Improved Online Localisation of CANDU Fuel Defects Using Ancillary Data Sources and Neural Networks

Christopher Wallace,[*,a] Curtis McEwan,[b] Graeme West,[a] William Aylward,[a] and Stephen McArthur[a]

[a] *University of Strathclyde, Department of Electronic & Electrical Engineering*
Glasgow, UK

[b] *Bruce Power*
Tiverton, Ontario, Canada

[*]Email: christopher.wallace@strath.ac.uk

Number of pages: 20
Number of tables: 0
Number of figures: 10

Abstract

This paper summarises a novel approach to improved localisation of fuel defects by fusing existing data sources and methods within a neural network model to make accurate and quantifiable identification earlier than existing processes. The approach is demonstrated through application to a CANDU reactor and utilises a small, manually labelled set of delayed neutron data augmented with neutronic power data, to train a neural network to estimate the probability of a fuel channel containing a defect. Results demonstrate that the model is often capable of identifying likely defects significantly earlier than existing methods and could support earlier decision making to enable a reduction in cost and time required to localise defects. The approach described has broader application to other reactor types given the general difficulty of detecting fuel defects via fission product measurement and the large quantities of ancillary parameters normally already recorded, which can be leveraged using machine learning techniques.

Keywords — Fuel Reliability, Condition Monitoring, Machine Learning, Decision Support

I. INTRODUCTION

I.A. Background

Monitoring of fuel integrity is an important aspect of safe operation of nuclear reactors. While the reliability of fuel assemblies used in modern reactors is typically very high and occurrence of defects relatively rare [1, 2, 3], the release of even small quantities of fission products or radioisotopes can have signficant impact on operation of plant. Leakage of radioactive material into coolant loops and potentially other areas of plant can result in elevated dose rates, particularly during maintenance or refuelling.

I.B. Overview

This paper introduces an approach of combining existing fuel defect detection data with ancillary data sources that can be used to augment, enhance and automate the localisation process. A case study is presented through the application of this approach for a method of defect localisation on CANDU reactors.

After discussion of the broader range of fuel defect detection and localisation technologies in the civil nuclear industry, this paper summarises the current state of the art in online CANDU fuel defect detection and identifies limitations of the existing approach. Improvements to the existing detection method are then proposed through the combination of neutronic bundle power estimates with the existing detector data in the form of a neural network which can estimate the probability that a channel contains defect fuel.

Testing results of the trained model are presented, demonstrating it to be capable of often making high confidence estimates of fuel defects signficantly earlier than existing methods. Finally the possible implementation of the model as a decision support tool to support existing fuel management is discussed.

The findings of the paper, in the context of fuel defect detection, have potential applicability for other reactor types and detection methods. The advent of powerful machine learning technologies in scenarios such as fuel defect detection enable the rapid alignment and analysis of multiple datasets and the identification of undiscovered or non-intuitive relationships which can be exploited to derive efficiencies, even where data is limited.

II. FUEL DEFECT DETECTION

Detection of fuel defects can be routinely performed during or after de-fuelling via sipping [4] using sealed enclosures or similar processes. On-line detection of fuel defects is typically more complex due to shielding, limited access to reactor internals and the high levels of background flux and fission product generation when at power. A range of techniques, influenced by respective reactor design, are used to detect fuel defects in an online scenario. Most Light-Water Reactors (LWR) and Pressurised Heavy Water Reactors (PWHR) employ some form of Fission Product (FP) detection system [3, 5, 6] however these systems are capable only of indicating the presence of a defect, not the location of the defect within the core. Figure 1 shows an example of the detection of fission products associated with a defect in the cladding of a fuel bundle in a CANada Deuterium Uranium (CANDU) reactor, which do not otherwise occur in, significant concentrations during normal operation.

II.A. Fuel Defect Localisation

There are very limited examples of dedicated systems available to locate fuel defects during online operation. Advanced Gas-Cooled Reactors employ a 'Burst Cartridge Detection' system [3, 7] which enables the detection of Beta activity from Kr and Xe within gas samples extracted from the primary circuit. This approach however requires a channel-by-channel search using a portable measurement device. In some LWR reactor types, localization of the defective fuel is possible by calculating the ratio of ^{134}Cs to ^{137}Cs during transient periods and correlating this with burnup[8].

Some Pressurised Heavy Water Reactors (PHWR), such as CANDU reactors, employ a technology known as 'Delayed Neutron Detection' [9] in order to identify delayed neutrons created through a known decay process that occurrs within fuel bundles. The presence of delayed neutrons indicates the presence of fission products and by making use of sample lines from fuel channels, and specifically measuring for shorter lived fission products, the location of a fuel defect can be determined

Methods have been proposed for embedding of particular signatures of noble gases within fuel assemblies [10] in order to enable identification of defects on detection of these signatures, or the use of gamma tomography[11] to identify characteristic gamma emmisions, however these technique

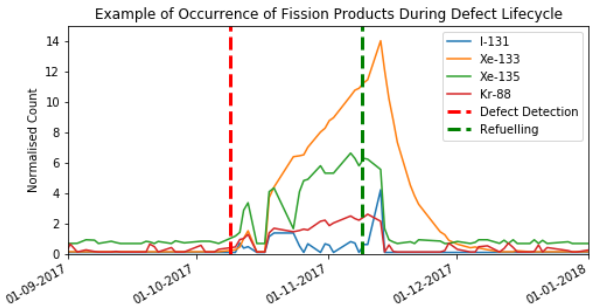

Fig. 1. An example of the detection of increased concentrations of particular fission products in the primary loop due to leakage from a fuel defect. The concentrations reduce when the defective fuel is removed from the core.

have not been attempted at scale on a commercial reactor. Other research has investigated a more sophisticated analysis of the generation of fission products [12] in order to detect fuel defects in an online setting.

III. CANDU DEFECT DETECTION AND LOCALISATION

III.A. Delayed Neutron Detection

Dedicated pipework was installed in some CANDU reactors at construction to enable sampling of primary loop coolant on a per-channel basis, designed to optimise a balanced measurement of delayed neutrons from ^{87}Br and ^{137}I. This configuration reduces as far as possible background noise from other activation products. The neutron count is recorded for each channel for a fixed period (typically 90 seconds) using a detector array, as shown in Figure 2. This method therefore enables the online detection of fuel defects. Feeder scanning provides an alternative off-line method for fuel defect localistion[13, 14].

The conceptual basis for the Delayed Neutron (DN) detection system is that a channel containing a fuel defect will emit an anomalously high number of delayed neutrons compared to other channels. The relatively low counting statistics and high levels of background count necessitate further analysis of the data in order to improve the sensitivity of the DN system.

Fig. 2. One row of a typical DN sensor array on a CANDU reactor, with the east face containing sample lines for half of the fuel channels with the remaining channels at the west face

III.B. Double Normalisation

The Double Normalisation (2XN) process [15] was introduced to address two factors which affect the measurement of delayed neutrons in order to detect fuel defects:

- Different levels of fuel burnup and activation of primary coolant circuit components introduces an inherent channel to channel variation in background count
- Background count levels during different measurements can vary to a level sufficient to mask anomalous changes in count

These issues are improved by performing a process of normalisation in two stages, specifically:

- **Scan Normalisation:** Each set of scan values (typically in sets of 30 channels at a time) are normalised by dividing the individual counts by the mean count for that scan set
- **Channel Normalisation:** Each delayed neutron count is divided by the historical mean count for that channel.

This can be given more formally, for a scan set S comprising neutron counts $[d_1, d_2 \dots d_n]$, a scan average $\bar{S}$ and historical averages $[\bar{d_1}, \bar{d_2} \dots \bar{d_n}]$ there exists an associated single, scan based, normalisation $S^{'}$ as:

$$S^{'} = [\frac{d_1}{\bar{S}}, \frac{d_2}{\bar{S}} \dots \frac{d_i}{\bar{S}}] = [d'_1, d'_2 \dots d'_i] \tag{1}$$

and a double normalised, scan-based and channel-based, normalisation $S^{''}$:

$$S^{''} = [\frac{d_1'}{\overline{d_1}}, \frac{d_2'}{\overline{d_2}} \dots \frac{d_n'}{\overline{d_n}}] = [d_1^{''}, d_2^{''} \dots d_n^{''}] \tag{2}$$

The 2XN data, an example of which is shown in Figure 3, is used as the primary method of detecting fuel defects. The DN data sets are typically grouped in 30 measurement (or channel) sets, based on the configuration of the detector array. It is expected that fuel defects will manifest as a sustained increase in double normalized count in one channel relative to other channels, in a manner that will be evident to the human expert. This process is subjective and depends on the prominence of the trend of defect 2XN data relative to the background of the other channels. In Figure 3 for example, an elevated 2XN count is observed for up to 4 scans before sufficient confidence could be assigned to the anomaly to schedule a refuelling however the defect was known to exist as early as 20 scans previously, as indicated by the red dashed vertical line in the lower plot.

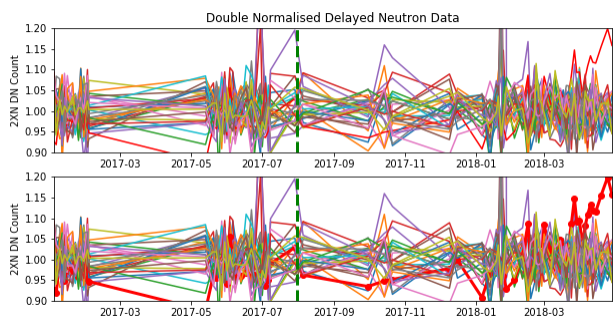

Fig. 3. The double normalised DN data for a fuel defect shown above as the engineer would observe it and below highlighting the defect channel, demonstrating the difficulty of detection, with the dashed vertical line indicating the initial detection of the defect the final data point indicating the localisation and refulleing.

III.C. DN Detection Issues

The process of scanning the DN system is time consuming and requires a human operator to activate and monitor the equipment. Scans are therefore taken only infrequently when there is no evidence of a fuel defect from fission product measurement (approximately weekly) and more frequently (approximately every 1-2 days) when a fuel defect is suspected to be present.

Further complicating the detection process, it is considered that sample tube fouling[12] and loss of detector sensitivity with ageing of plant are likely to be negatively affecting the detection of delayed neutrons.

III.D. Improved Defect Localisation

The existing fission product detection systems function well to detect the *occurrence* of defects (as shown in Figure 1) however the high levels of noise and limited data available via the DN system inherently limits the *localisation* of defects. A method is sought therefore to improve the ability to localise defects to a particular channel using existing data and equipment in order to provide decision support to operators. This method should, as far as possible, automate the detection process such that estimates of defect occurrence are quantifiable and remove as much subjectivity from the decision process as possible.

IV. AUTOMATED FUEL DEFECT CLASSIFICATION

IV.A. Ancillary Data Sources

It has been noted in discussions with CANDU operators that operating experience has recorded a possible relationship between the fuel power and occurrence of defects has been observed historically [16]. It has been proposed that significant changes in bundle or channel power which change the steam fraction within fuel channels could enable increased mobility of certain fission products from fuel cladding defects. This phenomenon has not been investigated or quantified however given the limits of the current 2XN analysis approach, it was proposed that neutronic power data may usefully augment the DN data should a suitable model be developed.

IV.B. Augmenting DN Data With Bundle Power Data

CANDU reactors periodically run simulations of neutronic power for each fuel bundle in the reactor core, based on fuel burnup and other operating parameters. These bundle power estimates are compared with thermal power measurements made via specific instrumented channels in order to confirm their validity. For this purpose the bundle powers are aggregated to determine a single channel power.

Fuel bundle power data for a CANDU reactor was aggregated to determine channel powers and aligned with DN data covering the same period in order to identify any apparent relationships. Figure 4 shows this data for approximately one year, however no relationship is immediately apparent.

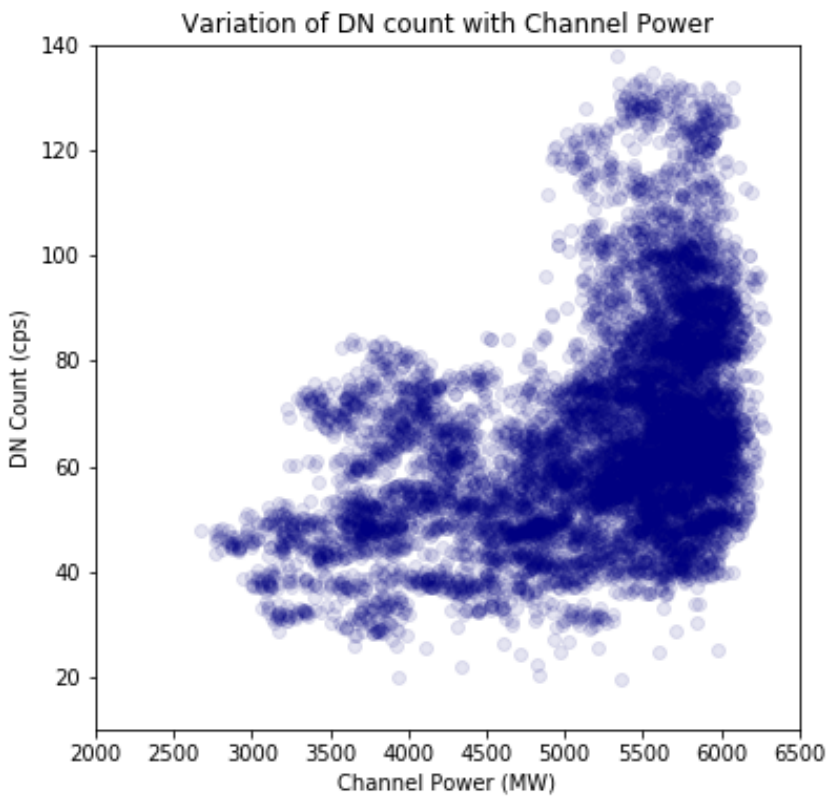

Fig. 4. Distribution of Channel Power with DN Count

The distribution shown includes all DN and channel power estimates over the period, including a number of fuel defects. If a relationship exists between the DN count and the bundle or channel power, it is possible that during normal operation the effect is masked by normal measurements from healthy fuel bundles. Figure 5 shows examples of two fuel defects from the relevant period,isolated to remove any masking effects of other channels, with measurements during the presence of a defect highlighted in red.

Once again however there is no obvious correlation between DN count and channel power

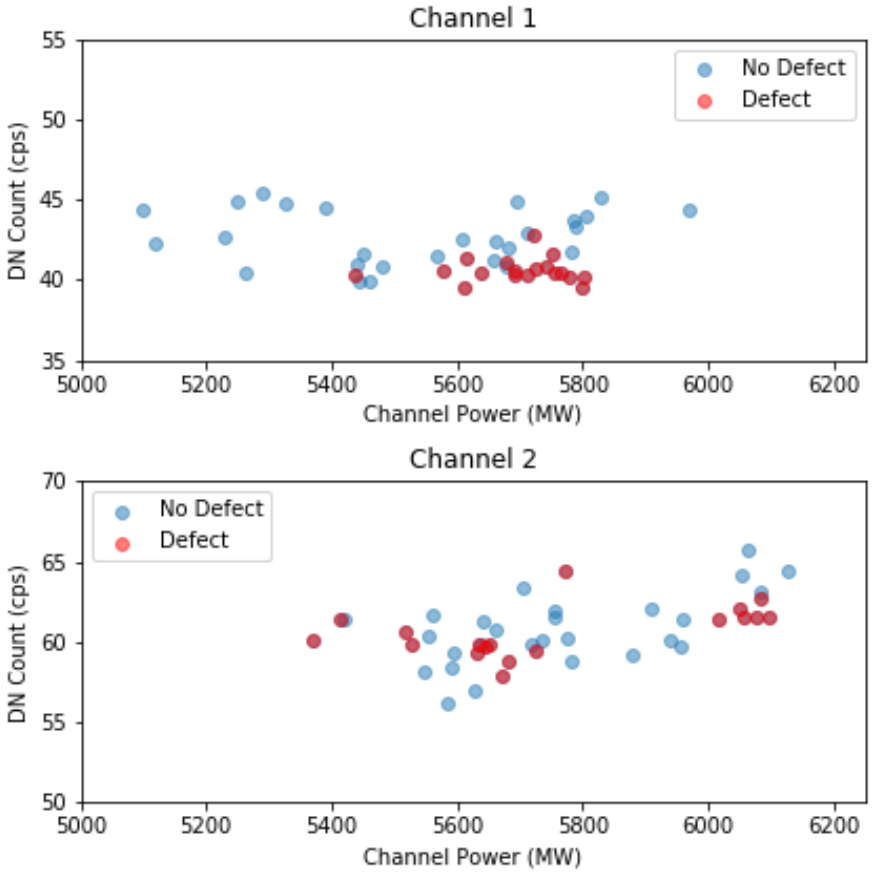

Fig. 5. Two examples of defect occurrences highlighting that absolute DN counts do not vary significantly outside normal envelopes of operation

either during normal operation or when a defect is present.

It is possible however that a more complex or non-linear relationship is present that is not easily visualised by such means. In order to test this hypothesis, a neural network model was constructed which combines DN count data and channel power data as input features and attempts to classify sequences of DN count and channel power as containing a defect or not.

IV.C. Data Structure and Features

Incorporating aspects of the existing process, specifically the identification of fuel defects by extended elevated 2XN value, and in order to attempt to replicate the visual detection process, it was determined that a time series of 2XN values would be used as an input feature to the model.

Experimentation with different window lengths, balancing the requirement for sufficient mea-

surements post-defect detection and the desire to identify defects as early as possible, determined that 8 DN scans was the optimum time series length for use in the model.

The corresponding channel power estimates that align with the DN scan dates (or where they do not align, the average of the adjacent channel power estimates) were also extracted for use an input features.

The input feature set for the model, $\mathbf{X}$, comprised of double normalised DN values d''_n and channel powers c_n is therefore given as:

$$\mathbf{X}=\begin{bmatrix} d''_1, & c_1 \\ d''_2, & c_2 \\ \dots & \dots \\ d''_8, & c_8 \end{bmatrix}$$

This data was extracted from the final eight scans prior to each confirmed defect in the labelled set. A visualisation of an example of the features, is shown in Figure 6, noting that the channel powers are also normalised in order to optimise training of the model. The output, $\mathbf{Y}$, consists of a vector of probability of the data containing a defect.

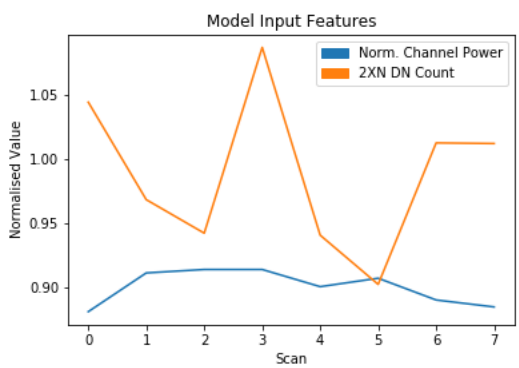

Fig. 6. Visualisation of a single instance of the model input features as DN and channel power time series

IV.D. Training and Testing Data

Given the relatively low occurence of fuel defects and the legacy equipment and software associated with the DN system, and therefore the overhead of gathering the relevant data, a small

labelled dataset was assembled. This data, comprising 30 defects, each of which was contained within a set of 30 measurements, was split in the following manner:

- **Training Data**: 600 labelled instances (including 20 defects)
- **Validation Data**: 150 labelled instances (including 5 defects)
- **Testing Data**: 150 labelled instances (including 5 defects)

It should be noted that the classes in this scenario are highly imbalanced due to the relatively low frequency of fuel defects in large populations however this is addressed during model training and accuracy interpretation.

IV.E. Neural Network Model

A recurrent neural network consisting of an input layer, two hidden layers and an output layer, as shown in Figure 7 with 16, 16, 10 and 1 units respectively using the ReLu [17] activation function. The ADAM [18] stochastic gradient descent algorithm was used to optimise a loss function based on the classification error, using the labelled training data. The model hyper-parameters were derived through testing and optimisation using random search of various numbers of layers and layer sizes. The use of a recurrent neural network architecture, specifically using Long Short-Term Memory units, enables the model to learn and exhibit temporal dynamic behavior.

In order to compensate for the significant class imbalance noted earlier, the loss function was configured to penalise false negative (missed defect) misclassifications more harshly than other errors.

Dropout layers between the fully connected hidden layers of the network were used to prevent over-fitting and to learn a more robust representation. The validation data was used used in order to tune the hyperparameters (e.g. layer number and size) and monitor training of the model but not directly to train the model.

IV.F. Testing and Training

The model was trained, as shown in Figure 8 until the validation loss function reached a minimum. The validation classification accuracy reached 100% after 500 epochs or complete passes of the dataset.

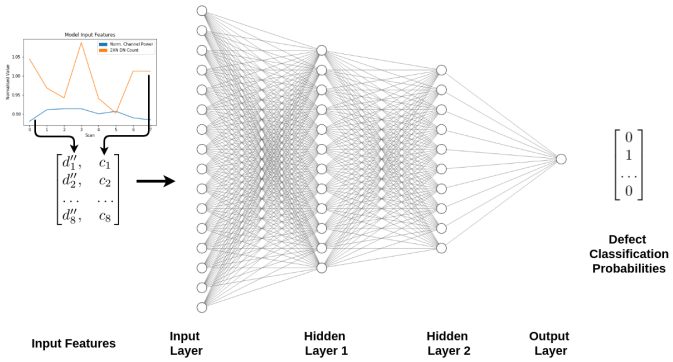

Fig. 7. The structure of the neural network model, with input features derived from 2XN DN count data and channel power on the left, and an output of a probability of defect, shown for multiple instances, on the right

The initial classification accuracy, at 96.5% is due to the class imbalance which would allow a model that classifies every instance of the data as 'defect free' to achieve this apparent performance, as the number of defects in the dataset comprised only around 3.5% of the total.

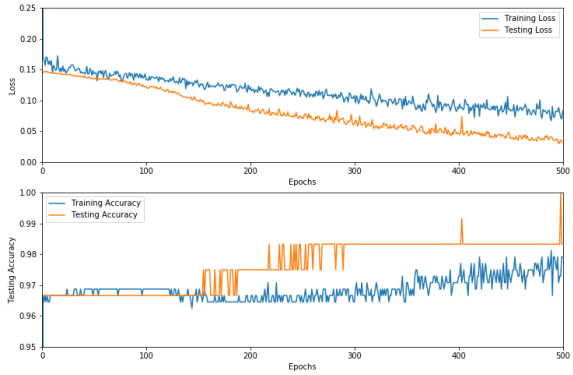

Fig. 8. Training of the model proceeds from a high level of accuracy given the class imbalances, improving from 96.5% to 100% on the validation set

It is important to note that the classification accuracy of 100% reflects only the ability of the model to correctly classify instances as 'defect' or 'non-defect' based on the final eight scans prior to the confirmation using the existing method. The classification accuracy derived from the model training does not directly indicate the ability of the model to identify defects earlier than the existing process.

V. RESULTS AND INTERPRETATION

In order to test the usefulness and accuracy of the model against previously unseen data, the test and validation datasets were processed by the model in order to generate predictions of defect probabilities.

Using the example shown in Figure 3, which was held in the test set of un-seen data, a series of predictions were made at each time step, using the window of previous DN and channel power data. These results, shown in Figure 9 illustrate the existing view of the 2XN process (with additional highlighting of the channel containing the defect) and the associated probability of a defect in the highlighted channel.

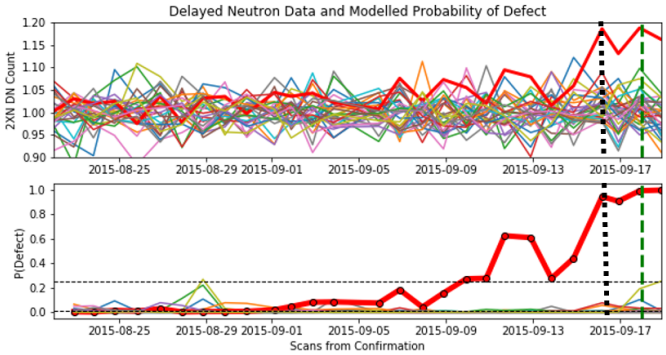

Fig. 9. Visualisation of the model results for unseen data. The defect channel is indicated in bold red in the upper and lower plots for convenience. The date at which the channel was first suspected by experts to contain a defect is shown by the vertical black dashed line and the model generated probability threshold is shown as a horizontal black dashed line. The date of refuelling is indicated by the green dashed vertical line

In order to provide to provide a quantifiable and comparable metric against which other datasets can be compared, it was determined through experimentation that a threshold of three standard deviations provides a level of significance against which the channel can be considered likely to contain a defect. This empirically derived threshold was found to work well on the available data. This threshold is shown as the black dashed lined in Figure 9.

It is envisioned that in application, a human expert would continue to utilise the 2XN process, while the probabilities generated by the model would be used for confirmatory purposes or to enable the identification of likely candidate channels. In the example shown in Figure 9, it can be seen that the model predicted a high probability of a defect several scans before the

existing 2XN process, potentially providing the opportunity to remove the defective fuel bundle earlier than would be achieved using the existing process.

V.A. Testing Results

In order to achieve a broader view of the performance of the model, this testing process was applied to the remaining testing and validation data. The benchmark for comparison, rather than the date at which the channel was refuelled, is the date at which the human expert identifed the channel as potentially containing a defect. Subsequent to this categorisation, these channels are watched with particular attention in order to confirm the suspected defect.

Using a set of results manually labelled by a human expert, it was possible, by comparison with the empirically derived threshold described earlier, to compare the difference in number of scans between human expert designation designation as 'potentially containing a defect' and the probabilities generated by the model.

Figure 10 illustrates the distribution of the number of days earlier than the existing process that the model was able to estimate the occurence of a defect with high probability as described above.

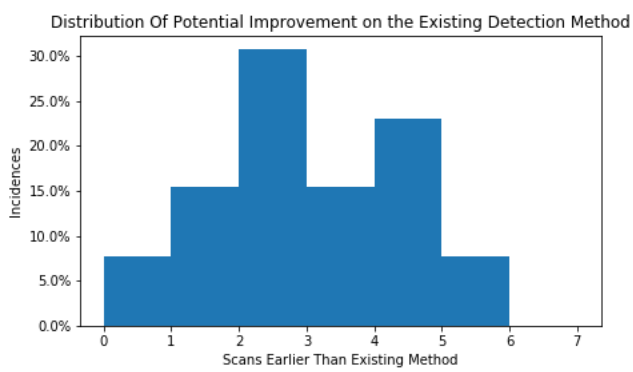

Fig. 10. The distribution of the number of scans earlier than the existing method with hich the model was able to indicate a significant probability of a defect.

These results demonstrate that in most cases the model can identify likely defects 2-4 scans earlier than the existing method and in some cases up to 6 scans earlier.

The results presented are required to be considered in the context of the decision making process undertaken by the engineer where a high degree of confidence is required in order to under-

take a refuelling. Given the relatively limited occurrence of the defects, and the correspondingly limited scope for back-testing of the model, it is likely that the most useful aspect of the model outputs is the downselection of the number of channels potentially containing a defect.

For example, even in test data where the model did not identify the correct channel until approximately the same point as the human expert, the channel had typically been in a small subset of 'most likely' channels for several scans prior to this point. Compared to the current process (Figure 3), where all channels are tightly grouped around the mean, this is already a significant advantage.

Further, it is possible that with testing on an larger number of defects, a more appropriate significance threshold could be identified. The results are therefore indicative of performance given a particular interpretation, rather than a definitive comparison of performance.

This view of the use of the model and comparisons of results reinforces the concept of decision support in this work, rather than an attempt to directly automate the analysis process.

VI. CONCLUSIONS

This work has presented a novel method of improving localisation speed of fuel defects for the particular case of a CANDU reactor. While this work has been developed on a relatively small dataset, it has highlighted the significant scope for time and cost saving possible from improved decision support. Further, the principles and structure of the work described have applicability for other reactor types, even beyond the specific case of fuel defect detection. The ability to leverage existing datasets using machine learning technologies to identify and exploit previously unobserved or un-intuitive relationships to drive efficiencies and understanding has broad scope, particularly in the nuclear industry where ageing plant and IT infrastructure have generally not considered such possibilities from design and construction.

VI.A. Future Work

Work is ongoing for the case presented here, specifically to extract a larger dataset through subsampling and to investigate innovative ways of artificially increasing the size of the dataset using data augmentation[19].

It is unclear as to the cause of the variation in time between detection and identification of

some defects compared to others however it is likely that this is due to a combination of different areas or extents of fuel cladding damage and the variation in local channel conditions. Work to improve the understanding of the relationship between fission product release rate, DN count and bundle power should improve this understanding and may even enable detection of fuel defects on a per-bundle basis, using records not yet available during this work.

ACKNOWLEDGMENT

This work was funded and supported by Bruce Power.

Template creation supported by the Department of Energy Computational Science Graduate Fellowship, provided under grant number DE-FG02-97ER25308.

REFERENCES

[1] K. Edsinger, J. Deshon, A. Kucuk, E. Mader, B. Cheng, S. Yagnik, and R. Daum, "Recent US Fuel Reliability Experience," *TOP Fuel, WWRFPM 2009* (2009)URL http://inis.iaea.org/Search/search.aspx?orig_q=RN:40094823.

[2] R. YANG, B. CHENG, J. DESHON, K. EDSINGER, and O. OZER, "Fuel R & D to Improve Fuel Reliability," *Journal of Nuclear Science and Technology*, **43**, *9*, 951 (2006); 10.1080/18811248.2006.9711181., URL https://www.tandfonline.com/doi/abs/10.1080/18811248.2006.9711181.

[3] "Management of Damaged Spent Nuclear Fuel," No. NF-T-3.6, International Atomic Energy Agency, IAEA Nuclear Energy Series.

[4] J.-Y. Park, M.-S. Shim, and J.-H. Lee, "Current Status of Integrity Assessment By Sipping System of Spent Fuel Bundles Irradiated in CANDU Reactors," *Nuclear Engineering and Technology*, **46**, *6*, 875 (2014); 10.5516/NET.09.2014.018., URL http://www.sciencedirect.com/science/article/pii/S173857331530156X.

[5] S. D. Kreider and A. Schneider, "The detection of failed fuel in LWRs [light water reactors] - a historical review," *Transactions of the American Nuclear Society*, **61 p. 46-47** (1990)URL http://inis.iaea.org/Search/search.aspx?orig_q=RN:22034865.

[6] B. J. Lewis, "Fundamental aspects of defective nuclear fuel behaviour and fission product release," *Journal of Nuclear Materials*, **160**, *2-3*, 201 (1988); 10.1016/0022-3115(88)90049-9., URL https://www.sciencedirect.com/science/article/pii/0022311588900499.

[7] D. K. Cartwright, "Burst Cartridge Detection In UK Power Reactors," *pp 175-212 of Detection and Location of Failed Fuel Elements. Vienna, International Atomic Energy Agency, 1968.* (1970)URL https://www.osti.gov/biblio/4753674.

[8] R. Beraha, G. Beuken, G. Frejaville, C. Leuthrot, and Y. Musante, "Fuel Survey in the Light Water Reactors Based on the Activity of the Fission Products," *Nuclear Technology*, **49**, *3*, 426 (1980); 10.13182/NT80-A17690., URL https://doi.org/10.13182/NT80-A17690.

[9] J. J. Lipsett and C. M. Tseng, "Delayed-Neutron Systems for Failed Fuel Location in CANDU Reactors," *IEEE Transactions on Nuclear Science*, **23**, *1*, 325 (1976); 10.1109/TNS.1976.4328265.

[10] B. Pollack, B. J. Lewis, and D. Kelly, "Viability Assessment of Noble Gas Bundle Tagging for Failed-Fuel Identification in CANDU Reactors," *Nuclear Technology*, **182**, *1*, 39 (2013); 10.13182/NT13-A15824., URL https://www.tandfonline.com/doi/full/10.13182/NT13-A15824.

[11] S. Holcombe, S. Jacobsson Svrd, and L. Hallstadius, "A Novel gamma emission tomography instrument for enhanced fuel characterization capabilities within the OECD Halden Reactor Project," *Annals of Nuclear Energy*, **85**, 837 (2015); 10.1016/j.anucene.2015.06.043., URL http://www.sciencedirect.com/science/article/pii/S0306454915003588.

[12] B. J. Lewis, P. K. Chan, A. El-Jaby, F. C. Iglesias, and A. Fitchett, "Fission product release modelling for application of fuel-failure monitoring and detection - An overview," *Journal of Nuclear Materials*, **489**, 64 (2017); 10.1016/j.jnucmat.2017.03.037., URL http://www.sciencedirect.com/science/article/pii/S0022311517302362.

[13] J. J. Lipsett and W. B. Stewart, "Failed Fuel Location in CANDU-PHW Reactors Using a Feeder Scanning Technique," *IEEE Transactions on Nuclear Science*, **23**, *1*, 321 (1976); 10.1109/TNS.1976.4328264.

[14] A. M. Manzer, "In-core assessment of defective fuel in CANDU-600 reactors," (1985)URL http://inis.iaea.org/Search/search.aspx?orig_q=RN:17017413.

[15] H. Schmeing and C. Hogg, "The Bruce B delayed neutron (DN) monitoring system for failed fuel detection: twenty years of operating experience, recent improvements, and planned upgrades," *25th Annual CNS Conference* (2004).

[16] M. Dobrean and E. Sorin, "Pickering NGS B - Unit 6 strategy to remove fuel defects during operation at high power," (2010).

[17] V. Nair and G. E. Hinton, "Rectified Linear Units Improve Restricted Boltzmann Machines," *Proceedings of the 27 th International Conference on Machine Learning*, 8 (2010).

[18] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *arXiv:1412.6980 [cs]* (2014)URL http://arxiv.org/abs/1412.6980, arXiv: 1412.6980.

[19] D. A. v. Dyk and X.-L. Meng, "The Art of Data Augmentation," *Journal of Computational and Graphical Statistics*, **10**, *1*, 1 (2001); 10.1198/10618600152418584., URL https://doi.org/10.1198/10618600152418584.