

## Image analysis to support DCVD verification

**Erik Branger**  
Uppsala University

**Sophie Grape**  
Uppsala University

**Markus Preston**  
Uppsala University

### ABSTRACT

The Digital Cherenkov Viewing Device (DCVD) is one instrument available to authority inspectors for verifying spent fuel assemblies in wet storage. The measurements result in images of the Cherenkov light emissions from the fuel assembly under study. This work presents research on applying image analysis and statistical methods to improve data quality and to extract more information from the measurements, extending the use of these methods beyond what is currently implemented in the DCVD software. The goal of this project is to apply template matching and statistical analysis to the images. However, before such techniques can be applied, effort is needed to ensure that the measurements are directly comparable.

Two main issues are investigated here, the first being the positioning of the Region-Of-Interest. By developing an automated Region-Of-Interest placer, a consistent and reproducible Region-Of-Interest placement can be achieved. The second is automatic identification of fuel type, to support a later comparison with a template. We demonstrate that a method based on Principal Component Analysis can be used to determine the fuel type. Finally, we present the first results regarding template matching, comparing a measured image to a template, aiming to identify regions in the image where the two differ. Such differences could be due to a partial defect located in that region, but also due to other reasons such as debris covering the fuel top. Automatically identification of such regions can in the future be used to focus inspector attention to features requiring expert judgement, supporting efficient use of the measurement data and inspector effort. The first results demonstrate the feasibility of the method, but also that more work is required before the method is robust.

Keywords: Nuclear fuel, partial defect verification, Cherenkov light, DCVD.

### 1. INTRODUCTION

One of the key tasks for international nuclear safeguards inspectors is the verification of spent nuclear fuel (SNF). Multiple instruments have been developed to aid the inspectors in making independent measurements and assessments of SNF[1]. Due to the intense radiation emitted by the fission products and minor actinides, direct assessments of the radiation emitted by the nuclear material is not feasible. Thus, SNF measurements typically aim to quantify the radiation emission by the fission products and minor actinides, to verify that it is consistent with operator declarations of burnup (BU), initial enrichment (IE) and cooling time (CT), or BIC collectively.

A commonly used safeguards technique is the measurement of Cherenkov light emissions by SNF in wet storage. This Cherenkov light is predominately created following decays of fission products, which emit beta and gamma particles. When the gamma rays Compton-scatter on electrons in the water, or when beta electrons pass through the cladding with sufficient energy, the high-speed electrons will produce Cherenkov light in the water. The Cherenkov light is the most intense in the ultraviolet (UV) region, hence instruments that measure Cherenkov light measures the UV component to have the highest sensitivity. The two most frequently used instruments for detection of Cherenkov light are the Improved Cherenkov Viewing Device (ICVD), an analogue device turning UV light into visible light that an inspector can see, and the Digital Cherenkov Viewing Device (DCVD) [1], which is capable of both displaying and

recording the measured UV light intensity. The next-generation XCVD and autonomous robotic RCVD [2] are under development, but are foreseen to have capabilities in line with the DCVD.

## **2. THE DCVD MEASUREMENT METHODOLOGY**

Both the ICVD and DCVD are used for gross defect verification, verifying that the item under study is a SNF as opposed to a non-fuel item or dummy. This verification is performed by qualitatively studying the collimation of Cherenkov light emitted by the SNF. The collimation of light is typically pronounced for a real SNF, and significantly less pronounced for a non-fuel object. The DCVD is also capable of performing partial defect verification, to verify that 50% or more of the SNF has not been diverted and replaced with non-radioactive rods. Two methods are used for this:

- The first one is based on quantitative Cherenkov light measurements. The quantitative verification is done by placing a Region-Of-Interest (ROI) over the image of the fuel being measured, and integrating the measured light intensity in the ROI to quantitatively determine the Cherenkov light intensity. From earlier simulation studies [3], it is known that a 50% partial defect will reduce the measured Cherenkov light intensity of a SNF by at least 30%. The partial defect verification methodology hence relies on a comparison between the measured Cherenkov light intensity and a predicted intensity, the latter which is based on operator declarations and simulations of the gamma and beta radiation emissions and the corresponding Cherenkov light production[4]. The methodology also requires that SNFs are grouped together according to fuel types with identical or very similar physical design, which allows a direct, relative comparison of the measured intensities in each group.
- The second is based on template matching, and is used to identify partial defects where no substitution has taken place, i.e. identification of missing rods in visible positions. A template, describing the position of the fuel rods, is applied to the image of the Cherenkov light intensity in the ROI, to determine the rod positions. For each fuel rod position, the software analyses whether the Cherenkov light intensity is low (i.e. consistent with the presence of a fuel rod) or a high (i.e. consistent with a missing rod), as Cherenkov light can be produced in the water located where the rod should have been. This method is primarily applied to measurements of BWR fuels, as such fuel designs typically have most fuel rods in visible positions.

### 2.1 Improvements to the image analysis methods

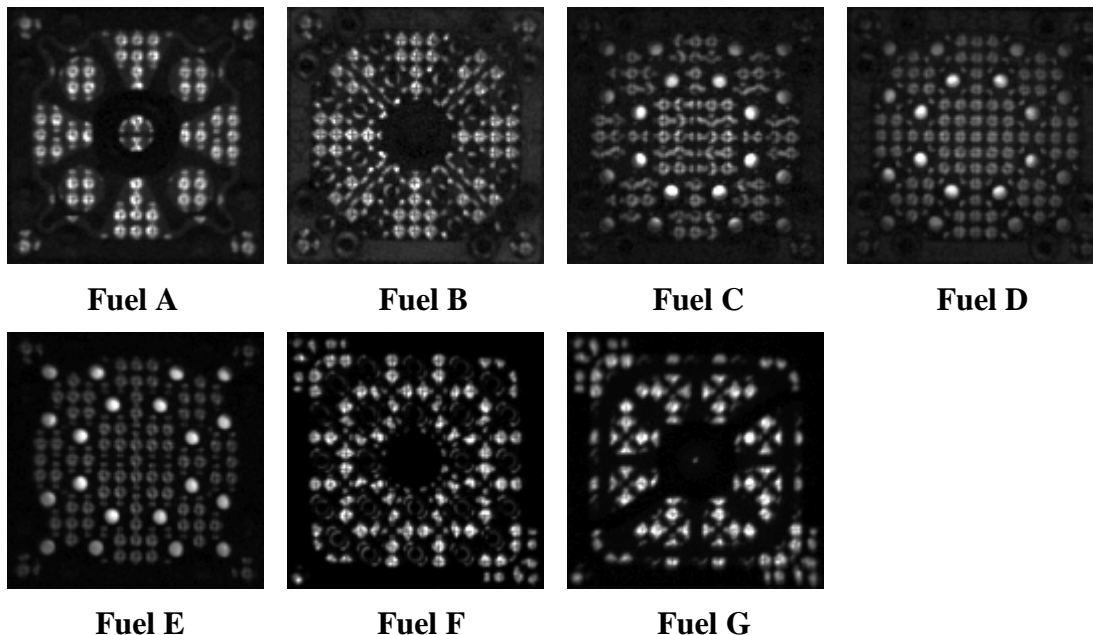
As part of an ongoing Action Sheet under EURATOM, research is ongoing to improve the DCVD image analysis to improve effectiveness and efficiency in SNF verification. Under the action sheet, Pacific Northwest National Lab and Los Alamos National Lab will investigate techniques such as neural networks for automated image analysis and Uppsala University will investigate classical image analysis techniques, multivariate analysis and statistical methods to develop automatic analysis of DCVD data. This work presents the first results from Uppsala University, aiming at applying template matching to identify partial defects. As part of developing the template matching, two issues were identified that needed to be solved before the template matching could be applied:

- The ROI was overall well placed in the available measurement data, but not with sufficient accuracy to allow automatic template matching. An automated ROI positioning method was therefore developed to reliably and repeatedly place the ROI over the fuel assembly that was measured. For the already implemented template matching, aimed at

identifying missing rods, an improved ROI placement can be used to improve the performance. Additionally, for quantitative measurements, the ROI placement can be a source of uncertainty, and systematically placing the ROI in the same position relative to the fuel in the image can help reduce this uncertainty.

- Multiple fuel types exist in the available data, and the correct template representing the fuel under investigation must be used in template matching procedure. An automatic determination of the fuel type was developed for this purpose. This can also be a tool to aid the inspector to ensure that the correct fuel assembly is being measured, and can ensure that the grouping according to fuel type is done correctly. Furthermore, some features of a SNF, such as the presence of a Rod Control Cluster Assembly (RCCA) inserted into a PWR fuel may not always be declared (and is thus unknown to the inspector). The presence of RCCAs impacts the detected Cherenkov light intensity to such an extent that it motivates fuels with a specific insert to constitute a fuel type of their own. Automatic identification and grouping of fuel would simplify the work of the inspector.

After the ROI had been accurately placed and the fuel type had been identified, template matching could be applied, and the first results are presented in this work. The template matching works by comparing a measured image to a reference image of the same fuel type to identify differences. Such differences could be due to partial defects, but could have other causes such as debris covering parts of the fuel. An automated method of highlighting regions where differences exist can aid an inspector, ensuring that the inspector expertise is focused on regions needing attention. The focus of the template matching studied here is for the case of relatively few rods replaced, which will not decrease the total Cherenkov light intensity significantly, or alter the overall light distribution, but may introduce local changes to the light distribution.



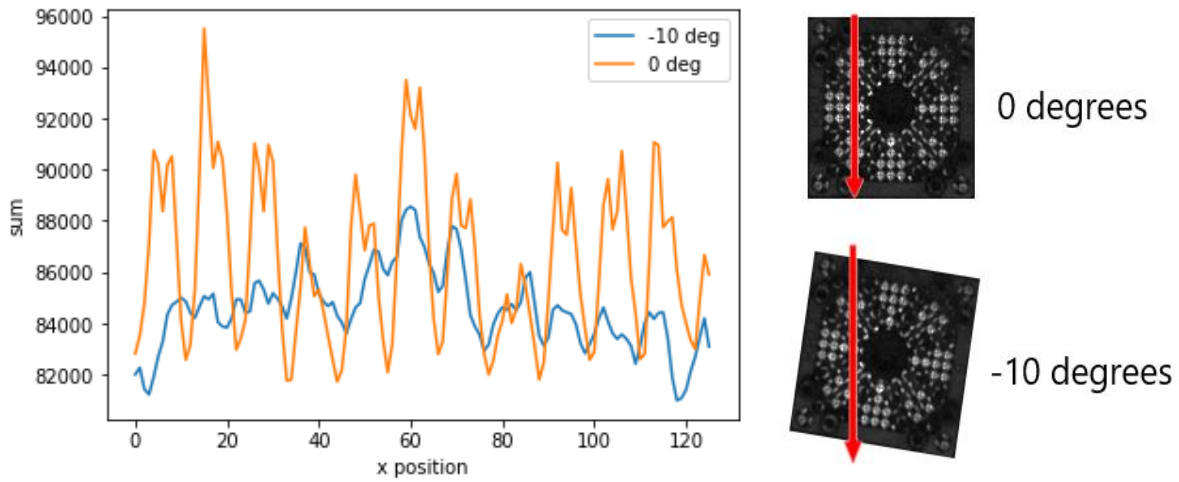
**Figure 1: One measurement from each of the seven fuel types used in this work. Fuel type A and B are PWR 16x16 with two different kinds of RCCA inserts. Fuel type C, D and E are PWR 16x16 without inserts and with subtly different physical design. Fuel type F and G are PWR 17x17 with two different types of RCCA inserts.**

As part of the action sheet, anonymized DCVD data has been provided to the participants, enabling the developed methods to be tested on measured data. For this study, a limited set of data consisting of 2800 images was used. The data consist of images of in total seven types of fuel assemblies. An example of one image from each fuel type is shown in Figure 1. The images show two types of PWR 17x17 fuel assemblies and five types of PWR 16x16 fuel assemblies. The PWR 16x16 fuels have two different RCCA insert types (A and B in Figure 1). Two PWR 17x17 fuel types have different RCCA inserts (F and G in Figure 1), whereas three PWR 16x16 fuel types lack inserts but have slightly different physical design (C, D and E in Figure 1). The data was selected to contain a sufficient number of good-quality images per fuel type (the smallest has 68 images) for statistical analysis to be performed. Additionally, they were chosen to include fuel types that are visually rather different (PWR 16x16 and 17x17, and noticeably different RCCA inserts), as well as fuel types that are challenging to distinguish from each other (the three PWR16x16 without inserts).

### **3. AUTOMATED ROI ADJUSTMENT**

For the data analyzed here, the inspectors have manually placed an ROI above the measured fuel assembly in connection to performing the measurements. Although the ROI placement is overall good, direct comparison of different images requires that the ROIs are systematically placed in the same way for each measurement. To allow a direct comparison between all measurements, the ROIs were rescaled to 128x128 pixels because the ROI size differed between measurement campaigns. With this image size, each fuel rod is about five pixels wide. Thus, if the ROI is misplaced by a few pixels, a direct comparison between the images becomes impossible. The situation is further complicated by that the fuel assembly in the ROI can be slightly rotated, with a rotation that can differ between assemblies. Therefore, to enable a direct comparison between measurements the rotation and position of the fuel needs to be automatically adjusted. The method should ideally work on all fuel types and thus all measurements used, and not be specific to a single fuel type.

To calculate the optimal rotation, the fact that the fuel rods form a square lattice is used. When the image is projected onto the x-axis, the rods will align for the correct rotation, as exemplified in Figure 2 for the image denoted “0 degrees”. Thus, at the x-coordinate of a column of rods, the projection will have a low value, due to the column of dark rods. For the bright spots between the rods, they will likewise line up and be projected onto a high value. For an incorrect rotation, the projection will cover both the rods and the bright spots, as exemplified in Figure 2 by the image denoted “-10 degrees”. In total, the correct rotation will show significant variation in the projection, while the incorrect rotation will be much flatter, as shown in the graph of Figure 2. To quantify this variation, the variance of the projection is calculated, where the maximum variance corresponds to the optimal rotation.



**Figure 2: An image projected on the x-axis, for a rotation of 0 degrees from optimal and a rotation of 10 degrees clockwise from optimal.**

To determine the center of the fuel assembly, the projection in Figure 2 can be used again. All studied fuel assemblies are symmetric when mirrored or rotated, and this symmetry extends to the projection. However, note that the symmetry may be broken if a substantial fraction of the fuel rods have been replaced with non-radioactive ones, but in such cases, the total intensity may also change sufficiently to be detected using such methods. If the projection on the x-axis is denoted  $P(i)$  for x-coordinate  $i$ , a symmetry score  $S(c)$  around the coordinate  $c$  can be defined as

$$S(c) = \sum_i (P(c - i) - P(c + i))^2$$


For the correct symmetry center  $c$ , it can be expected that  $P(c - i)$  and  $P(c + i)$  are similar due to the symmetry, hence their difference is small, and the total sum is thus also small. For other values of  $c$ , the difference will be systematically off from zero, hence the score will increase. Thus, the minimum value for  $S(c)$  indicates the position of the symmetry center. The same procedure, applied to the y-axis, can be used to locate the fuel center in the y direction. Once the optimal rotation is found and the fuel assembly center is located, the ROI can be adjusted in the original measurement, and the image inside the ROI showing only the fuel assembly can be extracted, as has been done in Figure 1.

Testing this method on the available data, a visual inspection of the adjusted ROI for all the images shows that the rotation was consistently good, and the optimal rotation was easily identified from the maximum variance of the projection. However the ROI location was off for about 30 images of the total 2800. The reason for this is that some images were captured with the DCVD slightly away from the aligned position, hence the light intensity maximum is shifted away from the fuel center. This, in combination with the periodic structure visible in Figure 2, results in an ROI center placement that was off by one or two rows of fuel rods. Although the failures are sufficiently infrequent that they can be manually compensated for, a more robust method may still further improve these results, and may be required for inspector use.

## 4. FUEL TYPE CLASSIFICATION

For the automatic fuel type classification, a method based on Principal Component Analysis (PCA) was used. This technique has frequently been applied to face recognition [5], where human faces share several similarities features, yet also have unique distinguishing features. When applied to a set of images, PCA will determine an orthogonal basis of images, from which the original images can be reconstructed. The lowest principal components will explain the largest variance in the data, while the higher components explain features that contribute little to the variance. In this way, the original data can be represented using the first few principal components, with a low loss of data.

The PCA technique was first applied to determine the feasibility of using it for image classification in the context of DCVD images. To remove the possible dependence on the total intensity, and the individual relative intensity distributions in the images, the images obtained after the ROI adjustment step were rescaled to have the same lowest pixel value of zero, and a mean pixel value of one. The images had previously all been rescaled to 128x128 pixels, which allows all images to be treated the same by the PCA analysis. After rescaling, the PCA method in SciKit-learn [6] was used to decompose the image data in its principle components. The result is plotted in Figure 3, for the first three principal components for the image data.



**Figure 3: The first three principal components for the DCVD data. The fuel type is indicated by the color, and is the same as in Figure 1.**

As can be seen in Figure 3, several clusters emerge in the first three PCA components. This shows the feasibility of using the PCA components for fuel type classification as the data clusters are well separated in most cases. The PWR 17x17 assemblies (F and G, in figure 3 shown with two red colors) are well separated, from both each other and the other fuel types, and can easily be identified from the first three components alone. The PWR 16x16 assemblies with RCCA

inserts (A and B, in figure 3 shown with dark blue colors) also form a cluster, being rather well separated from the PWR 16x16 without insert. It is not obvious from only the first three components if they can be reliably separated from each other, however. For the PWR 16x16 assemblies without inserts (C, D and E), in Figure 3, shown with light blue, green and yellow colors, they are the most numerous, and in the PCA analysis the lowest PCA components can be expected to relate to differences between these fuel assemblies. Figure 3 shows that they form the largest cluster, where the three overlap significantly, suggesting that it may be challenging to separate and classify them, especially if the first PCA components are the most sensitive to these differences. This is consistent with a visual analysis of Figure 1, where those three assembly types are visually rather similar, and difficult to distinguish even for an experienced inspector.

Since fuel type analysis using PCA seems feasible, a  $k$  nearest neighbor classifier ( $k$ NN) was trained on the PCA components, to determine the possible classification accuracy. As a starting point, five neighbors in the  $k$ NN algorithm and 20 PCA components were used. A random selection of 70% of the data was used for training, and 30% for testing. The results are that a correct classification accuracy of 99.1% could be obtained, showing that the PCA components give sufficient data to allow reliable classification. The fuel assemblies that were most often misclassified were the PWR 16x16 without inserts (C, D and E), which were typically classified correctly as a PWR 16x16 without insert, but incorrectly within the fuel types C, D and E. The misclassified fuel assemblies were also typically on the edge of the cluster in Figure 3, where the data is sparse. This sparse region is also where misaligned images are located, as they become outliers in the PCA analysis, which increases the difficulty of correct classification.


As this is a feasibility study, a full hyperparameter optimization was not performed, though a limited optimization was done. For this set, using fewer than around 15-20 PCA components worsens the results, but using more does not improve it. For a larger data set with more fuel types to identify, it is expected that more PCA components will be required. The performance is better for a relatively low number of neighbors such as 3-5. However, this is mostly due to that the classification was hardest for the sparse part of the data, where there were very few training points available, and where there were outliers. For a larger data set, or if outliers can be eliminated by a better ROI adjustment, using more neighbors may be preferable. In total, for reasonable combinations of hyperparameters tried here, the  $k$ NN could correctly classify between 98.6% and 99.4% of the fuel assemblies in the data used.

## **5. TEMPLATE MATCHING**

Once the ROI is systematically placed and the fuel type known, template matching can be applied to the measurements for partial defect detection purposes. A partial defect in a spent fuel assembly will result in a change in the Cherenkov light intensity, where the change is most pronounced near the rods that were removed and/or replaced with non-radioactive substitutes. For a substantial partial defect (more than 30% of rods removed), this can be detected from the total intensity. However, for a small partial defect, such as a few well-spaced rods, the change will however be subtle, and difficult to detect by an inspector. An automated method of detecting such changes could potentially be more sensitive, and could highlight to the inspector where the inspector attention and expert judgement should be focused. To locate such local changes, one possibility is comparing the measurement with a template measurement for reference.

Ideally, to test the performance of the template matching methods, data from fuel assemblies with partial defects should be investigated. However, such measurements are scarce, and simulations of such assemblies are currently too time-consuming to enable a large set of images to be simulated. As a substitute to investigate the possible performance of template matching methods, we instead use a feature present in fuel assemblies of design A, having an RCCA insert. Out of the 250 measurements of fuel type A, about 50 fuel assemblies have RCCA inserts where one or two control rods have been removed, leaving a bright spot at the guide tubes at those locations. This is a significant local deviation from fuels with complete RCCA inserts, and was used to assess the possibility of using template matching to identify such local differences.

To create a template, the average image of the about 200 type A fuels without missing control rods was calculated. The average was calculated after the ROI had been adjusted, to ensure consistent positioning of the fuel assembly within the ROI. The resulting image can be seen in in Figure 4a) For each measurement of the ca 50 type A fuels with a missing control rod, the measured image was scaled to match the template with respect to minimum and mean intensity, to ensure that the two images had a similar intensity distribution. Next, the template was subtracted from the rescaled measurement. To highlight the regions with the most significant deviations in light intensity, a threshold was applied. The result after applying this procedure is exemplified by Figure 4c). The threshold value was manually selected here such that the missing control rod always passed the threshold, but no contiguous regions without know differences passed the threshold.



**Figure 4: a) The template for the type A fuel assembly measurements with all control rods present. b) An example type A fuel assembly where one control rod is missing. The red arrow indicate the position of the missing control rod. c) The difference between the images, after a threshold has been applied to highlight the pixels with the largest difference. The bright spot due to the missing control rod is indicated with a red arrow.**

The results presented in Figure 4 highlights how template matching can be used to identify differences in light intensity between images, but also shows challenges needs to be addressed before a reliable and robust method can be made available. As seen in Figure c), where the bright region due to the missing control rod is clearly visible, several more pixels also had a difference above the threshold, scattered through the image. Analyzing the performance of the template matching procedure to the type A fuels with a missing control rod, several reasons for this difference in the bright part could be identified, that needs to be solved before the template matching becomes reliable and robust:



- Small differences between the physical design of the assembly being measured and the template exists, resulting in that certain pixels in the bright regions are brighter in the measurement than in the template. Further splitting the Group A fuels according to design (such as done for PWR16x16, which was split into types C, D and E) would be required to mitigate this error, but may be challenging due to how subtle the differences are.
- Ripples on the water surface locally slightly distort the measurement, resulting in that a bright region in the measurement is compared to a dark region in the template. After ROI adjustment, the distortion is modest, but the boundary between a bright and a dark region can still move by one or two pixels, resulting in scattered pixels differing, as seen in Figure 4c). This is partly solved with the improved best-shot function in newer DCVD software, or with image stacking in the XCVD, which can reduce the effects of ripples.
- The measurement in Figure 4b) has a focus that resolves the spacers holding the rods in place to a higher degree as compared to the template. If a fuel assembly is measured under conditions where there is some turbulence in the water, a blurring occurs, and the level of blurring may differ between measurements, making the comparison more difficult.

## **6. CONCLUSIONS AND OUTLOOK**

This work has investigated the use of classical image analysis and multivariate analysis methods to analyse DCVD images. The results indicate that such methods can be applied to DCVD data to provide more information about the fuel being measured. The methods are not specific to the DCVD, only to measurements of Cherenkov light done from above. Hence, the methods can be applied also to future Cherenkov viewing systems, such as the XCVD and the RCVD.

The quantitative Cherenkov light analysis as well as the template matching requires that an ROI is placed to extract only the fuel assembly under study. While a good placement is easy to achieve manually, a systematically identical placement over all fuel assemblies can reduce uncertainties in the quantitative measurements, and ensure that template matching becomes more robust. The method developed is based on the light intensity of the fuel assembly, and the assumption that the fuel rods are in a square lattice, and that the measurements are done with the instrument well-aligned over the fuel, to ensure a consistent, symmetric light distribution in the ROI. While the method performs well on the data analyzed, the data was of high quality and consistently acquired. For data of lesser quality, a more robust ROI adjustment method may be required, and should be developed.

Using PCA analysis and a  $k$ NN classifier, it was shown that the fuel type can be classified accurately for the used data set. Misclassification primarily occurred for the few misaligned images present, and for fuel types that visually look very similar. Based on these results, it seems feasible to extend the classifier to more types of fuels, making it usable for any measurement campaign. For BWR assemblies, these tend to show a higher variability between designs, and should be easy to classify. For PWR assemblies, where design differences between manufacturers may be more subtle, the results show that a good classification accuracy can still be achieved. Automated fuel type classification can be used both to inform the inspector that the fuel under study is of the declared type, reducing the risk of errors should the wrong fuel be measured or should the fuel declarations be incorrect. It can also be used to split fuel assembly

groups that are declared as a single type but were not. However, a correct classification will require a database of all available fuel types (and RCCA inserts for PWR assemblies) for the PCA method to train on. The data made available through the action sheet is expected to cover most, but not all, fuel types that an inspector can encounter.

Using template matching, it is possible to identify regions where a measurement differ significantly from a template. However, the results presented here show that more development is required before a template comparison is robust and reliable. Based on the performance shown here, missing rods, leaving a noticeable bright region, could likely be identified. Further development is however required to detect substituted rods, as the differences will be smaller and localized to a smaller region. Additional data is also required to develop and benchmark such methods, and will have to be obtained through simulations, as there are few fuel assemblies with significant partial defects available for measurement.

While this work has focused on classical image analysis and multivariate analysis, the results show that the measurements contain sufficient information that more data can be extracted from them. This highlight the feasibility of the work done by the other participants of the action sheet, who focus on neural net based image analysis. Such methods can more systematically identify and utilize features of the images for the analysis, typically with significantly less supervision than the methods applied here. For partial defect detection, the situation will however remain challenging, as there is little measured data on fuel assemblies with partial defects that can be used for training a classifier. The data set will need augmentation with simulated data of fuel assemblies with partial defects, and such data could also be used to benchmark the methods presented in this work.

## **ACKNOWLEDGEMENTS**

This work was funded by the Swedish Radiation Safety Authority (SSM) under agreement SSM2022-1105. The authors would like to thank EURATOM for making the DCVD measurements available to the authors under Action Sheet 67, as well as Action Sheet 67 participants Pacific Northwest National Lab and Los Alamos National Lab for ongoing collaboration and support.

## **REFERENCES**

- [1] IAEA, Safeguards Techniques and Equipment: 2011 Edition. IAEA, 2011, IAEA/NVS/1/2011 (ISBN:978-92-0-118910-3).
- [2] Robotics Challenge Winning Design Helps Speed up Spent Fuel Verification, IAEA press release, 18 March 2019. Available at <https://www.iaea.org/newscenter/news/robotics-challenge-winning-design-helps-speed-up-spent-fuel-verification> retrieved 04 April 2023.
- [3] J. Chen, D. Parcey, A. Gerwing, P. Carlson, R. Kosierb, M. Larsson, K. Axell, J. Dahlberg, B. Lindberg, S. Jacobsson Svärd, and E. Sundkvist, "Partial defect detection in LWR spent fuel using a Digital Cerenkov Viewing Device," in Institute of Nuclear Materials Management 50th annual meeting Tucson, Arizona, 2009, pp. 12-16.
- [4] E. Branger, Zs. Elter, S. Grape, M. Preston, Studies of the impact of beta contributions on Cherenkov light emissions by spent nuclear fuel. The ESARDA Bulletin, vol. 64, 2022. Available at <https://doi.org/10.3011/ESARDA.IJNSNP.2022.1> retrieved 04 April 2023
- [5] Parinya Sanguansat, Principal Component Analysis. InTech, 2012. ISBN 978-953-51-0195-6
- [6] Scikit-learn: Machine Learning in Python. Pedregosa et. Al, JMLR 12, pp. 2825-2830, 2011.