DISCRIMINATING BETWEEN IRRADIATED MYRRHA FUEL AND LIGHT WATER REACTOR FUELS USING GAMMA RAYS AND NEUTRONS

Markus Preston
Department of Physics and Astronomy,
Uppsala University, Sweden

Erik Branger
Department of Physics and Astronomy,
Uppsala University, Sweden

Alessandro Borella
SCK CEN,
Belgian Nuclear Research Centre, Belgium

Sophie Grape
Department of Physics and Astronomy,
Uppsala University, Sweden

Riccardo Rossa
SCK CEN, Belgian Nuclear Research Centre, Belgium

ABSTRACT

MYRRHA is an accelerator-driven system featuring a MOX-fuelled core cooled by lead-bismuth eutectic, which is under development at SCK CEN in Mol, Belgium. An initial plutonium content of 30% in the fuel is foreseen, which together with the fast neutron spectrum in the core results in considerably different spent-fuel properties compared to spent fuel from typical light water reactors. These differences have been studied through depletion simulations, and include how radionuclide densities depend on burnup, which radionuclides contribute to the gamma-ray and neutron emission, and the intensity of the emitted radiation. As a consequence, current techniques for safeguards verification of spent fuel via non-destructive assay may need to be updated, or new techniques developed, for use in safeguarding of spent MYRRHA fuel. To some extent, these differences may have consequences in the wider context of nuclear safeguards for Generation IV systems. The focus of this paper is to investigate to what extent the gamma-ray and neutron signatures could be used to differentiate irradiated MYRRHA fuel from irradiated MOX and UO_2 fuels from a light water reactor. The ability to discriminate between light-water-reactor MOX and UO_2 has been recognised as an important task in safeguards verification today, and this work extends this objective to a future nuclear energy system.

Keywords: Spent nuclear fuel, MYRRHA, Non-destructive assay (NDA), Classification

INTRODUCTION

In safeguards verification of spent nuclear fuel, different radiation signatures measured by non-destructive assay (NDA) instruments can give information about the properties of the fuel and be used to verify the completeness and correctness of operator declarations. In the verification process, the parameters initial enrichment (IE), burnup (BU) and cooling time (CT), which characterise spent fuel well, are often used and referred to [1]. Although UO_2 fuel is the most common type of fuel under safeguards today, a number of countries are using or have been using mixed oxide (MOX) fuel. For MOX fuel, the initial plutonium content (IPC) may be used instead of the IE. The isotopic mixture of the plutonium in the fresh fuel may affect the properties of the spent fuel [1], so that the meaning of the IPC is not as well defined as the IE is for UO_2.

Although similar NDA techniques can be used to verify spent UO_2 and MOX fuels [2], situations have been identified where radiation signatures from the two are so similar that it would be

difficult to discriminate between the two types of fuel. Because MOX fuel contains a higher percentage of plutonium than UO_2 fuel, it is important that such discrimination is possible. To ensure this, the SMOPY detector [3] was developed for use on spent light-water reactor (LWR) fuels. Furthermore, several methods for discriminating between the two fuel types using gamma rays have been proposed [4, 5, 6]. Of these, [6] proposed to use machine learning classification algorithms for discriminating between the two fuel types.

In this work, we extend the methodology proposed in [6], and include LWR-UO_2 and LWR-MOX fuels as well as fuel from an innovative research reactor; the lead-bismuth-cooled subcritical reactor in the upcoming MYRRHA facility. MYRRHA will be a fast-neutron reactor containing MOX fuel with an IPC of 30%, making the radionuclide composition of the spent fuel quite different from LWR UO_2 or MOX fuels (in LWR-MOX reactors, the IPC is typically well below 12% [7]). It is vital to understand which radiation signatures are most important for discriminating between the three types of fuel in a safeguards context.

MYRRHA AND ITS FUEL

The MYRRHA facility [8] is under development at the SCK CEN research centre in Mol, Belgium. It is an innovative system, where a MOX-fuelled reactor core cooled by lead-bismuth eutectic (LBE) will be coupled to a 600-MeV linear proton accelerator. When the accelerated protons interact with the LBE, spallation neutrons are produced. When MYRRHA is run in the "sub-critical mode", the accelerator is used to sustain the neutron chain reaction in the core. MYRRHA can also be run in a "critical mode", where the accelerator is not used and additional fuel assemblies are introduced into the core to achieve a critical core configuration. In this work, only spent fuel from the critical mode is considered.
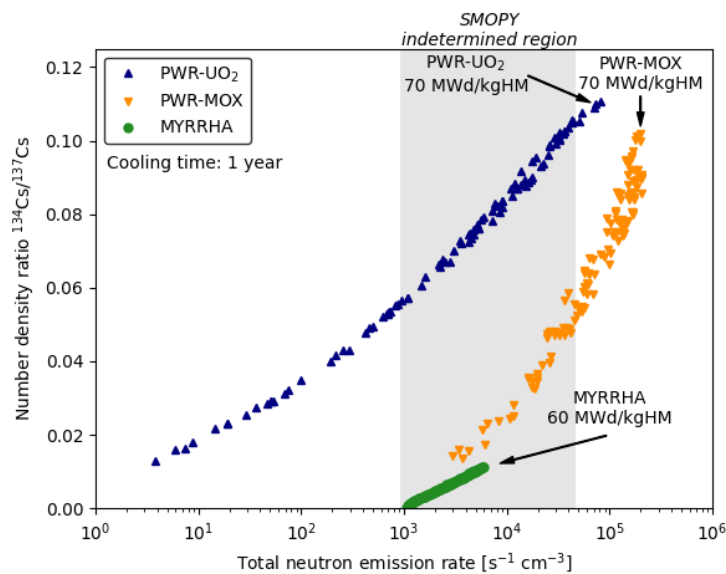
In [9], depletion simulations of MYRRHA fuel were used to show that the fast neutron energy spectrum in MYRRHA affects the production rate of various radionuclides in the fuel, and as a result also the characteristics of the radiation emitted from spent MYRRHA fuel will be different from spent LWR UO_2 and MOX fuels. It was found that the production of ^{134}Cs and ^{154}Eu, two gamma-ray emitters important for safeguards, is lower in MYRRHA fuel than in LWR fuels with comparable BU and CT. Similarly, the production of ^{244}Cm, the most important neutron emitter in spent LWR fuel, is lower in MYRRHA than in the LWR fuels. These differences were attributed to the decrease in neutron capture cross sections with increasing neutron energy, since these three nuclides are predominantly produced through one or several neutron captures. There are other differences as well: for example some nuclides have a slightly higher production rate in MYRRHA than in the LWR cases due to different fission yields and capture cross sections.

All data on the radionuclide inventory in the spent fuel in this work have been produced using Serpent2 [10], where Monte Carlo simulations of neutron transport are coupled with a depletion-equation solver. The LWR data come from depletion simulations of pressurised water reactor (PWR) UO_2 and MOX fuels described in [11] and the MYRRHA data come from depletion simulations described in [9]. The PWR-UO_2 data are available between IE of 1.5 wt-% and 6.0 wt-% (in steps of 0.1 wt-%), the PWR-MOX data between IPC of 4.0 wt-% and 10.0 wt-% (in steps of 0.2 wt-%) and the MYRRHA data at an IPC of 30 wt-%. The PWR-UO_2 and PWR-MOX data are available between BU of 5 and 70 MWd/kgHM (in steps of 0.5 MWd/kgHM), and the MYRRHA data between BU of 0.56 MWd and 60 MWd/kgHM (in steps of 0.56 MWd/kgHM). All three data use CT steps of 0.25 years up to 10 years and 0.5 years between 10 and 20 years. In this work, only fuel with a CT between 3 months and 20 years was included.

It is important to know what impact the differences between LWR and MYRRHA fuels in terms of radionuclide composition and radiation emission could have on safeguards verification of

spent fuel. The high IPC in MYRRHA fuel (30%) raises some safeguards concerns that have to be handled. Among the potential diversion scenarios that must be considered are replacement of spent MYRRHA fuel with spent LWR fuel and/or a lack of information due to lost continuity of knowledge. Because a MYRRHA fuel assembly is considerably smaller than a fuel assembly from a commercial reactor [9], we here consider a scenario where all geometric properties of the fuel have been removed. Such a situation would for example arise at the initial stages of a fuel reprocessing facility [12, 13], where spent fuel is shredded and dissolved. In that case, we would limit ourselves to the part of the facility prior to separation of fission products from uranium and plutonium. At these initial stages, direct NDA of uranium and/or plutonium is typically impossible due to the intense radiation emission from fission products and minor actinides.

As a first step, it is important to evaluate how useful existing NDA techniques for discriminating between UO_2 and MOX could be if adding MYRRHA fuel to the discrimination task. The SMOPY instrument [3] was developed to discriminate between low-BU MOX fuel and high-BU UO_2 fuel. Because the neutron emission from an LWR-MOX assembly is considerably higher than from an LWR-UO_2 assembly, the total neutron emission rate is first used to discriminate between the two. If the total neutron emission rate is neither very high nor very low, the assembly could be either low-BU MOX or high-BU UO_2. In that region, SMOPY uses the ratio $^{134}Cs/^{137}Cs$, obtained from a gamma-ray measurement, to discriminate between the cases. To evaluate the applicability of the SMOPY analysis methodology described in [3] to discriminate between these two types of LWR fuel and MYRRHA fuel, it was applied to the simulated data described above. The results, i. e. the total neutron emission rate and the $^{134}Cs/^{137}Cs$ number-density ratio, are shown in Fig. 1. These data were obtained directly from the simulated fuel libraries. For the LWR data, several IE/IPC values in the range available in the fuel library have been used. For the MYRRHA data, the only IPC used is 30%. Note that only data with a CT of one year have been used, as we assume that the SMOPY instrument requires some knowledge of the CT.

Figure 1: Scatter plot of the $^{134}Cs/^{137}Cs$ number-density ratio and the total neutron emission from spent PWR-UO_2, PWR-MOX and MYRRHA fuels, obtained from depletion simulations. The grey band shows the region in which the SMOPY instrument requires both the total neutron emission and the $^{134}Cs/^{137}Cs$ ratio to discriminate between UO_2 and MOX fuels from an LWR. The limits of this "indetermined" region were taken from [3].

It is clear that for the radiation signatures used by SMOPY, the MYRRHA fuel is considerably different from the LWR fuels. As shown in [9], both the total neutron emission rate and the ^{134}Cs/^{137}Cs ratio are expected to be lower in spent MYRRHA fuel than in spent LWR fuels. Therefore, if the SMOPY instrument is to be used in classification of MYRRHA spent fuel, the analysis method used by SMOPY will have to be extended to cover both the more complex case of three fuel types and the quite different (and lower) radiation emission from MYRRHA fuel. In this work, we take a more general approach by including additional radiation signatures as well as considering a scenario where the CT is not as well known.

METHOD

Input data

The MYRRHA fuel library contains 14,148 fuel samples, which is considerably less than both the PWR-UO$_2$ (789,406 samples) and PWR-MOX (531,991 samples) fuel libraries. To balance the three classes, 14,148 samples were drawn at random from each of the PWR libraries, resulting in a final dataset with 42,444 samples with one third from each of the fuel types.

To determine which types of radiation signatures are important for accurate discrimination, the classifiers were trained and tested for four scenarios: a) when only relative gamma-ray activities are measured, b) when also the absolute ^{137}Cs gamma-ray activity is measured, c) when a) and the absolute neutron emission rate are measured and d) when all of the above are measured.

In scenario a), only relative gamma-ray activities are used. In [9], the radionuclides ^{95}Nb, ^{95}Zr, ^{144}Ce, ^{106}Ru, ^{134}Cs, ^{125}Sb, ^{154}Eu and ^{137}Cs were identified as being the most important contributors to the gamma-ray emission above 500 keV in spent MYRRHA fuel based on depletion simulations of MYRRHA fuel. For each of these nuclides, one emission line was selected based on its energy and intensity, and the corresponding activity calculated using the branching ratio. Finally, each of these gamma-ray activities was divided by the ^{137}Cs 662-keV gamma-ray activity, resulting in seven activity ratios. This is typically done because the ratio is less sensitive to the measurement geometry than the activity itself [14]. As a result, scenario a) does not include information about the absolute activity of the fuel.

In scenario b), the relative gamma-ray activities of scenario a) were used, but also the absolute ^{137}Cs 662-keV gamma-ray activity. Because the ^{137}Cs content is approximately proportional to the BU of the fuel, this parameter is expected to yield additional information of interest. For long-CT fuel, the absolute ^{137}Cs 662-keV gamma-ray activity is expected to dominate the gamma-ray emission rate above 500 keV for all three fuel types [9].

Scenario c) includes the relative gamma-ray activities of scenario a) together with the absolute neutron emission rate. The neutron emission rate, which includes emission from spontaneous fission and (α, n) reactions in the spent fuel, depends on the BU of the fuel. This dependence was shown to be less pronounced for MYRRHA fuel than for PWR fuels in [9].

Finally, scenario d) includes the relative gamma-ray activities, the absolute ^{137}Cs gamma-ray activity and the absolute neutron emission rate. This is expected to provide the highest accuracy, at the expense of a more complicated measurement setup, calibration and analysis.

It is important to point out that we only consider the radiation emission in the spent fuel, and not the detection. Depending on the detector type and geometry and the geometry of the fuel sample, effects such as scattering, attenuation and multiplication will affect the radiation signatures. Nonetheless, this work serves as a starting point for identifying the underlying signatures of interest for this NDA task. In order to account for a real-life effect that does not require knowledge of the geometry of the setup, we have incorporated a threshold on the gamma-ray

activity ratios. If the activity ratio A_X/A_{137} (the gamma-ray activity of nucleus X relative to the gamma-ray activity of ^{137}Cs) was less than 0.1% for a given fuel sample, that feature was set to zero for that sample. This was done to imitate a realistic measurement where some radionuclide activities will be too low to be reliably detected. ^{137}Cs is the most long-lived of the selected gamma-ray emitters and is therefore assumed to be detectable in all fuels. Although this method of thresholding says nothing about the time required to achieve sufficient statistics, it can be reasonably assumed that in any case good statistics on ^{137}Cs will be a prerequisite. No threshold was applied to the absolute ^{137}Cs gamma-ray emission or the absolute neutron emission rate. Such thresholds will be considered in the future, when a realistic NDA setup is considered. In this work, we only assume that the measurement time is adjusted so that the absolute ^{137}Cs activity and/or absolute neutron emission rate can be determined when necessary.

Two important differences in methodology with respect to [6] are worth mentioning. First, in [6], the data were grouped into sets of similar CT and the classification trained and tested independently on each set. This means that at least some information about the CT of the fuel is required for classification. In reality, a potential diversion scenario could consist of replacing short-cooled MYRRHA fuel with long-cooled PWR fuel. If these fuels are not in the same "CT set", the algorithm in [6] will not work. Here, we therefore use all data with CT up to 20 years in the training and testing of the classifiers. This is a more challenging task due to the additional variations introduced by the CT, but is a more realistic case that should be considered. The second difference between this work and [6] is that gamma-activity ratios were not considered as features in [6]. This further complicates the classification in scenario a) of this work, because much information about BU and CT is lost when using relative, rather than absolute, activities.

In order to train and test the classification algorithms, 60% of the dataset was drawn at random to be used in training and the remaining 40% to be used only for testing. The same training and testing data were used by all classification algorithms, to remove any variations due to using different data. The dataset was balanced, meaning that it contained equal parts PWR-UO$_2$, PWR-MOX and MYRRHA data.

Principal Component Analysis

As in [6], we have used both principal component analysis (PCA) and different classification algorithms in this work. PCA was used to examine how the different input features (i. e. the radiation signatures) explain different directions of variation in the data, and could be used as a basis for selecting an appropriate subset of the input features in future work. Both the PCA and the classification algorithms were implemented using the Python library Scikit-learn [15].

PCA is an unsupervised algorithm which selects the N directions (the PCA eigenvectors) in the N-dimensional input space that account for the highest variance. Prior to using the PCA, each input feature was re-scaled to a normal distribution with a mean of zero and a standard deviation of one based on the training data. This scaling could then be applied to any test data, and was used on all data before classification (as some of the classifiers also require re-scaled features).

PCA can be used to identify important directions in the data, directions which allow separation between fuel types, and input features which contribute more or less to those directions. It is important to point out that the spread in the data is due to both the fuel type and the underlying variables IE/IPC, BU and CT. If multiple features follow similar directions in the data, one or more of these features could be discarded to yield more robust classification. This use of the PCA, i. e. feature selection, was not used in this work. It was however found that the two most short-lived radionuclides (^{95}Nb and ^{95}Zr) contribute very similarly to the principal components. It was also found that ^{154}Eu is the most important gamma-emitting radionuclide for separating the

three fuel types, followed by ^{134}Cs. These conclusions agree with the results in [9], and a more detailed study could be made in the future to reduce the number of input features using a PCA.

Classifier training and hyperparameter optimisation

In this work, three classification algorithms were selected: linear discriminant analysis (LDA) [16], k-nearest neighbour (kNN) [17], and Gaussian support vector machine (SVM) [18]. In [6], these algorithms were shown to be able to accurately discriminate between LWR UO$_2$ and MOX fuels for CT up to 20 years. LDA is a parametric classifier assuming that the data can be separated by a hyperplane, kNN is a non-parametric classifier working directly in the input space and the Gaussian SVM is a non-parametric classifier that projects the data into a higher dimension before classification. These algorithms work in quite different ways, and the selection allows a more general evaluation of which classifier type is suited for the discrimination task.

Each of the three classifiers has its own set of hyperparameters, which in part will determine the performance of the algorithm. To find the optimal hyperparameter settings, Bayesian hyperparameter optimisation from the scikit-optimize [19] library was used. Because the dataset was balanced, the accuracy was used as the performance metric. In Bayesian hyperparameter optimisation, the prior distributions for the different hyperparameters are defined and the parameter space searched for the configuration yielding the best accuracy. Table 1 lists the used hyperparameter priors and the optimal values found in the present work. For the hyperparameter optimisation, five-fold cross validation within the training dataset was used, thereby ensuring that the test data was independent of the training procedure.

Table 1: Optimal hyperparameter settings for each classifier and each input-feature scenario, determined using Bayesian optimisation.

Hyperparameter	Prior	Scenario a)	Scenario b)	Scenario c)	Scenario d)
kNN					
n_neighbors	[2, 10]	4	4	2	2
Weights	[uniform, distance]	distance	distance	distance	uniform
Algorithm	[ball_tree, kd_tree, brute, auto]	auto	brute	brute	brute
Leaf size[1]	[1, 50]	40	N/A	N/A	N/A
p	[1, 2]	1	1	1	1
LDA					
shrinkage	Real(0, 1)	0.380	0.354	0	0.0145
solver	[lsqr, eigen]	eigen	lsqr	lsqr	lsqr
Gaussian SVM					
C	Real($1 \cdot 10^{-6}$, $1 \cdot 10^{9}$)	$2.93 \cdot 10^{6}$	$2.66 \cdot 10^{6}$	$2.58 \cdot 10^{4}$	$3.79 \cdot 10^{4}$
gamma	Real($1 \cdot 10^{-6}$, 20)	0.121	0.000500	0.470	$2.44 \cdot 10^{-5}$

1   The *leaf size* parameter is not applicable when the nearest neighbours are found using a brute-force search (i. e. *algorithm*=brute).

Although a full discussion of the hyperparameters is outside the scope of this work, a few remarks can be made about the hyperparameter optima in Table 1. The optima for scenarios a) and b) are relatively similar, as are the optima for scenarios c) and d). The optimum values for

hyperparameters *n_neighbors* (for kNN), *shrinkage* (for LDA) and *C* (for Gaussian SVM) are larger for scenarios a) and b) than they are for scenarios c) and d). These parameters control, in different ways, how stringent the training procedure should be. By adjusting these hyperparameters, the risk of overfitting may be reduced. The fact that the *n_neighbors* and *shrinkage* hyperparameters are smaller for scenarios c) and d) indicates that adding certain input features improves the classification enough that the classifier can stay close to the training data without risking overfitting. Interestingly, the SVM hyperparameter *C* is smaller for scenarios c) and d), which generally indicates that the classifier tolerates more mis-classifications during training to yield a more general classifier. That is, when using the SVM classifier, adding certain input features requires making the SVM model more general by decreasing *C*.

Classifier testing

The performance of the trained classifiers was evaluated using the test data as input. To evaluate the robustness of the classification to noise, different levels of Gaussian noise were applied to the input features of the test data. The evaluated noise levels were between 0% and 20% of the input-feature values. Again, the attainable measurement accuracy will depend on the measurement geometry and detector characteristics. Also, different levels of accuracy for gamma-ray and neutron measurements is realistic. Finally, the actual noise level will be connected to the measurement time and therefore also the detection threshold. Modelling the gamma-ray and neutron *detection* is outside the scope of this work, but the procedure used here is a good first-order approximation of a generic measurement setup.

There are limitations with using the accuracy as a performance metric, since there are three classes (PWR-UO$_2$, PWR-MOX and MYRRHA fuels). For example, the accuracy does not show if the classifier mainly has difficulties when discriminating between only two of these three fuel types. It is of course interesting to determine between which fuel types mis-classification is most likely to occur. To this end, confusion matrices were constructed for each detection scenario.

RESULTS AND DISCUSSION

Fig. 2 shows the accuracies of the discrimination for the three classification algorithms and the four detection scenarios a)-d).
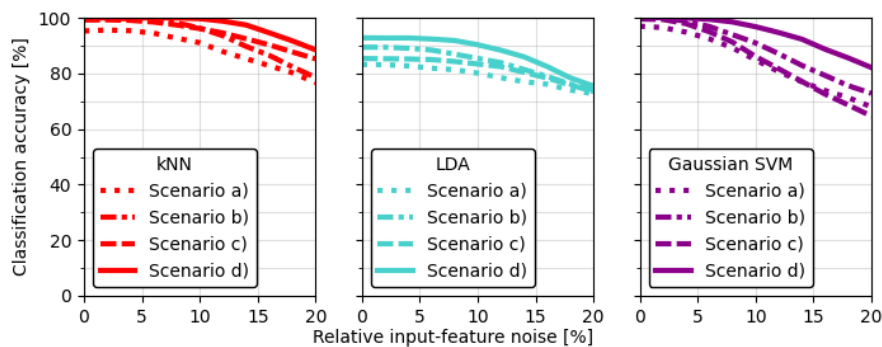

Figure 2: The classification accuracy for the three trained classifiers. Scenario a) includes only the relative gamma-ray activities, scenario b) adds the absolute [137]Cs gamma-ray activity, scenario c) adds the absolute neutron emission rate to scenario a), and scenario d) includes all of the above features.

A few things are worth noting:

- As expected, increasing the noise on the input features results in an accuracy decrease. The kNN and LDA classifiers are more robust to noise than the Gaussian SVM classifier, with the kNN showing the best overall accuracy.

- Including additional features (the absolute [137]Cs gamma-ray activity and the absolute neutron emission rate) improves the discrimination capabilities. Scenario b) typically yields a higher accuracy than scenario c), meaning that it is preferable to add the absolute [137]Cs gamma-ray activity rather than the absolute neutron emission rate.

- For some values of the input-feature noise, the lines cross each other. This can be due to both the fact that some features may be more robust against noise and that the added noise introduces statistical differences between the test data for the different scenarios.

- Including the relative gamma-ray activities and the absolute [137]Cs gamma-ray activity (i. e. scenario b) results in accuracies close to 100% when using the kNN and Gaussian SVM classifiers. Further adding the absolute neutron emission rate (i. e. scenario d) makes the classification more robust to noise and results in the most accurate results.

- The accuracy of the LDA is, for any given scenario, considerably lower than it is for the other two classifiers. This marks a contrast with the results in [6], where the difference was much smaller. The difference could be due to the fact that the addition of MYRRHA fuel creates a multiclass classification problem, which is more difficult for the linear LDA classifier. As in [6], the performance of the kNN and SVM classifiers is similar.

As discussed in the introduction, this work is especially motivated by the fact that the plutonium content in MYRRHA fuel is higher than that it is in typical LWR fuel. Replacing spent MYRRHA fuel with spent LWR fuel is therefore a potential diversion scenario that should be considered. To make a first survey of which types of mis-classification can happen in the four detection scenarios, confusion matrices were constructed. The confusion matrix relates the number of classifications in each fuel type to the number of true samples of each fuel type. Because kNN was identified as the best-performing classifier of the methods used, Fig. 3 shows the confusion matrices in the cases of 0%, 10% and 20% noise for the kNN classifier, for detection scenario d).
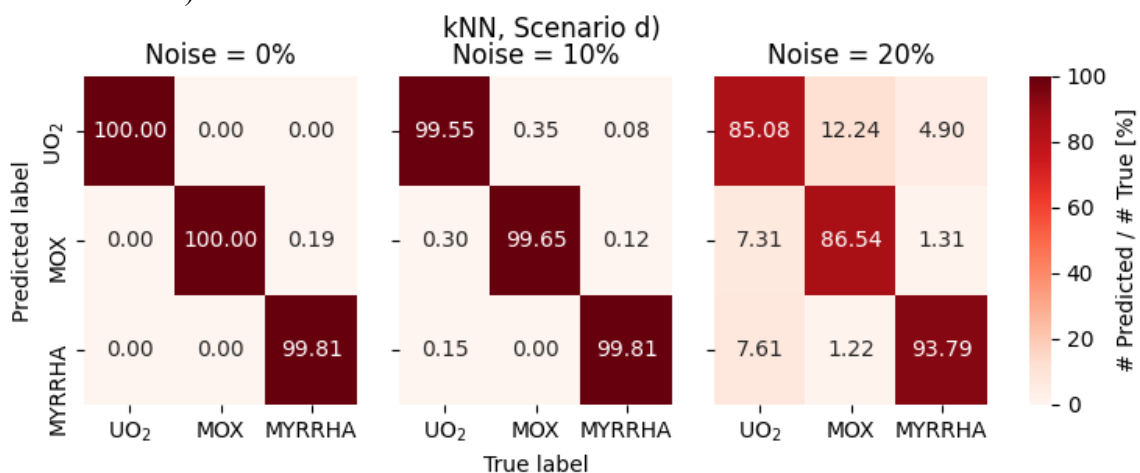
Figure 3: Confusion matrices showing the percentage of the true fuel types that were correctly classified when using the kNN classifier and scenario d). The impact of the added Gaussian noise on the classification performance is visible.

The confusion matrices show that in the absence of added noise, nearly all data are perfectly identified using the kNN classifier. The situation is similar for the Gaussian SVM. As the noise increases, the frequency of mis-classification increases, including the fuel from MYRRHA. Nonetheless, at 10% added noise the accuracy is still above 99% for all three fuel types. At 20% added noise, the frequency of mis-classifications increases, particularly among the LWR data. At this noise level, the risk for mis-classifying MYRRHA fuel as PWR-UO$_2$ appears to be higher than mis-classifying MYRRHA fuel as PWR-MOX fuel. The situation is the same in scenario d) also with the LDA and SVM classifiers, and the only clear exception to this trend occurs for scenario a), where the risk for mis-classifying MYRRHA as PWR-MOX is greater. The reason is most likely that once absolute radiation signatures (absolute ^{137}Cs and/or absolute neutron) have been added, clearer information about the fuel BU is available. When that is the case, features such as ^{154}Eu/^{137}Cs and the absolute neutron emission rate provide more clear differentiation between MYRRHA and PWR-MOX than between MYRRHA and PWR-UO$_2$. Based on these results, it is clear that scenario d) is needed to allow accurate classification at noise levels of 20%. At lower noise levels, scenario b) could be sufficient.

CONCLUSIONS AND OUTLOOK

The results in this paper show that using machine learning classification methods could be a promising methodology for discriminating between LWR fuel and MYRRHA fuel. Classification accuracies of up to 100% have been obtained for data including PWR-UO$_2$, PWR-MOX and MYRRHA fuel with CT up to 20 years. The results are also relevant in the more general context of safeguards of future nuclear energy systems.

Comparing with the results in [6], the results presented here show that adding MYRRHA fuel does make the discrimination task more difficult than the case with only LWR UO$_2$ and MOX. It should however be kept in mind that the relative gamma-ray activities alone were not used as features in [6], making scenario a) (only the relative gamma-ray activities) more challenging in itself. Nonetheless, it seems that only using relative gamma-ray activities as radiation signatures is not sufficient, and adding the absolute ^{137}Cs gamma-ray activity and/or the absolute neutron emission rate is required to reach accuracies of up to 100%. A topic for future work could be to start with the full set of input features and remove features one at a time to find the optimal trade-off between number of features and accuracy.

The parameters IE/IPC, BU and CT are often used to characterise spent nuclear fuel, and can be used to describe the variations in the radiation signatures *for a particular fuel type*. For example, the neutron emission increases with BU. Therefore, these parameters can be useful in describing the regions in the data where overlap between the different fuel types are most likely. In the present analysis, the parameters IE/IPC, BU and CT may be viewed as hidden variables within each fuel-type class. A topic for future work is therefore to identify which regions of IE/IPC, BU and CT result in the most mis-classifications. Such an analysis could help shed more light on the risk of mis-classifying UO$_2$ fuel as MYRRHA fuel shown in Fig. 3. Furthermore, the potential influence of variations in the irradiation history on the classification performance should be investigated.

Finally, and perhaps most importantly, detection of the radiation emitted from the spent fuel must be modelled to accurately estimate realistic counting statistics, thresholds and systematic effects such as self-attenuation. For this, the actual application should be well-defined, for example whether the fuel to be verified is an assembly, is shredded or dissolved. This is also a topic for future work. From this work, it can be concluded that efforts should be made to minimise the noise on the measured radiation signatures to provide accurate classification. This consideration should be taken into account when considering a realistic detection scenario.

ACKNOWLEDGEMENTS

REFERENCES

[1] A. M. Bolind, The use of the BIC set in the characterization of used nuclear fuel assemblies by nondestructive assay, Ann. Nucl. Energy 66 (2014) 31-50.

[2] S. Vaccaro et al., Advancing the Fork detector for quantitative spent nuclear fuel verification, Nucl. Instrum. Methods A 888 (2018), 202-217.

[3] A. Lebrun, M. Merelli, J-L. Szabo, M. Huver, R. Arlt, and J. Arenas-Carrasco, SMOPY a new NDA tool for safeguards of LEU and MOX spent fuel, International Atomic Energy Agency Report IAEA-SM-367/14/03 (2003).

[4] C. Willman, A. Håkansson, O. Osifo, A. Bäcklin, and S. Jacobsson Svärd, A nondestructive method for discriminating MOX fuel from LEU fuel for safeguards purposes, Ann. Nucl. Energy 33 (9) (2006) 766-773.

[5] M. L. Dennis and S. Usman, Feasibility of ^{106}Ru peak measurement for MOX fuel burnup analysis, Nucl. Eng. Design 240 (10) (2010) 3687-3696.

[6] M. Åberg Lindell, P. Andersson, S. Grape, C. Hellesen, A. Håkansson, and M. Thulin, Discrimination of irradiated MOX fuel from UOX fuel by multivariate statistical analysis of simulated activities of gamma-emitting isotopes, Nucl. Instrum. Methods A 885 (2018), 67-78.

[7] Nuclear Energy Agency – Organisation for Economic Co-operation and Development, Physics of plutonium recycling volume 1 – Issues and perspectives, OECD Documents (1995).

[8] H. Aït Abderrahim, P. Baeten, D. De Bruyn, and R. Fernandez, MYRRHA – A multi-purpose fast spectrum research reactor, Energ. Convers. Manage. 63 (2012) 4-10.

[9] M. Preston, A. Borella, E. Branger, S. Grape, and R. Rossa, Analysis of radiation emission from MYRRHA spent fuel and implications for non-destructive safeguards verification, Ann. Nucl. Energy 163 (2021) 108525.

[10] J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho, The Serpent Monte Carlo code: Status, development and applications in 2013, Ann. Nucl. Energy 82 (2015) 142-150.

[11] Zs. Elter, L. Pöder Balkeståhl, E. Branger, and S. Grape, Pressurized water reactor spent nuclear fuel data library produced with the Serpent2 code, Data in Brief 33 (2020) 106429.

[12] E. A. Hakkila et al., Materials management in an internationally safeguarded fuels reprocessing plant, Tech. Rep. LA-8042, Vol. II, Los Alamos National Laboratory, USA (1980).

[13] P. C. Durst et al., Advanced safeguards approaches for new reprocessing facilities, Tech. Rep. PNNL-16674, Pacific Northwest National Laboratory, USA (2007).

[14] D. Reilly, N. Ensslin, and H. Smith, Jr., Passive nondestructive assay of nuclear materials, Tech. Rep. LA-UR-90-732, Los Alamos National Laboratory, USA (1991).

[15] F. Pedregosa et al., Scikit-learn: machine learning in Python, J. Mach. Learn. Res. 12 (2011) 2825-2830.

[16] R. A. Fisher, The use of multiple measurements in taxonomic problems, Ann. Eugenic. 7 (2) (1936) 179-188.

[17] E. Fix and J. L. Hodges, Jr., Discriminatory Analysis. Nonparametric Discrimination: Consistency Properties, Int. Stat. Rev. 57 (3) (1989) 238-247.

[18] B. E. Boser, I. M. Guyon, and V. N. Vapnik, A training algorithm for optimal margin classifiers, in: COLT '92: Proceedings of the fifth annual workshop on Computational learning theory (1992) 144-152.

[19] T. Head, M. Kumar, H. Nahrstaedt, G. Louppe, and I. Shcherbatyi, scikit-optimize/scikit-optimize (version v0.8.1), 2020 [Online] Available: http://doi.org/10.5281/zenodo.4014775