observation and this increases the possibility of misclassification. The models using features with the ionization chambers reach in general larger classification accuracies (e.g. more than 99% for model ID 8) compared to models using features with the fission chambers (e.g. 57% for model ID 7), and this confirms that the gamma-ray detectors are more sensitive to the replacement of fuel pins (Rossa, 2020-b).

Models using only features from central detectors (i.e. model IDs 1-3) reach a maximum classification accuracy of 87% whereas models using only features from external detectors (i.e. model IDs 4-6) reach a maximum classification accuracy of 83%. However, the use of detectors outside the fuel assembly presents advantages from the practical use of the instrument during a safeguards inspections, and from the possibility to apply the same instrument to different fuel assembly geometries.

The combination of features from different detector positions (i.e. model IDs 7-9) lead to a general increase of the classification accuracy. The machine learning model combining features from the ionization chambers (model ID 8) reach a classification accuracy larger than 99% while the model combining all features (model IDs 9) have accuracy larger than 95%. The results show that the increase of the number of features does not necessarily lead to an increase of the classification accuracy, but rather the detector type is the main responsible for the classification accuracy.

Table 3: Classification accuracies of the machine learning models as a function of the number of neighbors. Models highlighted in green are further analyzed with the confusion matrix.

Nr. of neigh.	Machine learning model identification								
	1	2	3	4	5	6	7	8	9
6	37%	81%	87%	43%	81%	83%	57%	>99%	96%
7	38%	81%	86%	44%	82%	83%	58%	>99%	97%
8	39%	78%	84%	44%	80%	81%	58%	99%	94%
9	39%	77%	83%	44%	81%	81%	57%	98%	92%
10	39%	76%	80%	44%	79%	78%	56%	94%	90%
11	38%	73%	79%	44%	76%	76%	56%	89%	87%
12	38%	70%	76%	43%	73%	74%	55%	81%	83%
13	39%	67%	73%	44%	68%	71%	55%	75%	78%
14	39%	63%	70%	44%	62%	67%	54%	68%	72%
15	39%	60%	69%	44%	59%	65%	54%	65%	71%

3.2 Confusion matrixes of selected kNN models

Three models were selected from the results of Table 3 for further investigation with the confusion matrix. The first model (so-called central model) was the model with largest classification accuracy using only features from central detectors positions (i.e. model IDs 1-3), the second model (so-called external model) was the model with largest classification accuracy using only features from external detectors positions (i.e. model IDs 4-6), and the third model (so-called combination model) was the model with largest classification accuracy using features from both central and external detectors positions (i.e. model IDs 7-9).

The confusion matrix was constructed for each of the selected models and the results are shown in Figure 3. The results for the central model show that most of the misclassifications occur among neighboring class labels (e.g. observations with class label 3 misclassified as class label 2). However, there are a few cases where the misclassification is rather large (e.g. 6 cases with >50% replaced pins are classified as having <20% of replaced pins). Most critical from a safeguards perspective, there are 15 cases where defect fuel assemblies are misclassified as complete fuel assembly.

Focusing on the results of the external model, most of the misclassifications occur for diversion with up to 20% of fuel pins. Several misclassifications, as in the case of the central model, are among neighboring class labels. A total of 39 cases of defect fuel assemblies were classified as complete fuel assemblies.

The results from the combination model show that all but two cases are correctly classified, with only exceptions of cases with class label 4 being classified with class label 3. The confusion matrix thus clearly show the advantage of combining features from different detector positions.

Focusing on the misclassifications where defect assemblies are classified as complete fuel assemblies (first column of the confusion matrix), Table 4 gives an overview of the diversion scenarios that were misclassified. The list for the central model shows that most of the misclassifications occur for scenarios with 12 replaced fuel pins, but also one case of diversion of 77 pins was undetected. The list for the external model show that other diversion scenarios were misclassified in this case, with number of replaced pins ranging from 4 up to 48. For both models the diversion scenarios misclassified as complete fuel assemblies have the replaced pins mostly on the outer row of the fuel assembly (e.g. scenarios with 12 replaced pins), or in positions between the central and external detector positions (e.g. scenarios with 24 and 48 replaced pins). The reason for the misclassification will be further investigated in future work.

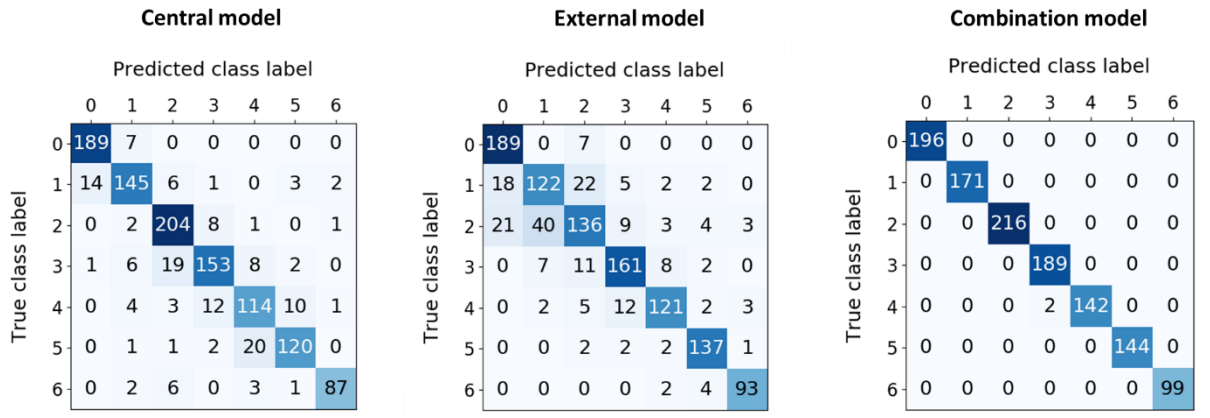

Figure 3: Confusion matrixes for selected kNN models. The matrix on the left refers to the model using only features from central detectors positions, the matrix in the center refers to the model using only features from external detectors positions, and the matrix on the right refers to the model using features both from central and external detectors positions.

Table 4: Overview of misclassifications where defect fuel assemblies are classified as complete fuel assemblies. The table on the left refers to the central model whereas the table on the right refers to the external model.

Central model	
ID diversion scenario	Number of misclassifications
D01201	6
D01202	5
D01204	1
D01205	2
D07701	1

External model	
ID diversion scenario	Number of misclassifications
D00401	9
D01501	1
D02406	8
D02801	7
D03601	5
D04802	9

4 Conclusion

Machine learning models using the kNN approach were applied for the classification of fuel assemblies with replaced pins. The classification was performed in terms of the percentage of fuel pins being replaced.

Several models were developed considering different detector responses and different detector positions. The results show that the detector type plays an important role in the model accuracy, and the ionization chambers appear to be more sensitive to the fuel replacement compared to fission chambers. Moreover, the classification accuracy increases by using detectors placed in the central guide tube positions and outside the fuel assembly.

Particular focus was placed on the diversion scenarios where defect fuel assemblies were misclassified as complete fuel assemblies by the machine learning model. The misclassification mostly occurred for diversion scenarios with <20% replaced pins. Further investigations will be carried out to explain the origin of the misclassification.

Future work will also include the comparison between the optimized models and other models using other machine learning approaches such as decision trees and neural networks. The extension of the dataset to other diversion scenarios will also be considered.

References

Bachmann A. M., et al., 2021. “Comparison and uncertainty of multivariate modelling techniques to characterize used nuclear fuel“. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, Volume 991.

Bishop C. M., 2006. “Pattern recognition and machine learning”. Springer, ISBN-13: 978-0387-31073-2.

Borella, A. et al., 2017. “Signatures from the spent fuel: simulations and interpretation of the data with neural network analysis.” ESARDA Bulletin 55, Pages 29–38.

Borella, A. et al., 2018. “Simulated observables for spent fuel non-destructive assay”. Proceedings of the 2018 INMM annual meeting.

Gauld, I.C., et al., 2011. ”Origen-ARP: automatic rapid processing for spent fuel depletion, decay, and source term analysis”. ORNL/TM-2005/39 Version 6.1 Sect. D1.

Grape S., et al., 2020. “Determination of spent nuclear fuel parameters using modelled signatures from non-destructive assay and Random Forest regression“. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, Volume 969.

Ham, Y.S., et al., 2009. “Development of a safeguards verification method and instrument to detect pin diversion from pressurized water reactor (PWR) spent fuel assemblies phase I study”. Lawrence Livermore National Laboratory technical report LLNL-TR-409660.

Ham, Y.S. et al., 2015. “Partial defect verification of spent fuel assemblies by PDET: principle and field testing in interim spent fuel storage facility (CLAB) in Sweden. Proceedings of the 2015 ANIMMA conference.

Hellesen C., et al., 2017. “Nuclear spent fuel parameter determination using multivariate analysis of fission product gamma spectra”. Annals of Nuclear Energy, volume 110, Pages 886-895.

Iggane, M. et al., 2012. “Self-training using a k-Nearest Neighbor as a base classifier reinforced by Support Vector Machines”. Int. J. Comput. Appl. 56 (6).

International Atomic Energy Agency (IAEA), 1970. “Treaty on the non-proliferation of nuclear weapons”. INFCIRC/140.

International Atomic Energy Agency (IAEA), 1972. “The structure and content of agreements between the Agency and States required in connection with the treaty on the nonproliferation of nuclear weapons”. INFCIRC/153 (corrected).

International Atomic Energy Agency (IAEA), 1997. “Model protocol additional to the agreement(s) between state(s) and the International Atomic Energy Agency for the application of safeguards”. INFCIRC/540.

International Atomic Energy Agency (IAEA), 2002. “IAEA safeguards glossary 2001 edition”. International nuclear verification series No. 3.

International Atomic Energy Agency (IAEA), 2011. “Safeguards techniques and equipment: 2011 edition”. International nuclear verification series no. 1 (rev. 2).

International Atomic Energy Agency (IAEA), 2020. “The safeguards implementation report for 2019”. GOV/2020/9.

Kourou K., et al., 2015. “Machine learning applications in cancer prognosis and prediction”. Computational and Structural Biotechnology Journal, Volume 13, 2015, Pages 8-17.

Mitchell, T.M., 1997. “Machine Learning”, McGraw-Hill editor.

Murphy K. P., 2012. “Machine learning: a probabilistic perspective“. Massachusetts Institute of Technology, ISBN 978-0-262-01802-9.

Ngai E. W. T., et al., 2011. “The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature”. Decision Support Systems, Volume 50, Issue 3, Pages 559-569.

Pedregosa F., et al., 2011. “Scikit-learn: Machine Learning in Python”. Journal of Machine learning Research, volume 12, pages 2825-2830.

Orton C. R., et al., 2011. “Proof of concept simulations of the multi-isotope process monitor: An online, nondestructive, near-real-time safeguards monitor for nuclear fuel reprocessing facilities”. Nuclear Instruments and Methods in Physics Research A 629 (1) 209–219.

Rossa R., 2016. “Advanced non-destructive methods for criticality safety and safeguards of used nuclear fuel”. Ph.D. dissertation, Université libre de Bruxelles. Available online at: oai:dipot.ulb.ac.be:2013/238202/Holdings.

Rossa R., et al., 2019. “Use of machine learning models for the detection of fuel pin replacement in spent fuel assemblies”. ESARDA Bulletin No.58.

Rossa R., et al., 2020-a. “Development of the SCK CEN reference datasets for spent fuel safeguards research and development“. Data in Brief 30 (2020) 105462.

Rossa R., et al., 2020-b. “Comparison of machine learning models for the detection of partial defects in spent nuclear fuel”. Annals of Nuclear Energy, Volume 147.

Sitaraman, S. et al., 2007. “Characterization of a safeguards verification methodology to detect pin diversion from pressurized water reactor (PWR) spent fuel assemblies using Monte Carlo techniques”. Proceedings of the 48th INMM Annual Meeting.