

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Investigation of a Methodology for the Detection of Diversions in Spent Nuclear Fuel

Moad S. Al-dbissi

DEPARTMENT OF PHYSICS
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2024

Investigation of a Methodology for the Detection of Diversions in Spent Nuclear Fuel.
MOAD S. AL-DBISSI
ISBN 978-91-7905-987-3

Acknowledgements, dedications, and similar personal statements in this thesis, reflect the author's own views.

© Moad S. Al-dbissi, 2024.

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr 5453
ISSN 0346-718X

Division of Subatomic, High Energy and Plasma Physics
Department of Physics
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Telephone +46(0)31-772 1000

Cover:

A schematic showing the overall methodology of using ANNs for the identification of partial defects in SNF assemblies based on the measurements of the neutron flux and its gradient in the empty guide tubes of the assembly.

Chalmers Digital Print
Göteborg, Sweden 2024

Investigation of a Methodology for the Detection of Diversions in Spent Nuclear Fuel
MOAD AL-DBISSI
Division of Subatomic, High Energy and Plasma Physics
Department of Physics
Chalmers University of Technology

Abstract

One of the main tasks in nuclear safeguards is the inspection of Spent Nuclear Fuel (SNF) to detect possible diversions of their special nuclear material content, e.g., ^{235}U and ^{239}Pu . These inspections verify the declared SNF via passive measurements of characteristic signatures such as the emissions of neutrons and gamma rays. The current PhD research investigates different aspects for the development of a novel non-intrusive methodology that can enhance safeguards inspections of SNF assemblies, and it includes two main parts. In the first part, simulations are performed to evaluate the feasibility of measuring the neutron flux and its gradient inside the empty guide tubes of a SNF assembly with a miniaturized detector made of an array of optical fiber-based neutron scintillators. In addition, experiments are carried out to characterize these types of neutron scintillators. The results of this preparatory work show that neutron flux gradient measurements in SNF assemblies may be a viable option and provide insights for the construction of a prototype of a detector for the purpose. In the second part of the research, the application of machine learning models based on Artificial Neural Networks (ANNs) is studied to process measured SNF signatures and reconstruct the arrangement of the fuel pins in an assembly. The objective of this part is two-fold. On one hand, ANN models are explored for the task of determining possible diversion patterns from SNF signatures collected inside the accessible guide tubes. On the other hand, the advantage of using the neutron flux gradient as input to the algorithm is evaluated. The training and testing of the ANN models are performed with synthetic datasets generated from Monte-Carlo simulations of a typical PWR SNF assembly, considering the intact configuration and different degrees and patterns of diversion. The results show that the models effectively predict diversions and characterize most of them to a good extent. In addition, the use of the neutron flux gradient, which is not analyzed during standard inspections, is proven to be advantageous.

KEYWORDS: nuclear safeguards, spent nuclear fuel, neutron scintillator, flux gradient detector, machine learning, artificial neural networks

List of Publications

This thesis is based on the work contained in the following papers.

Paper I

Aldbissi, M., Vinai, P., Borella, A., Rossa, R., Pázsit, I. (2022). Conceptual design and initial evaluation of a neutron flux gradient detector. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 1026, 166030. <https://doi.org/10.1016/j.nima.2021.166030>.

Paper II

Aldbissi, M., Vinai, P., Borella, A., Rossa, R., and Pázsit, I. (2022). Evaluation of the performance of a neutron gradient detector for partial defect testing in spent nuclear fuel assemblies. INMM 63rd annual meeting, August 2022. <https://resources.inmm.org/annual-meeting-proceedings/evaluation-performance-neutron-gradient-detector-partial-defect-testing>.

Paper III

Borella, A., Aldbissi, M., Boogers, E., Pázsit, I., Rossa, R., Van Der Meer, K., Vinai, P., Wagemans, J. (2022). Neutron and gamma-rays sensitivity of new fiber-mounted scintillation neutron detectors for spent fuel pin diversion measurements. IAEA Symposium on International Safeguards: Reflecting on the Past and Anticipating the Future, Vienna, Austria, October 31 – November 4, 2022.

Paper IV

Aldbissi, M., Borella, A., Pázsit, I., Rossa, R., Vinai, P. (2022). Optimizing neural networks to detect replaced spent fuel pins using the partial defect tester. IAEA Symposium on International Safeguards: Reflecting on the Past and Anticipating the Future, Vienna, Austria, October 31 – November 4, 2022.

Paper V

Aldbissi, M., Rossa, R., Borella, A., Pázsit, I., Vinai, P. (2023). Identification of diversions in spent PWR fuel assemblies by PDET signatures using Artificial Neural Networks (ANNs). Annals of Nuclear Energy, 193, 110005. <https://doi.org/10.1016/j.anucene.2023.110005>.

Paper VI

Aldbissi, M., Vinai, P., Rossa, R., Borella, A., Pázsit, I. (2023). Using machine learning for the detection of missing fuel pins in spent nuclear fuel assemblies based on measurements of the gradient of the neutron flux. INMM/ESARDA Joint Annual meeting, May 2023. <https://resources.inmm.org/annual-meeting-proceedings/using-machine-learning-detection-missing-fuel-pins-spent-nuclear-fuel>.

Paper VII

Aldbissi, M., Pázsit, I., Rossa, R., Borella, A., Vinai, P. (2024). On the use of neutron flux gradient with ANNs for the detection of diverted spent nuclear fuel. (Submitted to Annals of Nuclear Energy).

The author's contribution to the included papers

Paper I: Aldbissi, M. developed all the Monte-Carlo models for the gradient detector and the hypothetical test setup, performed all the simulation work, and generated the plots and the tables included in the paper. The development of the theoretical concepts, the interpretation of the results, and the writing of the paper was done in collaboration with the co-authors.

Paper II: Aldbissi, M. developed all the Monte-Carlo models for the spent nuclear fuel assemblies, performed all the simulation work, drafted the first version and generated the plots and the tables included in the paper. The development of the theoretical concepts, the interpretation of the results, and the writing and reviewing of the paper was done in collaboration with the co-authors.

Paper III: Aldbissi, M. participated in the preparation of the fiber-based scintillators, setting up and performing the experiments and contributed to the calculations, interpretation of the results and the review of the paper.

Paper IV: Aldbissi, M. developed all the Artificial Neural Network (ANN) models, performed all the optimization, training and testing of the ANNs, drafted the first version and generated the plots and the tables included in the paper. The development of the theoretical concepts, the interpretation of the results, and the writing and reviewing of the paper was done in collaboration with the co-authors.

Paper V: Aldbissi, M. developed all the ANN models, performed all the optimization, training and testing of the ANNs, drafted the first version and generated the plots and the tables included in the paper. The development of the theoretical concepts, the interpretation of the results, and the writing and reviewing of the paper was done in collaboration with the co-authors.

Paper VI: Aldbissi, M. generated the dataset of spent nuclear fuel assemblies via Monte-Carlo simulations, developed all the ANN models, performed all the optimization, training and testing of the ANNs, drafted the first version and generated the plots and the tables included in the paper. The development of the theoretical concepts, the interpretation of the results, and the writing and reviewing of the paper was done in collaboration with the co-authors.

Paper VII: Aldbissi, M. generated the dataset of spent nuclear fuel assemblies via Monte-Carlo simulations, developed all the ANN models, performed all the optimization, training and testing of the ANNs, drafted the first version and generated the plots and the tables included in the paper. The development of the theoretical concepts, the interpretation of the results, and the writing and reviewing of the paper was done in collaboration with the co-authors.

Acknowledgments

I would like to express my deepest gratitude to my main supervisor Assoc. Prof. Paolo Vinai for his continuous guidance, unconditional support and positive encouragement throughout the years of this PhD program and most importantly for his valuable scientific contributions to this research.

I would like to thank my co-supervisors, Prof. Imre Pázsit at Chalmers University of Technology, and Dr. Alessandro Borella and Dr. Riccardo Rossa at The Belgian Nuclear Research Centre (SCK CEN), for all their guidance and support and for all their contributions to the project.

I would also like to thank Prof. Christophe Demazière for covering the role of the examiner of this PhD program and for his contributions in reviewing this thesis.

Special thanks go to Michèle Coeck for facilitating the collaboration with SCK CEN, and for Klaas van der Meer and Catrinel Turcanu for hosting me at SCK CEN for half the period of my PhD studies.

The experimental work at Chalmers University of Technology was made possible by the help and effort of Anders Nordlund, Lennart Norberg and Lasse Urholm, and at SCK CEN by the help of Eric Boogers, Luc Borms and Bart Van Houdt.

We are grateful to Prof. Tsuyoshi Misawa, Prof. Yasunori Kitamura and Dr. Yoshiyuki Takahashi at the Institute for Integrated Radiation and Nuclear Science, Kyoto University (KURNS), for providing us with the two optical fiber-based neutron scintillators.

Special thanks also go to Prof. Anna Erikson for accepting the task of being the opponent of the PhD defense, to Dr. Matthew Eaton, Prof. Tsuyoshi Misawa, Assoc. Prof. Roumiana Chakarova, and Prof. Giovanni Volpe for agreeing to be in the grading committee.

The computational work was enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC) at the Chalmers Centre for Computational Science and Engineering (C3SE), and at the National Supercomputer Centre (NSC), Linköping University (LiU), partially funded by the Swedish Research Council through grant agreement no. 2018-05973.

The project was financially supported by SCK CEN under grant agreement PO4500047684, the Swedish Radiation Safety Authority under agreements SSM2021-709, SSM2021-786 and SSM2023-4389, the Ringhals power plant and from the Euratom research and training program 2014-2018 under grant agreement No 847594 (ARIEL).

I would also like to express my gratitude to my family for their continuous love, support and encouragement during the years of my studies.

List of Acronyms

AdaGrad	Adaptive Gradient.
AdaM	Adaptive Moment.
AI	Artificial Intelligence.
ANNs	Artificial Neural Networks.
BU	Burn Up.
BWR	Boiling Water Reactor.
CF	Calibration Factor.
CT	Cooling Time.
DA	Destructive Assay.
DAS	Data Acquisition System.
DCVD	Digital Cherenkov Viewing Device.
DTs	Decision Trees.
FD	Fork Detector.
IAEA	International Atomic Energy Agency.
IE	Initial Enrichment.
kNNs	k-Nearest Neighbors.
KURNS	Kyoto University Institute for Integrated Radiation and Nuclear Science.
LLNL	Lawrence Livermore National Laboratory.
LWR	Light Water Reactor.
MBAs	Material Balance Areas.
NDA	Non-Destructive Assay.
NGD	Neutron Gradient Detector.
NMAC	Nuclear Material Accountancy and Control.
NPT	Nuclear non-Proliferation Treaty.
PDET	Partial Defect Tester.
PGET	Passive Gamma Emission Tomography.
PM	Photo-Multiplier.
PWR	Pressurized Water Reactor.
ReLU	Rectified Linear Unit.
SNF	Spent Nuclear Fuel.
UN	United Nations.

Contents

1	Introduction	1
1.1	Nuclear safeguards	1
1.2	Spent nuclear fuel	2
1.3	Objectives of the research	5
1.4	Structure of the thesis	6
2	Neutron gradient detector	7
2.1	Conceptual design	7
2.2	Quantitative analysis	9
2.3	Evaluation in SNF assemblies	13
2.4	Experimental testing	15
2.4.1	Experiment at the hot-cell laboratory	16
2.4.2	Experiment at the BR1 research reactor	17
2.4.3	Experiment at the LNK facility	20
3	ANNs for the identification of diverted SNF	23
3.1	Artificial Neural Network (ANNs)	24
3.2	Dataset with PDET signatures	25
3.3	Percentage of diverted fuel using PDET signatures	27
3.3.1	Comparison of machine learning algorithms	27
3.3.2	Optimized ANN model	29
3.4	Identification of diversion patterns using PDET signatures	31
3.5	Identification of diversion patterns using NGD signatures	36
3.5.1	Dataset with the neutron flux gradient	36
3.5.2	The gradient as input feature	37
3.5.3	Refinement of the analysis	41
4	Conclusions	47
4.1	Summary of the research	47
4.2	Future work	50
4.2.1	Neutron gradient detector	50
4.2.2	Artificial Neural Networks (ANNs)	50
4.2.3	Training dataset	51
4.3	Ethical considerations	51
4.3.1	Data privacy	51

4.3.2	Latent bias	51
4.3.3	Explainability	52
4.3.4	Proliferation risks	52
4.3.5	Human oversight	52
	References	53
	Papers I - VII	57

List of Figures

1.2.1 Schematics of a close-up of a 17x17 PWR fuel assembly.	3
1.2.2 Typical composition of LWR fresh fuel (a) and spent fuel after in-core irradiation (b).	4
2.1.1 Conceptual design of the Neutron Gradient Detector (NGD).	8
2.2.1 The spatial dependence of the radial component of the gradient with and without the presence of the detector.	10
2.2.2 Estimated direction of the gradient vector at different locations along the x -axis.	11
2.2.3 Estimated flux gradient with imperfect efficiencies, positions and rotations during the calibration	12
2.3.1 Serpent model of a PWR SNF assembly with the NGD placed inside the guide tubes.	13
2.3.2 Intact fuel assembly; the scaled thermal flux distribution (left) and the scaled gradient measurements in the guide tubes (right).	14
2.3.3 Example of a diversion scenario where a number of fuel pins are replaced with stainless-steel dummy pins.	14
2.3.4 Diverted fuel assembly; the scaled thermal flux distribution (left) and the scaled gradient measurements in the guide tubes (right).	15
2.4.1 A LiF-ZnS(Ag) fiber-based scintillator.	16
2.4.2 Set-up of the experiment in the hot-cell laboratory at Chalmers.	16
2.4.3 Schematic of the BR1 research reactor [1].	17
2.4.4 Aluminum rods used to insert the scintillators in the reactor cavity.	18
2.4.5 Neutron count rates from scintillators A and B with respect to the reactor power.	19
3.1.1 Schematic diagram of a feed-forward artificial neural network	24
3.2.1 Examples of diversion scenarios included in the dataset.	26
3.3.1 Confusion matrices based on the comparison between the ANN, DT and kNN models.	28
3.3.2 Confusion matrix of the optimized ANN model.	29
3.3.3 Relative differences between probabilities associated with true and wrongly predicted class labels.	30
3.3.4 Diversion patterns predicted as intact by the ANN model (false negatives).	31
3.3.5 Number of misclassifications with respect to BU and IE.	31
3.4.1 Probability distribution of the correctly classified fuel pins.	33
3.4.2 Probability distribution of the misclassified fuel pins.	33

3.4.3 Examples of misclassified diversion scenarios with low number of incorrect fuel pins.	35
3.5.1 General form of the confusion matrix used to characterize the predictions of the ANN.	37
3.5.2 Example of a diversion scenario for which the results of the models are consistent with the global trends of Table 3.4.	39
3.5.3 Example of a diversion scenario for which the results of the models deviate from the global trends of Table 3.4.	40
3.5.4 Confusion matrix of the model N (a), confusion matrix of the model $N+G_m+G_d$ (b), and the confusion matrix of the model $N+G_x+G_y$ (c).	42
3.5.5 Example A of a diversion case from the dataset and how it was reconstructed by the three ANN models.	43
3.5.6 Simulated responses of the thermal neutron flux and the gradient of the neutron flux from an intact fuel assembly and the diversion pattern from example A. . .	44
3.5.7 Example B of a diversion case (left) and its reconstruction via model $N+G_x+G_y$ (right).	44
3.5.8 Simulated responses of the thermal neutron flux and the gradient of the neutron flux from example B.	45
3.5.9 Additional diversion patterns for the updated dataset.	45
3.5.10 Example B (left) and its reconstruction via updated model $N+G_x+G_y$ (right). .	45

List of Tables

2.4.1 Mean value and standard deviation of the neutron count rates from the scintillators.	17
2.4.2 Thermal neutron flux at the center of the cavity, at different power levels.	19
2.4.3 Sensitivities and relative efficiencies of scintillators A and B, at different power levels.	20
2.4.4 Gamma-ray sensitivity at different thresholds for the two scintillators.	21
3.2.1 Values of BU, IE and CT included in the PDET dataset.	27
3.3.1 Percentage of diverted fuel pins and prescribed class labels.	27
3.4.1 Number of fuel assemblies with x incorrect fuel pins.	34
3.5.1 Performance metrics with respect to the different combinations of input features.	38
3.5.2 Optimized performance metrics for the three ANN models: N, N+G _m +G _d , and N+G _x +G _y	41

Chapter 1

Introduction

Nuclear technology presents a spectrum of societal benefits, encompassing both its capacity to serve as a reliable and low-carbon energy source and its potential to foster technological innovation. The advancements driven by nuclear technology impact several fields such as energy production, medicine, materials science and engineering, and industrial processes, thereby underscoring its multifaceted role in the sustainable development of our society.

As an energy source, nuclear power plants contribute substantially to meeting electricity demands and offering a stable base-load supply that complements intermittent renewable sources. One aspect of their attractiveness is related to the reduced greenhouse gas emissions, so they can aid in mitigating climate change. Presently, nuclear energy is utilized in over 30 countries covering approximately 10% of the global electricity supply through the operation of 411 nuclear power plants worldwide. In addition, 58 reactors are currently under construction [2].

However, the potentially perilous ramifications stemming from any misapplication of nuclear technology, e.g., for the construction of nuclear weapons, remain indisputable. It is noteworthy that the origins of nuclear technology are rooted in weaponry development, and it was only subsequently that its potential for civilian utilization and power generation became apparent. The dual nature of this technology, which can be applied for both peaceful and military purposes emphasises the importance of creating robust safeguards to prevent nuclear proliferation and ensure the safe and secure use of nuclear materials.

In this chapter a historical background and an introduction to nuclear safeguards are given in section 1.1, a description of spent nuclear fuel and the importance of safeguarding it are discussed in section 1.2, the objectives of the doctoral research are introduced in section 1.3, and the structure of the thesis is provided in section 1.4.

1.1 Nuclear safeguards

In the aftermaths of the atomic bombing of Hiroshima and Nagasaki during the second world war in 1945, there was a growing awareness of the need to control and monitor nuclear materials to prevent their misuse. The creation of The United Nations (UN) in the same year highlighted the importance of international cooperation in managing nuclear technology. In 1953, the Atoms

For Peace initiative was introduced, advocating the peaceful use of nuclear technology while controlling its proliferation. In 1957, the International Atomic Energy Agency (IAEA) was established as an autonomous agency under the UN to ensure the safe and secure application of nuclear technology [3]. Later, the Nuclear Nonproliferation Treaty (NPT) came into force in 1970, which became a cornerstone of the nuclear nonproliferation regime [4]. It aimed to prevent the spread of nuclear weapons and promote nuclear disarmament.

International treaties and agreements such as the NPT are one of the main categories that constitute a state's legal and institutional framework regarding nuclear nonproliferation [5]. In order to assure that states honor their international obligations, the NPT established a system of safeguards to verify compliance with its provisions, overseen by the IAEA. This system of safeguards comprises a set of technical and non-technical measures that are applied in nuclear facilities in order to detect any misuse of nuclear materials or technology either by the state itself or by other actors (e.g., terrorist groups). The main objective of safeguards is to account for nuclear materials at all times at the facility and to detect potential diversions in a timely manner. The IAEA approach to safeguards is based on a system of Nuclear Material Accountancy and Control (NMAC) complemented by containment, surveillance, and monitoring measures. Nuclear material accountancy deals, amongst others, with the registration of nuclear materials and a system of bookkeeping checks and balances that provides an accurate record of the stocks and movements of nuclear materials in time and space [6].

Nuclear facilities are commonly divided into two types from a safeguards application point of view, i.e., item facilities and bulk facilities. In item facilities, all nuclear materials are kept in item form and the integrity of the item remains unaltered during its time at the facility [7]. In this case, the IAEA safeguards are based on item accountancy procedures, Non-Destructive Assay (NDA) measurements, and verification of the integrity of the items [7]. Examples of item facilities are most power reactors, critical assemblies, and storage installations for spent fuel. On the other hand, bulk facilities are facilities where nuclear materials are held, processed, or used in bulk forms [7]. Bulk facilities may be organized for safeguards purposes into multiple Material Balance Areas (MBAs) by separating activities related to the storage and assembly of discrete items from the ones involving storage or processing of bulk materials [7]. In such facilities, flow and inventory values declared by the operator are verified by the IAEA via independent measurements and observations. Examples of bulk handling facilities are plants for conversion, enrichment, fuel fabrication, spent fuel reprocessing, and storage facilities for bulk materials.

1.2 Spent nuclear fuel

Currently, the majority of nuclear power reactors in operation are Light Water Reactors (LWRs), i.e, Pressurized Water Reactors (PWRs) and Boiling Water Reactors (BWRs). As of 2022, out of 411 operational nuclear reactors, 301 (73.2%) are PWRs and 42 (10.2%) are BWRs [2]. LWRs are item facilities in which the fuel typically consists of uranium dioxide (UO_2) pellets sealed within zirconium alloy tubes to form fuel rods. The fuel rods are then clustered in assemblies which are arranged in the reactor core. The assemblies can differ in size, shape and design depending on the reactor type. Figure 1.2.1 shows a schematic of a standard 17x17 nuclear fuel assembly used for PWRs. The reactor core of a PWR contains 150 to 200 fuel assemblies and the one of a BWR contains 400 to 800 fuel assemblies. Approximately one-third of these assemblies

are replaced every 1-2 years according to the refueling process. This results in a large number of discharged assemblies, the so-called Spent Nuclear Fuel (SNF) assemblies, along the operating lifetime of a plant.




Figure 1.2.1: Schematics of a close-up of a 17x17 PWR fuel assembly.

LWRs are thermal systems, i.e., most of the fission reactions that generate energy are induced by thermal neutrons. The uranium used in the fresh fuel for these reactors is enriched so that the concentration of the fissile isotope ^{235}U is increased from the natural level of 0.7% up to 3-5%, see Figure 1.2.2 (a). This provides the necessary amount of ^{235}U to sustain the nuclear fission chain reaction in a LWR core. In a fission reaction, the ^{235}U nuclei is split by a neutron releasing, together with a substantial amount of energy (~ 200 MeV), two (or more) lighter nuclei known as fission products which are highly unstable and radioactive, and on average 2-3 additional neutrons which can serve to continue the nuclear chain reaction. As the fuel "burns" in the core during operation, the ^{235}U is consumed and both fission products and minor actinides accumulate, see Figure 1.2.2 (b). Therefore, the ability of the fuel to sustain the fission chain reaction diminishes over time. When the fuel is depleted, it is removed from the reactor and replaced with new fresh fuel, a process known as "refueling". The Spent Nuclear Fuel (SNF) is then stored underwater in cooling pools located at the facility to dissipate its decay heat and protect against radiation release. Depending on the strategy for handling SNF, the assemblies are later moved into an interim storage which can be either wet (bigger cooling pools than the ones on site) or dry (concrete casks). Eventually, the spent fuel is transferred to a reprocessing facility (closed fuel cycle) or is encapsulated in copper canisters for a permanent repository disposal (open fuel cycle) [8].




Figure 1.2.2: Typical composition of LWR fresh fuel (a) and spent fuel after in-core irradiation (b).

SNF is particularly sensitive from a safeguards perspective because of its residual fissile material content such as ^{235}U and ^{239}Pu , where the residual ^{235}U is the remnant of the initial enrichment while the residual ^{239}Pu is the byproduct of neutron capture reactions in ^{238}U . According to the requirements from the IAEA, while in wet storage, SNF assemblies from LWRs are regularly inspected to verify that their content corresponds to the declaration provided by the power utilities. Each SNF assembly has a specification as a fresh assembly from the manufacturer, an irradiation history, and an associated burn-up and cooling time from the power utility, which allow to estimate the isotopic content of the assembly. Verifying the total fissile content of an assembly is a very time and resource demanding task which can only be performed by Destructive Assay (DA) methods. Therefore, for practical reasons, it suffices to rely on NDA techniques to ensure that no fuel material has been illicitly removed from the assembly. In the field of nuclear safeguards, any statistically significant deviations between the declared amount of nuclear materials and the amount determined by the verification measurements are known as "Defects". Three levels of defects must be considered according to the IAEA [7]:

- "Gross Defect" which refers to a defect in an item, e.g., a fuel assembly, that has been completely falsified to the maximum extent possible so that all or most of the declared material is missing.
- "Partial Defect" which refers to a defect in an item that has been falsified to such an extent that some fraction of the declared amount of material is still present.
- "Bias Defect" which refers to a defect in an item that has been slightly falsified so that only a small fraction of the declared amount of material is missing.

Currently, several NDA techniques are used to detect defects in SNF assemblies. The Digital Cherenkov Viewing Device (DCVD) [9], the Fork Detector (FD) [10] and the Passive Gamma

Emission Tomography (PGET) [11] are examples of techniques that are approved for inspections by the IAEA and have been extensively applied for many years [12]. These techniques rely on the passive measurement of characteristic signatures of the spent fuel such as the emissions of neutrons, gamma rays, or Cherenkov light. The neutrons emitted in SNF mainly come from the spontaneous fission of minor actinides such as ^{242}Cm , ^{244}Cm and ^{252}Cf [13]. The gamma radiation is due to the radioactive decay of fission products such as ^{134}Cs and ^{137}Cs . The emissions of neutrons and gamma rays play a crucial role in the characterisation of spent nuclear fuel since their levels depend on the specific fuel composition and irradiation history.

The interpretation of the measurements performed with NDA techniques are mainly based on data analysis and statistical methods that are carried out by the inspectors. In addition, the analysis usually aims at the detection of coarse diversions of nuclear material. Some techniques, e.g., the PGET, can provide pin-level resolution, however they are laborious since they require the movement of the fuel assemblies from their storage positions in the SNF pools.

1.3 Objectives of the research

The main objective of the PhD research presented in this thesis is to investigate different aspects in support of the development of a novel non-intrusive methodology that can enhance safeguards inspections of SNF assemblies from LWRs and the detection of partial defects.

The overall concept of the methodology consists of two steps. In the first step, the thermal neutron flux and its gradient are measured via miniaturized angular-sensitive detectors that can be inserted in the empty guide tubes of the SNF assembly. Following this procedure, the assembly does not need to be moved from its storage location. In the second step, an algorithmic processing of the measurements is used to characterize possible anomalies in terms of their extent and their location. The approach has the potential of reducing the amount of expert judgement required for the interpretation of the measurements and providing more detailed estimations, and therefore it can facilitate the decision process of the safeguards inspectors.

To enable the measurements of the neutron flux and its gradient inside the SNF assembly, simulations are performed in this PhD research to evaluate the feasibility of detecting the neutron flux simultaneously at different locations inside the empty guide tubes of a SNF assembly with a miniaturized detector made of an array of optical fiber-based neutron scintillators, and thus deriving the local neutron flux gradient. In addition, experimental work is carried out to characterize these types of neutron scintillators and obtain insights for the construction of a prototype of the detector.

For the processing of the measured SNF signatures to retrieve the arrangement of the fuel pins in the assembly, the application of machine learning models based on Artificial Neural Networks (ANNs) is studied. The objective of this part of the PhD research is two-fold. On the one hand, ANN models are explored for the task of detecting the fuel pins and determining possible diversion patterns via SNF signatures collected at different points inside the system (corresponding to the accessible guide tubes). On the other hand, the advantage of providing information on the neutron flux gradient as input feature to the algorithm is evaluated.

1.4 Structure of the thesis

The structure is as follows. Chapter 2 summarizes Papers I, II and III and is arranged in two parts. The first part discusses the conceptual design and the Monte-Carlo evaluation of a miniaturized gradient detector suitable for the methodology under investigation and made of an array of tiny optical fiber-based neutron scintillators. The second part of Chapter 2 describes the experimental work that was conducted at Chalmers and at SCK CEN to evaluate optical fiber-based neutron scintillators for the gradient detector. Chapter 3 summarizes Papers IV, V, VI and VII and it concerns the development, training and testing of ANN models for the identification of diversions in SNF, using two different synthetic datasets, namely a dataset of simulated measurements of neutron flux and gamma emission rates (provided by SCK CEN) and a dataset of simulated measurements of neutron flux and its gradient (developed in the context of this project). Chapter 4 provides conclusions, an outlook for the continuation of the work, and the ethical considerations related to the research.

Chapter 2

Neutron gradient detector

The feasibility of using the neutron current or the gradient of the neutron flux for localisation problems in nuclear reactor cores (e.g., finding the position of a static neutron source, the tip of a partially inserted control rod, a vibrating fuel assembly, or a vibrating control rod) has been demonstrated in [14, 15]. Methods for measuring them have also been investigated in, e.g., [16, 17]. The neutron current and the gradient are related to each other and contain more information compared to the scalar flux. Consequently, methods based on such angularly-sensitive quantities have the potential for a more accurate detection of missing or replaced fuel pins in SNF assemblies. The research presented in this thesis is focused only on the gradient since it can be measured more easily.

A dedicated detector can be constructed with several optical fiber-based scintillators with small neutron sensitive volumes. These types of scintillators were first developed and successfully tested in [18, 19], and efforts have been made to further study them in [20]. Recent work has also shown the suitability of using a cluster of these scintillators for high-resolution neutron flux measurements and for the characterization of highly localized gradients [21].

In this chapter, preparatory work is carried out to study the proposed detector and its feasibility for the verification of SNF. The conceptual design, which combines 4 optical fiber-based neutron scintillators is presented in section 2.1. The performance of the detector is investigated via Monte-Carlo simulations, first in a water tank with a neutron source, see section 2.2, and then inside the guide tubes of a simulated SNF assembly, see section 2.3. Experiments are also conducted to characterize the optical fiber-based neutron scintillators, see section 2.4.

2.1 Conceptual design

The scalar neutron flux represents the density or flow rate of neutrons per unit area per unit time and it has a spatial distribution. The gradient on the other hand provides information on how this distribution changes with respect to space. The gradient is a vector that points in the direction of the steepest increase of the neutron flux and whose magnitude indicates the rate of change. In a three-dimensional space, each component of the vector corresponds to the rate of change of the neutron flux with respect to one of the spatial coordinates (x , y , z).

For the application at hand, the fuel pins in SNF assemblies are assumed to be either intact, fully removed, or replaced with dummy pins (mock-up fuel pins made with appropriate surrogate materials without any fissile content), so only their radial position in the horizontal plane is of interest and the problem is considered two-dimensional. The gradient can then be evaluated at some fixed axial elevation where only the x and y Cartesian components are needed.

Accordingly, a detector can be constructed from several small optical fiber-based neutron scintillators to measure the scalar flux in several positions concurrently over a two-dimensional plane and thus determine the gradient of the neutron flux. There exist several different options for small fiber-based scintillators in terms of the neutron converter and scintillation material (LiCaF, boron loaded plastic scintillator, etc.). We restrict the present study to the type of detectors which we have at hands and which were also used in previous works [20], namely LiF as neutron converter and ZnS(Ag) as the scintillation material. The scintillators are mainly sensitive to neutrons in the thermal energy range (~ 0.025 eV).

The diameter of the individual fiber-based scintillators can be as small as about 1 mm, however, the diameter of the gradient detector will be inevitably larger. By aiming at performing measurements within a PWR fuel assembly, one can use the instrumentation guide tubes of the assembly, which are about 1 cm in diameter.

A design of a Neutron Gradient Detector (NGD), with the mentioned size limitation, is proposed as follows. Four optical fiber-based scintillators are mounted in a cylindrical aluminum holder, that has a diameter of 1 cm and a height of 5 cm, according to a rectangular pattern, see Figure 2.1.1. Aluminum is chosen for the cylinder because of its easy manufacturing properties, and low neutron absorption cross section. The fiber-based scintillators are 1 mm in diameter and their tip (3 mm in height) is covered with a LiF/ZnS(Ag) mixture that acts as the converter/scintillation material. The fibers are coated with a thin layer of Teflon for protection against external light.




Figure 2.1.1: Conceptual design of the Neutron Gradient Detector (NGD).

The complete detector containing the four scintillators can be inserted into an instrumentation guide tube and moved to a suitable axial position. Then the two detector pairs at diagonally opposite positions, perpendicular to each other, can be used to measure the x and y components of the flux gradient. The Cartesian components can later be used to calculate the magnitude and direction of the gradient vector.

2.2 Quantitative analysis

Whereas it is intuitively clear that, in theory, the detector design described in subsection 2.1 is suitable to determine the flux gradient, it is useful to assess its performance by detailed simulations. For this purpose, the code Serpent is used [22]. Serpent is a multi-purpose three-dimensional continuous-energy Monte-Carlo particle transport code developed at VTT, the Technical Research Centre of Finland. The code is designed for traditional reactor physics applications, for multi-physics reactor calculations, and for neutron and photon transport calculations in radiation, fusion and medical physics problems. Serpent also includes numerical capabilities that allow parallel computing on clusters and multi-core workstations.

To assess the performance of the detector via Serpent, a hypothetical test case is considered. The test case was chosen to be similar to that used in earlier works, i.e., a neutron source in a water tank [17, 16]. The reason is partly that it is a simple setup, with an azimuthally symmetric flux distribution in the horizontal plane, in which the results can be easily interpreted. And partly, because such an experiment will be possible to replicate in the future, when the detector will actually be fabricated. The test case consists of a cylindrical Aluminum tank 1 m in height and 1 m in diameter filled with water, with a ^{252}Cf source, 2 cm in diameter, in the middle.

The first goal of the quantitative analysis is to investigate how the presence of the detector affects the accuracy of the estimation of the gradient. Similarly to the case of the ordinary neutron detectors, the presence of the detector might alter the neutron flux distribution. The consequences of such a flux disturbance are usually not significant when measuring the scalar flux. In the current application, the gradient is derived from the difference between the neutron flux values that are measured by the four scintillators, which are placed relatively close to each other. The potential distortion introduced with the four scintillators might have a bigger impact on the determination of the flux gradient than on the scalar flux. In addition, a systematic underestimation of the gradient might arise from a self-shielding effect, i.e., the scintillators at the higher flux position might shield against the neutron current pointing to the scintillators at the lower flux position.

The strategy to quantify these effects relies on two sets of simulations. The first simulation does not include the detector and the "unperturbed" thermal neutron flux is calculated in hypothetical measurement positions where the gradient detector can be inserted. In the second step, simulations are made for the case that the detector occupies the positions previously selected, one at a time, and the reaction rates in the scintillators are calculated. The gradient obtained from the difference of the reaction rates of the diagonally opposite scintillator pairs is then compared with the gradient of the neutron flux obtained without the detector.

The comparison between the gradient from the unperturbed flux and the gradient from the reaction rates is not completely trivial. At any single measurement point, the magnitude of the




Figure 2.2.1: The spatial dependence of the radial component of the gradient with and without the presence of the detector.

gradient will be quantitatively different for the flux and the reaction rate, since they correspond to physically different quantities. Nevertheless, the two are proportional to a scaling factor which can be considered as a constant for a given energy distribution. The scaling factor does not depend on the actual value of the gradient, and hence on the measurement position. If the space dependence of the two gradients is proportional to a constant scaling factor, then it is a demonstration of the equivalence between the two gradients and the negligible effect of the presence of the detector.

The magnitude of the neutron flux gradient without the detector and that obtained from the reaction rates estimated in the scintillators of the detector are compared at different distances from the neutron source, see Figure 2.2.1. The constant scaling factor between the two curves was obtained by the least squares method and used for scaling the gradient from the unperturbed flux. The space dependence of the two gradients are very close, indicating that the distortion effect of the proposed detector design is negligible.

The suitability of the detector to estimate the direction of the gradient vector is also investigated. A general case with non-zero components of the gradient is considered, i.e., the detector is positioned such that none of the two scintillator pairs lie on a radial line from the neutron source. Several simulations were performed with the detector at different positions along the x axis. The results from the calculated reaction rates show that the direction angle of the gradient is estimated correctly, see Figure 2.2.2 together with Figure 2.2.1. One advantage of the direction of the gradient is that it does not require any normalisation, i.e., it can be directly compared to




Figure 2.2.2: Estimated direction of the gradient vector at different locations along the x -axis.

the expected (true) value of the direction.

In the simulations so far, it was assumed that all four scintillators are equal and thus have the same efficiency. A quantitative analysis of the effect of different sensitivities of the scintillators, and the methods for correcting them, are also of interest.

For this application, the absolute efficiencies of the four scintillators are not needed, only the efficiencies relative to each other, and a method to compensate for them. As mentioned earlier, the flux gradient is determined to a constant scaling factor, whose value is not of interest. However, since scintillators with different efficiencies do not measure the same neutron flux values at a specific position, changes in the orientation of the scintillators might affect the scaling factor. The purpose of the correction is thus to make sure that a constant scaling is preserved, irrespective of the orientation of the detector. The correction method used in the current study amounts to an in-situ calibration of the relative efficiencies, which can be even performed in a field measurement, and has to be executed only once.

In the case of a detector with four scintillators of varying efficiencies, the correction method is based on rotating the detector by 90° , 180° and 270° from its original orientation. Accordingly, each scintillator will occupy each of the four angular positions once. The average value of the four reaction rates at each measurement position is then calculated, and the two components of the gradient are determined by taking the difference of the average reaction rates in the diagonally opposing positions. Such a correction method will lead to an unbiased estimate of the gradient and to the relative efficiencies of the four scintillators. The relative efficiencies can be used to correct the measurements in other points, without the need for additional rotations




Figure 2.2.3: Estimated flux gradient with imperfect efficiencies, positions and rotations during the calibration

of the detector.

An illustrative example of these results is shown in Figure 2.2.3. The detector is placed at 3 cm from the neutron source and the reconstruction of the direction of the gradient is compared between three scenarios. In the first scenario (represented by the blue arrow) the four scintillators have the exact same material composition and the detector is placed accurately at 30° as its initial orientation. In the second scenario (red arrow) the four scintillators have the same material composition but an error is assigned to the ideal orientation angle, i.e., the detector is placed at 27° . In the third scenario (yellow arrow) the four scintillators have different efficiencies (different atomic fractions of ^6Li) and non-perfect rotations were performed during the calibration process. Both the magnitude and the direction of the gradient vector can be reconstructed despite the varying efficiencies of the scintillators and the uncertainty in the initial positioning of the detector or in the rotation angles during the calibration process.

The results from the quantitative analysis of the gradient detector via Monte-Carlo simulations are discussed in more details in Paper I in the appendix.

2.3 Evaluation in SNF assemblies

The next step in assessing the performance of the NGD is to study the response of the detector when placed inside SNF assemblies. The objective here is to evaluate the changes in the simulated measurements of the gradient because of possible diversions as compared to an intact fuel assembly, and thus to have a better understanding on how these changes can be used to localise the diversions.

For this reason a model of a 17x17 PWR SNF assembly is developed using the Monte-Carlo code Serpent. The assembly consists of 264 fuel pins, with Zircaloy cladding and an initial enrichment of 3.5 w%, and 25 empty guide tubes where the NGD can be placed, see Figure 2.3.1.




Figure 2.3.1: Serpent model of a PWR SNF assembly with the NGD placed inside the guide tubes.

The Serpent simulations were performed in two-steps. The first step is a burn-up simulation of the declared fresh fuel assembly, which consists of an irradiation cycle that continues until a final burn-up value of 40 MWd/kgU is achieved, followed by a decay cycle that replicates a cooling time of 5 years in the spent nuclear fuel pool. The second step is a fixed-source simulation that is performed with the fuel composition obtained from the burn-up simulation, which is then distributed consistently with the diversion patterns of interest, in order to estimate the thermal neutron flux and its gradient in the guide tubes of the assembly. The fixed-source simulation is performed using 5×10^9 neutron histories to balance the achieved statistical uncertainty and the computational time. The statistical error of the Monte-Carlo calculation of the thermal neutron flux is 0.1% on average.

The spatial distribution of the thermal flux and the gradient calculated in the guide tubes of an intact fuel assembly with no diversion are plotted in Figure 2.3.2. The values of the flux and the gradient are independently scaled by dividing by the maximum value. The gradient is constructed using the simulated neutron flux from the scintillators of the detectors in the guide tubes. The two pairs of scintillators at diagonally opposite positions, perpendicular to each other, are used to estimate the two Cartesian components of the gradient, respectively. For practical reasons, instead of the two Cartesian components, the magnitude (absolute value) and the direction of the gradient vector are plotted because they are physical concepts that are easier to interpret. For an intact assembly, the largest absolute values of the gradient are found at the guide tubes close to the four corners of the assembly and the lowest value (close to zero) is in the central guide tube. The direction of the gradient vector points from lower to higher values of the thermal neutron flux inside the guide tubes.




Figure 2.3.2: Intact fuel assembly; the scaled thermal flux distribution (left) and the scaled gradient measurements in the guide tubes (right).

It is then expected that a diversion of fuel pins (partial defects) affects the radial distribution of the thermal neutron flux in the assembly and therefore the magnitude and direction of the gradient of the neutron flux. To illustrate such an effect, an example of a diversion scenario is created by replacing two complete columns of fuel pins (34 in total) with stainless-steel dummy pins at the left-most side of the system, see Figure 2.3.3. The scaled spatial distribution of the thermal flux and the scaled gradient calculated in the guide tubes of the diverted fuel assembly are shown in Figure 2.3.4. The thermal flux has a significant decrease in the region of diversion because the removal of the fuel pins reduces the neutron emission in that region. Consequently, significant variations can be observed both in magnitude and direction of the gradient, especially in the guide tubes closest to the diversion.




Figure 2.3.3: Example of a diversion scenario where a number of fuel pins are replaced with stainless-steel dummy pins.




Figure 2.3.4: Diverted fuel assembly; the scaled thermal flux distribution (left) and the scaled gradient measurements in the guide tubes (right).

The gradient of the neutron flux, which provides the direction in which the neutron flux changes and the rate of such a change, has richer information than the scalar neutron flux. Based on the deviations in the gradient measurements between the intact case (used as reference) and a diversion case, relevant details can be retrieved to help identify and localize missing fuel pins in SNF assemblies. The results regarding the evaluation of the NGD in SNF assemblies via Monte-Carlo are discussed in more details in Paper II in the appendix.

2.4 Experimental testing

Two optical fiber-based neutron scintillators, on which the concept of the NGD is based, were available during the course of this PhD project, courtesy of Kyoto University Institute for Integrated Radiation and Nuclear Science (KURNS), Japan. The experimental testing and characterization of these scintillators is a first step towards the future construction of the NGD.

Each of the two scintillators consists of a ~ 2 m long and ~ 1 mm thin plastic optical fiber whose tip is covered with a LiF-ZnS(Ag) material as shown in Figure 2.4.1. The tip of the fiber acts as the neutron sensitive part, where LiF is the neutron converter according to the reaction



and ZnS(Ag) is the scintillation material. The two products of the reaction in Eq. (2.4.1), an α particle and a tritium atom, interact with the ZnS inorganic scintillator grains mixed into the same matrix, and a scintillation light (photons) is produced. The scintillation light is the result of the de-excitation of the luminescence centers in the ZnS molecules [23]. The generated photons travel through the single optical fiber which is coupled to a Photo-Multiplier (PM) tube where the photons are converted into electrons. The neutron detection system including the scintillator, optical fiber and PM tube, is then connected to a Data Acquisition System (DAS).




Figure 2.4.1: A LiF-ZnS(Ag) fiber-based scintillator.

2.4.1 Experiment at the hot-cell laboratory

The first set of experiments were performed in the hot-cell laboratory at the Department of Physics, Chalmers University of Technology. The sensitive parts of the scintillators were placed one at a time near an Americium-Beryllium (Am-Be) neutron source surrounded by polyethylene plates, see Figure 2.4.2. The Am-Be source has an activity of 5 curie and an emission rate of 1.1×10^7 n/s and is stored in a small steel canister (3 cm in radius and 6 cm in height) for safe handling. The neutron source was surrounded by polyethylene plates from all sides to act as a moderator for slowing down the neutrons emitted from the source. The thermalisation of the neutrons emitted from the source is needed since the detectors are mainly sensitive to neutrons in the thermal energy range.




Figure 2.4.2: Set-up of the experiment in the hot-cell laboratory at Chalmers.

Ten measurements were performed with each scintillator, with a measurement time of 10 s each. The neutron counts were recorded after each measurement, and the mean value and the standard deviation of the neutron count rates for the two scintillators (denoted as A and B, respectively) are reported in Table 2.4.1. Scintillator A tends to provide a slightly higher neutron count rate compared to scintillator B. These kind of deviations are expected, since the neutron-sensitive tips and the coupling between the optical fibers and the PM tubes were made by hand. Nevertheless, the two scintillators are proven to provide relatively close count rates for thermal neutrons.

Table 2.4.1: Mean value and standard deviation of the neutron count rates from the scintillators.

Scintillator	Count rate (/10 s)	
	Mean value	Standard deviation
A	808.5	18.5
B	791.3	16.8

2.4.2 Experiment at the BR1 research reactor

Further characterization of the two scintillators was carried out in the BR1 research reactor at the Belgian Nuclear Research Centre SCK CEN [1] in order to determine their sensitivity to thermal neutrons, their calibration factors and their relative efficiencies.

At the top of the BR1, a spherical cavity with a radius of 50 cm is available for irradiation experiments and calibration of detection instruments under a well-defined Maxwellian thermal neutron flux, see Figure 2.4.3. The neutron flux at the center of the cavity is constantly monitored using a calibrated fission chamber. The two scintillators were inserted into the cavity one at a time using customized aluminum rods as the one shown in Figure 2.4.4. The fibers were taped to the inside of the aluminum rods with their tips at the bottom of the rod. The rods were then inserted into the cavity from the top. When the rod is fully inserted, its bottom is located exactly at the centre of the cavity.




Figure 2.4.3: Schematic of the BR1 research reactor [1].




Figure 2.4.4: Aluminum rods used to insert the scintillators in the reactor cavity.

At first, background measurements were performed with both scintillators while inserted in the cavity with the reactor being turned-off. The background count rate was negligible for both scintillators. The reactor was then turned-on and the neutron count rate was recorded for each scintillator, at the centre of the cavity, at different reactor power levels. The measurements were performed after the reactor reached criticality at each power level and the measurement time was chosen to be 10 minutes each. Figure 2.4.5 shows how the neutron count rate from each scintillator increases with increasing the power level of the reactor. As already observed in the experiments in the hot-cell laboratory at Chalmers University of Technology, scintillator A tends to provide higher neutron count rates compared to scintillator B.

The conventional thermal neutron flux in the center of the cavity can be obtained by multiplying the corrected count rate of the monitor fission chamber with a calibration factor (CF), i.e.:

$$CF = \frac{\phi_{Th}}{N_{FC}} = (2.60 \pm 0.03)10^4 (cm^{-2}) \quad (2.4.2)$$

where, ϕ_{Th} is the thermal neutron flux in $(cm^{-2}s^{-1})$ and N_{FC} is the count rate from the monitor fission chamber in (s^{-1}) .




Figure 2.4.5: Neutron count rates from scintillators A and B with respect to the reactor power.

The count rates from the fission chamber and the correction factor were provided by the reactor operators and the values of the thermal neutron flux at each power level are listed in Table 2.4.2. The ratio between the count rate from the neutron scintillators and the conventional thermal flux at the center of the cavity represents the sensitivity of the scintillators (S_i) in (cm^2):

$$S_i = \frac{N_i}{\phi_{Th}}, \quad i = A \text{ or } B \quad (2.4.3)$$

where, N_i is the count rate in (s^{-1}) for scintillators A or B.

Ideally this ratio should be the same at all power levels. Table 2.4.3 shows that the sensitivity values for each scintillator are close to each other with an average level of uncertainty of 1-2%. The sensitivities given in Table 2.4.3 can serve as a calibration factor, similar to the one from the calibrated fission chamber, so that for each power level the thermal neutron flux at one point can be obtained from the neutron count rate of the scintillators at the same point.

Table 2.4.2: Thermal neutron flux at the center of the cavity, at different power levels.

Power (kW)	ϕ_{Th} ($10^5 \text{cm}^{-2} \text{s}^{-1}$)
0.1	1.04(2)
0.3	2.86(5)
0.7	6.76(9)
1	9.62(13)
3	29.12(39)

Table 2.4.3: Sensitivities and relative efficiencies of scintillators A and B, at different power levels.

Power (kW)	S_A (10^{-3} cm^2)	S_B (10^{-3} cm^2)	E_r (%)
0.1	1.15(2)	1.08(2)	94(2)
0.3	1.30(2)	1.19(2)	91(2)
0.7	1.25(2)	1.19(2)	95(2)
1	1.24(2)	1.16(2)	94(2)
3	1.23(2)	1.14(1)	92(2)

The relative efficiency (E_r) in (cm^2) of scintillator B compared to that of scintillator A is calculated as:

$$E_r = \frac{S_B}{S_A} = \frac{N_B}{N_A} \quad (2.4.4)$$

where, S_B and S_A are the sensitivities in (cm^2) of scintillators B and A, respectively.

The values of the relative efficiency in percentage at different power levels are included in Table 2.4.3 and are very close to each other. Consistently with the results in the experiment carried out at Chalmers University of Technology and with the count rates of Figure 2.4.5, scintillator A shows a higher efficiency.

2.4.3 Experiment at the LNK facility

Detectors used within spent nuclear fuel assemblies are exposed to both a neutron flux and a gamma dose rate. Therefore, it is important to test the sensitivity of the scintillators, optical fibers and PM tubes to gamma rays as well. If the whole detection system is not properly shielded, gamma rays may cause undesired contributions in terms of photons and thus an inaccurate estimation of the neutron flux. If the sensitivity to gamma rays is correctly identified, the detector could be calibrated accordingly by setting a proper discrimination threshold to discard the contribution of the gamma rays in the recorded count rate. To investigate this aspect, a third set of experiments was performed in the Laboratory for Nuclear Calibrations (LNK) at SCK CEN.

In these experiments, the two scintillators were exposed to gamma sources, namely ^{137}Cs and ^{60}Co , with different dose rates that varied in a range between 5 mGy/h and 192 Gy/h. The scintillators did not show any response to gamma rays for dose rates below ~ 70 Gy/h. A sensitivity to gamma-rays was measured at a dose rate of 77.2 Gy/h and 192 Gy/h. The results when using different channel thresholds are reported in Table 2.4.4.

Again, scintillator A provides slightly higher results in comparison with scintillator B. The results show that the sensitivity can be reduced by increasing the threshold. At a threshold of 100, no sensitivity to gamma-rays was measured.

This experiment proves that the neutron scintillators do have a sensitivity to gamma radiation. In a spent fuel assembly the gamma dose rate can be up to 1000 Gy/h [24] which is much higher than 192 Gy/h, so further testing is required to study the behavior of the scintillators under gamma radiation.

The results from the experimental testing of the two scintillators and their sensitivities to thermal neutrons and gamma rays are discussed in Paper III.

Table 2.4.4: Gamma-ray sensitivity at different thresholds for the two scintillators.

Dose rate (Gy/h)	Scintillator	Sensitivity in c/s per 100 Gy/h		
		Threshold 10	Threshold 20	Threshold 100
77.2	A	1.8(1)	0.3(≤ 1)	—
192	A	1.3(1)	0.4(≤ 1)	—
77.2	B	0.4(1)	0.1(≤ 1)	—
192	B	0.8(1)	0.2(≤ 1)	—

Chapter 3

ANNs for the identification of diverted SNF

The general strategy for the verification of the integrity of SNF assemblies is to acquire measurements of observable quantities, such as the neutron flux or gamma emission rates, and determine whether the outcome of the measurements is consistent with the declared configuration of the assemblies or not. An underlying assumption is that there is a one-to-one correspondence between the spatial distribution of the observables and the actual pin pattern of the fuel assembly, whether intact or not. The identification of defects in the unknown system (the SNF assembly under inspection) via the observable quantities that originated from the system configuration itself is a so-called “inverse task”. In the current PhD research, machine learning models based on Artificial Neural Network (ANNs) are investigated for solving the inverse task and thus for the identification of partial defects from NDA measurements.

The suitability of ANNs in solving inverse tasks in nuclear engineering problems was pointed out some time ago [25]. In addition, recent research efforts have shown the potential of machine learning to enhance the processing of measured data in nuclear systems and extract more details of their configuration. For example, machine learning algorithms were used to quantify the percentage of replaced fuel pins in SNF assemblies [26, 27], to predict characteristic parameters of SNF such as Burn-Up (BU), Initial Enrichment (IE) and Cooling Time (CT) [28], to detect and localize missing radioactive sources within a small grid [29], to detect anomalies in the actinide inventories for a SNF reprocessing facility [30], and to identify and localize perturbations in nuclear reactor cores from neutron flux measurements [31, 32, 33].

In this chapter, ANNs are introduced in section 3.1. A first synthetic dataset is described in section 3.2. The dataset includes the detector responses of the Partial Defect Tester (PDET) and it was used to investigate two tasks. The first was to develop ANNs that can classify SNF assemblies into different categories based on the percentage of diverted fuel pins, see section 3.3. The second was to study the potential of ANNs in identifying the presence/absence of the individual fuel pins in SNF assemblies and hence reconstruct possible diversion patterns, see section 3.4. A second dataset was created using the simulated measurements of the scalar neutron flux and its gradient and it was used to investigate the effects of using the gradient of the neutron flux for the characterization of diversion patterns, see section 3.5.

3.1 Artificial Neural Network (ANNs)

Artificial Neural Networks (ANNs) are an advanced approach for machine learning and deep learning tasks. They can model non-linear relationships and thus learn to identify patterns in complex and large sets of data. Examples of their application can be found in systems for image, voice and text recognition [34]. ANNs are attractive for the investigation of partial defects in SNF assemblies because they may enable a more detailed evaluation of the system configuration which is needed, e.g., for a precise localization of the possible missing fuel pins.

Different architectures and strategies can be used to construct ANNs. In the current context, a feed-forward ANN [35] has been chosen. Typically, this consists of interconnected neurons, arranged in an input layer, one or more hidden layers, and an output layer, see Figure 3.1.1.




Figure 3.1.1: Schematic of a feed-forward artificial neural network [35].

To provide correct and usable predictions, a neural network needs to be trained, i.e., it requires a proper tuning with the problem under study. Then, sets of true cases of interest with their respective inputs and outputs are used to make the neural network learn the relationship between inputs and outputs.

The network receives a set of inputs (x_i , $i=1\dots n$) via the input layer. The input quantities are combined according to the connections between the neurons and moved progressively through the hidden layers until they reach the output layer where the final result is delivered. Each connection between the nodes in different layers has an associated weight that determines the strength of the connection and influences the output of the neurons. These weights are parameters that the neural network learns during the training process. The network is optimized for the given task through hyper-parameters such as, the number of neurons and the number of hidden layers, the activation functions, the loss function and the optimization algorithm (optimizer).

The neurons are usually selected so that their number in the input layer is equal to the number of input features and thus corresponds to the dimensionality of the input. The number of neurons in the output layer is associated with the size of the output and the nature of the task, i.e., classification, regression, etc. The number of hidden layers and the number of neurons in each of them needs to be tuned in order to optimize the performance of the algorithm.

Each neuron in the hidden layer(s) and in the output layer has an activation function which takes the weighted sum of inputs coming into the neuron and produces an output that is passed into the next layer, the produced output in the output layer being the final result. Activation functions are non-linear mathematical functions, by introducing non-linearity to the model they allow it to learn more complicated patterns. Different activation functions can be chosen, depending on the problem to process. Some popular activation functions [36] are: the Sigmoid, the Hyperbolic Tangent (Tan-h), the Rectified Linear Unit (ReLU) and the Softmax.

When the final output is produced, the network compares its predictions to the actual target values using a loss function, e.g, the Logarithmic Loss (Cross-Entropy) and the Categorical Cross-Entropy [37]. The goal during the training process is to minimize this loss. This is done by the network through a process called backpropagation which calculates the gradient of the loss with respect to the initial weights and updates them using an optimization algorithm. The optimizer iteratively adjusts the weights and biases to guide the network towards finding the optimal set of parameters that result in the most accurate predictions. Some popular optimizers [38] are: the Gradient Descent, the Stochastic Gradient Descent, the Nesterov Accelerated Gradient, the Adaptive Gradient (AdaGrad) and the Adaptive Moment Estimation (AdaM).

For a proper training of the ANN, the batch-size and the number of epochs are other key-parameters that needs to be carefully selected. The batch-size represents the number of training samples that the ANN needs to process before updating its internal parameters. The number of epochs determines the number of complete passes the network makes through the entire dataset. These parameters have no default/preferable values, therefore they have to be optimized with respect to the problem at-hand in order to avoid over- or under-fitting of data.

The accuracy of an ANN model can be estimated via a N-fold cross-validation process. Accordingly, the whole dataset is shuffled and divided into N random batches. N-1 of these batches are used for the training, while the remaining one is used as a testing dataset. The process is repeated N times so that each of the N batches can serve as testing dataset. The final estimate of the model accuracy is taken as the average of the accuracy on the testing datasets respectively calculated in the N repetitions. The result will then be less biased and the prediction capability of the model can be evaluated in a fairer manner when compared with other types of models. However, the cross-validation is computationally expensive since it requires the development of a separate model for each repetition.

3.2 Dataset with PDET signatures

A first investigation of the capabilities of ANNs for solving the inverse task of detecting partial defects in SNF assemblies relied on a synthetic dataset which was previously developed at SCK CEN [26, 27]. The dataset includes the responses from the detectors of the Partial Defect Tester (PDET) simulated via the Monte-Carlo N-Particle (MCNP) code [39] in intact 17x17 PWR spent fuel assemblies and a variety of hypothetical diversion scenarios. The simulated assemblies consist of 264 fuel pins and 25 empty guide tube positions.

The PDET is a non-destructive safeguards inspection tool which includes a set of neutron fission chambers and gamma-ray ionization chambers that can be inserted in the guide tubes of SNF assemblies of PWRs and measure the passive emission of neutrons and gamma-rays from the

spent fuel [40, 41]. The first prototype of the instrument was developed by Lawrence Livermore National Laboratory (LLNL) [42] and then further developed by SCK CEN [27].

The dataset contains 196 intact fuel assemblies. Each of them is unique in terms of its Initial Enrichment (IE), Cooling Time (CT) and Burn-Up (BU). There are 107 modelled diversion patterns, both symmetrical and asymmetrical, and have a minimum of 4 up to a maximum of 180 fuel pins replaced by stainless steel dummy pins, see Figure 3.2.1. Each of the 107 diversion scenarios is repeated 9 times, but with different conditions of IE, CT and BU, so the overall number of fuel assemblies with missing fuel pins is 963. The values of IE, CT and BU, which are used for the intact and diverted cases in all the possible combinations, are summarized in Table 3.2.1.

Theoretically, the one-to-one correspondence between the neutron or gamma-ray flux and the defect configuration exists only for fuel assemblies of the same IE, BU and CT values. Accordingly, it would be a more straightforward procedure if the training dataset contained intact cases and diversion scenarios for assemblies with the same IE, BU and CT. However, in reality SNF assemblies subjected to safeguards inspections can have many different combinations of those three parameters and thus creating a separate dataset based on each set of possible values would be a laborious task. Using the PDET dataset, which includes different values of IE, CT and BU for the training can therefore showcase the ability of ANN models to process mixed data. The study provides some indications of the significance of the IE, CT and BU, as it is discussed later in section 3.3.




Figure 3.2.1: Examples of diversion scenarios included in the dataset.

Table 3.2.1: Values of BU, IE and CT included in the PDET dataset.

	Intact fuel assemblies	Diversion scenarios
Burn-Up (MWd/kgU)	5, 10, 15, 20, 30, 40, 60	10, 30, 60
Initial Enrichment (w%)	2, 2.5, 3, 3.5, 4, 4.5, 5	2, 3.5, 5
Cooling Time (years)	1, 5, 10, 50	5

3.3 Percentage of diverted fuel using PDET signatures

The first task approached was to develop an ANN that can categorize SNF assemblies with respect to a set of classes based on the percentage of diverted fuel pins. Seven classes are prescribed and are reported in Table 3.3.1. The class label 0 is for intact fuel assemblies and the class labels 1 to 6 indicate fuel assemblies with progressively higher numbers of replaced fuel pins. The simulated responses from the PDET are the input features to the network, and the class label based on the percentage of replaced fuel pins represents the response (output). This type of problem corresponds to a supervised, multi-class classification problem.

Table 3.3.1: Percentage of diverted fuel pins and prescribed class labels.

Percentage of replaced pins (x)	Class label
$x = 0$	0
$0 < x \leq 10\%$	1
$10\% < x \leq 20\%$	2
$20\% < x \leq 30\%$	3
$30\% < x \leq 40\%$	4
$40\% < x \leq 50\%$	5
$x > 50\%$	6

Accordingly, a feed-forward network with 1 hidden layer was built using the TensorFlow [43] and the Keras [44] open-source software libraries. The neurons that belong to the input and hidden layers are activated with the ReLU function due to its simplicity, effectiveness and efficient convergence rates. The neurons in the output layer are activated with the SoftMax function which can provide, in a multi-class problem, the probability of an input to belong to each of the specified output classes.

The weights and the learning rate of the network are optimized using the AdaM optimizer and the evaluation of the error (loss) of the algorithm during the optimization process is obtained from the Categorical Cross-Entropy loss function, which can handle multiple classes.

3.3.1 Comparison of machine learning algorithms

First, an ANN model was trained and tested using only the neutron flux responses available from the PDET dataset and the capabilities of the ANN model were compared with two different non-parametric supervised learning methods, namely the Decision Trees (DTs) [45] and the k-Nearest

Neighbors (kNNs) [46]. The DT and kNN models used for this comparison were developed at SCK CEN [26].

The training and testing of the ANN model was performed via a 5-fold cross-validation process. The number of neurons in the input layer was set to 25 corresponding to the number of input features (the neutron flux responses from the 25 guide tube positions). The number of neurons in the output layer was set to 7 corresponding to the number of output class labels. Parameters such as the number of epochs, batch size and number of neurons in the hidden layer were manually tuned-in to achieve best performance. Similarly, the DT and kNN models were trained and tested via a 5-fold cross-validation process using the same set of data as the ANN model. Parameters such as the tree depth for the DT, and the number of k-neighbors for the kNN were also manually optimized. The tuning process of the three models is described in more details in paper IV.

Results show that for an optimal choice of the parameters, the ANN (with 96 neurons in the hidden layer) and the DT (with tree depth equal to 10) have a similar classification accuracy (90.8% for the ANN and 86.5% for the DT), while the kNN (with a number of k-neighbors of 5) has a worse performance, the classification accuracy being equal to 69.2%. More insights can be obtained from the confusion matrix of each algorithm, see Figure 3.3.1. The confusion matrix provides a summary of the number of correct and incorrect predictions of one algorithm, for each class included in the database. The results shown in the confusion matrices are based on the testing of the last batch (out of the 5 batches used for the cross-validation process). The ANNs showed the best classification results among the three algorithms with the DTs providing a comparable performance.




Figure 3.3.1: Confusion matrices based on the comparison between the ANN, DT and kNN models.

3.3.2 Optimized ANN model

Given the good performance in comparison with other machine learning algorithms, a second, further optimized model was developed for the task of categorizing SNF assemblies based on the percentage of diverted fuel pins. Using the same network settings, the new model was trained and tested using both the neutron and the gamma signatures from the PDET dataset, therefore showcasing the ability of the model to process combinations of different detector responses at the same time. Min-max normalization was separately performed on the set of neutron emissions and the set of gamma-ray emissions so that they are on the same scale and their processing is more consistent.

The ANN is trained and tested according to a 5-fold cross-validation process. The number of neurons in the input layer was set to 50, corresponding to the neutron flux and gamma emission rates obtained from each of the 25 guide tube positions. The number of neurons in the output layer was set to 7. A grid search optimization was performed to determine a number of epochs equal to 2000, a batch size equal to 10, and a number of neurons for the hidden layer equal to 50.

The ANN now reaches a classification accuracy of 96.5%. More insights are provided by the confusion matrix that summarizes the correct and incorrect predictions for each class, see Figure 3.3.2. The predictions presented in the confusion matrix are the aggregated testing results of the 5 batches from the the cross-validation process.

		Predicted Label						
		0	1	2	3	4	5	6
True Label	0	188	8	0	0	0	0	0
	1	18	153	0	0	0	0	0
	2	0	4	208	2	2	0	0
	3	0	2	0	184	2	1	0
	4	0	0	0	1	143	0	0
	5	0	0	0	0	0	144	0
	6	0	0	0	0	0	0	99

Figure 3.3.2: Confusion matrix of the optimized ANN model.

The confusion matrix shows that the misclassified fuel assemblies fall into one or two class higher or lower than their true class. The ANN estimates the probability that a fuel assembly belongs to any of the prescribed classes and assigns the fuel assembly to the class with the highest probability. In the current misclassifications, the correct class has always the second highest probability, and the relative differences between the probabilities of the true and predicted labels are below 5% in most cases, see Figure 3.3.3.

The majority of the misclassifications are in classes 0 and 1. This bias might be expected since the number of intact fuel assemblies in the dataset is bigger than the number of fuel assemblies




Figure 3.3.3: Relative differences between probabilities associated with true and wrongly predicted class labels.

with a specific diversion pattern and since the class with true label 1 (maximum partial defect of 10%) is the closest to the class of intact fuel assemblies with true label 0. On one hand, the cases with partial defects belonging to class 1 but identified as intact cases (false negatives), are of severe concern because diverted material goes undetected. Examples of these misclassified fuel assemblies are shown in Figure 3.3.4 and they share common characteristics, i.e., the removal is symmetric, in a checkered-like pattern, and more focused on the outer edges (in positions that are further from the guide tubes where the observables are measured). On the other hand, intact cases predicted in class 1 (false positives/false alarms) are a less sensitive issue, even though they would require additional resources and time for clarification during an inspection.

The analysis also indicates that higher numbers of misclassified diversions are found at BU of 10 MWd/kgU and IE of 2 w%, BU of 10 MWd/kgU and IE of 5 w%, and BU of 30 MWd/kgU and IE of 2 w%, see Figure 3.3.5. The tendency could be related to the nature of the dataset used to train the model. As described in section 3.2, the dataset covers more variations in BU, IE and CT for the intact cases than for the diversion scenarios. In addition, the intact fuel assemblies with low BU (between 5 and 30 MWd/kgU) are larger in number than the ones with high BU (between 30 and 60 MWd/kgU), while their values of IE are evenly distributed between 2 and 5 w%. Further investigations are needed to understand the effects of these parameters on the performance of the ANN and to identify possible sources of bias.




Figure 3.3.4: Diversion patterns predicted as intact by the ANN model (false negatives).

		Initial enrichment (w%)		
		2	3.5	5
Burn-up (MWd/kgU)	10	7	3	6
	30	6	3	2
	60	3	1	1

Figure 3.3.5: Number of misclassifications with respect to BU and IE.

3.4 Identification of diversion patterns using PDET signatures

To obtain more detailed characterizations of partial defects in SNF assemblies, a second, more complex task was tackled. Here, ANNs were applied to determine the exact configuration of fuel pins within the assembly under inspection. Each individual fuel pin is considered and its probability of being replaced is predicted by processing the neutron and gamma signatures from the PDET data. If the probability of a fuel pin to be identified is between 0.5 and 1, the fuel pin is labelled as missing, while, if the probability is less than 0.5, the fuel pin is labelled as present. This type of problem can be considered a multi-label binary classification. Similarly to the model described in subsection 3.3.2, the set of neutron and gamma emissions are separately normalized using min-max scaling.

Accordingly, a feed-forward network with 1 hidden layer was built using the same open-source software libraries mentioned in section 3.3, i.e., TensorFlow and Keras. The neurons that belong to the input and hidden layers are activated with the ReLU function and the neurons in the output layer are activated with the Sigmoid function. The latter is a typical choice for outputs that are non-mutually exclusive such as in the current application, where each pin is treated independently and can be either present or replaced. The weights and the learning rate of the network are optimized with the AdaM optimizer and the error (loss) of the algorithm in the optimization process can be adequately evaluated with the Binary Cross-Entropy loss function.

The number of neurons in the input layer was set to 50 corresponding to the neutron flux and gamma emission rates from the 25 guide tube positions. The number of neurons in the output layer was set to 264 corresponding to the number of outputs (one for each fuel pin in the assembly). A grid search optimization was performed and resulted in a number of epochs equal to 1000, a batch size equal to 25, and a number of neurons for the hidden layer equal to 300.

Given a 5-fold cross validation process, the ANN model was able to reconstruct the exact arrangement of the fuel pins in 667 out of the 1159 fuel assemblies available in the PDET dataset, which corresponds to a fraction of 57.5%. The predictions are correct for 97.4% of the intact fuel assemblies (191 out of 196) and 49.4% of the diversion scenarios (476 out of 963). As expected, the performance is better with the cases without partial defects because in the training the intact configurations are more numerous than any of the specific diverted configurations. Yet, the algorithm can detect 94.8% of all the incomplete fuel assemblies as diversion scenarios despite the incorrect number/location of the replaced fuel pins.

As mentioned earlier, the model estimates the probability of a fuel pin to be replaced within the fuel assembly. If the probability is higher than the threshold value of 0.5, the fuel pin is labelled as missing, otherwise as present. The analysis of the distribution of the probabilities for the fuel pins that are correctly and wrongly predicted can provide insights into the behavior of the model.

The distribution of the probabilities for all the correctly predicted fuel pins is shown in Figure 3.4.1. Two large peaks are found close to the probability values of 0 and 1, and reflect the high confidence in the correct results of the model.

The probability distribution of the wrongly predicted fuel pins is shown in Figure 3.4.2. A bigger portion of the misclassifications (64.5%) has probability less than 0.5 and thus consists of replaced fuel pins predicted as intact, i.e., false negatives. The distribution has two peaks near the probabilities of 0.5 and 0. The misclassifications with probabilities around 0.5 are characterized by a low level of confidence because the difference in probability with the other label (which is the true label) is small. The misclassifications with probabilities close to 0 (thus, far from the threshold) have a higher level of confidence. The tendency to make more misclassifications in favor of false negatives and with higher levels of confidence may depend on the training dataset, where the fraction of the intact fuel pins is large (78.4%). The construction of a more balanced dataset to avoid this type of bias is not straightforward. Adding more diversion scenarios does not necessarily increase the weight of the replaced pins since fuel assemblies with realistic partial defects still have a significant number of intact fuel pins.




Figure 3.4.1: Probability distribution of the correctly classified fuel pins.




Figure 3.4.2: Probability distribution of the misclassified fuel pins.

The misclassified fuel assemblies are 492 in total, and the associated errors may involve one or more fuel pins, see Table 3.4.1. The majority of these fuel assemblies have a relatively low number of incorrect fuel pins (between 1 and 20) and are therefore reconstructed correctly to a significant extent. Figure 3.4.3 shows examples of predicted diversion patterns with different numbers of misclassified fuel pins. Although the predictions are not entirely accurate, they can provide a useful indication of the main region of the real diversion, with the exception of the last case with 67 misclassified fuel pins.

Table 3.4.1: Number of fuel assemblies with x incorrect fuel pins.

Number of misclassified pins (x)	Number of cases	Fraction of the dataset (%)
$x = 0$	667	57.5
$1 \leq x \leq 10$	238	20.5
$10 < x \leq 20$	160	13.8
$20 < x \leq 30$	46	4.0
$30 < x \leq 40$	19	1.6
$40 < x \leq 50$	8	0.7
$50 < x \leq 60$	6	0.5
$x > 60$	15	1.3




Figure 3.4.3: Examples of misclassified diversion scenarios with low number of incorrect fuel pins.

3.5 Identification of diversion patterns using NGD signatures

Based on the positive performance of ANNs for the task of reconstructing the configuration of the fuel pins in SNF assemblies using the detector responses from the PDET, a novel aspect was explored, i.e., the use of the neutron flux gradient as an input feature for training the ANN model and solving the inverse task. This physical quantity is not considered in regular SNF verification, but it has richer information than the scalar neutron flux and thus is expected to enhance the identification of anomalies in SNF assemblies. Its measurement can be obtained from a detector such as the NGD described in chapter 2.

3.5.1 Dataset with the neutron flux gradient

A new synthetic dataset was created in this PhD project and it contains the simulated measurements of the scalar neutron flux and its gradient in the empty guide tubes of a 17x17 PWR SNF assemblies. The simulations were performed using the Serpent code in two steps, i.e, a burn-up simulation followed by a fixed-source simulation, see section 2.3.

The dataset contains one case of an intact fuel assembly (without defects) and 107 cases with diversion patterns, which are symmetrical or asymmetrical and have a minimum of 4 up to a maximum of 180 fuel pins replaced by stainless steel pins. The choice of the diversion patterns is consistent with the ones used in the PDET dataset (see Figure 3.2.1 for examples).

As discussed in section 3.2, SNF assemblies subjected to safeguards inspections can have different values of IE, BU and CT and therefore it would be more practical to have one dataset with different combinations of these parameters (similar to the case of the PDET dataset). However, in order to avoid any potential biases related to these parameters (see section 3.3) and to focus on the effects of the gradient, the fuel assemblies in this new dataset were simulated with only one set of values, i.e, an IE of 3.5 w%, a BU of 40 MWd/kgU, and a CT of 5 years.

Since only one version of each diversion pattern is considered, the ANN model is always tested over scenarios that are not seen in the training phase. Therefore, the ability of the network to make predictions with respect to unknown data can be better assessed. This aspect is relevant because a training dataset with all the possible diversion patterns would require unfeasible computational resources.

Considering the 25 empty guide tube positions in the simulated assemblies, the calculated system responses for each configuration in the dataset are 75, i.e., 25 values of thermal neutron flux, 25 values of the magnitude of the gradient (absolute value), and 25 values of the angle of the gradient vector (direction). An alternative version of the dataset was also created, in which the two Cartesian components (x and y) of the gradient vector are used instead of the magnitude and the direction. In both versions of the dataset, min-max normalization was performed on each set of input features separately so that they are on the same scale.

3.5.2 The gradient as input feature

An ANN capable of processing the neutron flux and the gradient information from the new dataset was developed with the same settings and functionality as the one described in section 3.4, i.e., predicting the presence or absence of individual fuel pins in SNF assemblies.

To highlight the effects of the gradient of the neutron flux on the identification of replaced fuel pins, a set of models were trained using different combinations of the available input features, that is:

- The magnitude and direction of the thermal neutron flux gradient ($G_m + G_d$)
- The magnitude of the gradient (G_m)
- The direction of the gradient (G_d)
- The thermal neutron flux (N)
- The neutron flux and the magnitude and direction of its gradient ($N + G_m + G_d$)
- The neutron flux and the magnitude of its gradient ($N + G_m$)
- The neutron flux and the direction of its gradient ($N + G_d$)

The number of neurons in the input layer was equal to the number of input features used in each model. The number of neurons in the output layer was set to 264 corresponding to the number of outputs. A grid search optimization was performed to determine the number of neurons in the hidden layer, the number of epochs and the batch size for each model independently.

The models were trained using a 6-fold cross-validation process and the results were scored based on the number of fuel pins that have been identified correctly in all the fuel assemblies available from the dataset. The predictions by each model are then characterized into 4 categories: The 'True Negatives' are all the correctly predicted intact fuel pins, the 'True Positives' are all the correctly predicted missing fuel pins, the 'False Positives' are the intact fuel pins that are wrongly predicted as missing and the 'False Negatives' are the missing fuel pins wrongly predicted as intact, see Figure 3.5.1.

ANN Model		Real Label	
		Intact (0)	Missing (1)
Predicted Label	Intact (0)	True Neg.	False Neg.
	Missing (1)	False Pos.	True Pos.

Figure 3.5.1: General form of the confusion matrix used to characterize the predictions of the ANN.

The performance of the ANN models is quantified with 4 metrics, i.e., the pin-accuracy, the precision, the recall and the F1 score. The pin-accuracy corresponds to the percentage of the correctly predicted fuel pins (the sum of the true positives and true negatives) out of the total number of fuel pins, considering all the fuel assemblies in the dataset. The precision is defined as the fraction of correctly predicted missing pins (the true positives) over all the pins predicted as missing (the sum of true and false positives). The recall is equal to the fraction of correctly

predicted missing pins (true positives) over the total number of missing pins in the dataset (equivalent to the sum of true positives and false negatives). The F1 score is the harmonic mean of the precision and recall values.

Table 3.5.1 shows the comparison between the different ANN models. The model that uses the neutron flux and the magnitude of its gradient ($N+G_m$) has the best performance in all four metrics. The model that uses only the direction of the gradient vector (G_d) has the lowest performance in terms of pin-accuracy and precision. The model that relies only on the thermal neutrons (N) has the lowest performance in terms of recall.

The precision reflects the general ability of the model to avoid false predictions of both the intact and replaced fuel pins. The use of both the thermal neutron flux and the magnitude of its gradient (either separate or combined) result in better precision values. The direction of the gradient vector as an input feature always has a negative effect on the precision of the ANN.

The recall value depends on the number of false negatives and thus is an indication of the ability of the model to correctly predict replaced fuel pins in the assembly. The models that use the gradient (either in magnitude, direction or both) have greater recall values and hence can better detect replaced fuel pins, while the model based only on the thermal neutron flux has the lowest recall value.

Table 3.5.1: Performance metrics with respect to the different combinations of input features.

Metric	Gradient Detector Responses						
	$N + G_m$	$N + G_m + G_d$	$G_m + G_d$	G_m	N	$N + G_d$	G_d
Pin-accuracy	0.82	0.81	0.80	0.80	0.79	0.77	0.76
Precision	0.66	0.63	0.62	0.64	0.64	0.56	0.55
Recall	0.60	0.59	0.59	0.52	0.43	0.44	0.44
F1	0.63	0.61	0.60	0.57	0.51	0.49	0.49

None of the models can fully reproduce any of the diversions. This is expected because the size of the dataset is relatively small. As an example, Figures 3.5.2 and 3.5.3 show two configurations with partial defects and their reconstruction via the models that use the neutron flux and the magnitude of its gradient ($N+G_m$), the neutron flux along with the magnitude and direction of its gradient ($N+G_m+G_d$), and the neutron flux (N), respectively.

Figure 3.5.2 shows an example where the results of the models reflect the global trends reported in Table 3.5.1. The $N+G_m$ model provides the closest prediction to the real pattern as indicated by the values of the evaluation metrics for the specific case. The N model and the $N+G_m+G_d$ model both have values of precision, recall and F1 score equal to zero because they do not identify any of the missing pins correctly (i.e., no true positives). In addition, the N model gives less false positives in comparison to the $N+G_m+G_d$ model, which is consistent with the general finding that the N model tends to have slightly higher precision value than the $N+G_m+G_d$ model.

A case that does not follow the global trend is also included, see Figure 3.5.3. The N model provides the best reconstruction of the diversion in terms of all the evaluation metrics. Such result might be related to different factors, e.g., the specific characteristics of the diversion pattern combined with the knowledge learned by the algorithm from similar scenarios in the training process.




Figure 3.5.2: Example of a diversion scenario for which the results of the models are consistent with the global trends of Table 3.4.




Figure 3.5.3: Example of a diversion scenario for which the results of the models deviate from the global trends of Table 3.4.

3.5.3 Refinement of the analysis

Considering the findings from the previous subsection, the training process of the ANN was further optimized. A 10-fold cross-validation process was adopted. A larger number of folds has the advantage of increasing the number of observations in the training phase, which can lead to a better performance. In order to make a 10-fold cross-validation process viable, two new diversion patterns were added to the dataset so that the total number of cases is 110 and 10 equal batches can be created. In addition, the entire cross-validation process is repeated 5 times with a different initial shuffling of the dataset in order to reduce any potential bias that can result from a specific distribution of the diversion patterns between the training and testing phases.

Accordingly, two of the models described in subsection 3.5.2 were trained again, i.e, the model that relies only the thermal neutron flux (model N) and the model that relies on the the neutron flux and the magnitude and direction of the gradient (model $N+G_m+G_d$). The choice for retraining these two models in particular was to gain more insights about using the information from the gradient as additional input features for training the ANN compared to relying on the scalar neutron flux solely. The model that processes only the gradient of the neutron flux (G_m+G_d) was not considered since it has a similar performance to model $N+G_m+G_d$. Since the scalar neutron flux is needed to derive the gradient it is available to be used as input to the algorithm at no additional cost.

One observation from the previous subsection was the negative effect on the precision from the direction of the gradient. In order to try to overcome this, a new model was considered (model $N+G_x+G_y$). The model was trained using the alternative version of the gradient dataset which relies on the neutron flux and the two Cartesian components of the gradient instead of the magnitude and direction.

After the training process was completed, the three models were scored based on the number of fuel pins that have been identified correctly in all the fuel assemblies available from the dataset. The results of each ANN model are summarized in the confusion matrices shown in Figure 3.5.4 and the performances are quantified using the same 4 metrics used in the previous analysis, i.e., the pin-accuracy, the precision, the recall and the F1 score, see Table 3.5.2. The results shown in the confusion matrices are taken from the cross-validation process with the best performance out of the five repetitions while the values of the performance metrics reported in 3.5.2 are averaged over the five repetitions. The relatively low standard deviations on the performance metrics reflect that the initial random shuffling of the dataset prior to the cross-validation has a minor effect on the performance of the network.

Table 3.5.2: Optimized performance metrics for the three ANN models: N, $N+G_m+G_d$, and $N+G_x+G_y$.

Metric	N		$N+G_m+G_d$		$N+G_x+G_y$	
	Mean value	Std. ($\pm$)	Mean value	Std. ($\pm$)	Mean value	Std. ($\pm$)
Pin accuracy	0.80	0.01	0.80	0.01	0.85	0.01
Precision	0.66	0.01	0.63	0.02	0.72	0.02
Recall	0.43	0.01	0.55	0.02	0.70	0.01
F1	0.52	0.01	0.58	0.02	0.70	0.01

N		Real Label	
		Intact (0)	Missing (1)
Predicted Label	Intact (0)	20087	4206
	Missing (1)	1572	3175

(a)

N + G _m + G _d		Real Label	
		Intact (0)	Missing (1)
Predicted Label	Intact (0)	19268	3322
	Missing (1)	2391	4059

(b)

N + G _x + G _y		Real Label	
		Intact (0)	Missing (1)
Predicted Label	Intact (0)	19641	2217
	Missing (1)	2018	5164

(c)

Figure 3.5.4: Confusion matrix of the model N (a), confusion matrix of the model N+G_m+G_d (b), and the confusion matrix of the model N+G_x+G_y (c).

The results show that model N (only the thermal neutron flux) and model N+G_m+G_d (thermal neutron flux together with the magnitude and direction of its gradient) have similar values in terms of pin-accuracy. The first model is better to predict the intact pins, see Figure 3.3.1a. The second model is better with missing pins, see Figure 3.3.1b. Model N+G_x+G_y (thermal neutron flux together with the x and y components of the gradient) has a higher pin accuracy since it predict correctly a larger number of missing fuel pins (true positives), see Figure 3.3.1c.

Model N has a slightly higher precision value than model N+G_m+G_d. Model N predicts less missing fuel pins correctly (less true positives), but it gives less errors in terms of intact pins (less false positives). The second model provides a higher number of correct missing pins (more true positives), but it over-predicts pins as missing (more false positives). According to the precision metric, model N+G_x+G_y performs better since it identifies a higher fraction of true positives (correct missing fuel pins) over false positives (misclassified intact fuel pins).

In terms of recall, the models that use the gradient (either in magnitude and direction or the two components) have larger recall values and hence can better detect replaced fuel pins in comparison to the model based only on the thermal neutron flux. An under-estimation of replaced fuel pins in SNF assemblies (higher number of false negatives) such as the case of the model that relies only on the thermal neutron flux is undesirable from a safeguards perspective since it can lead to diverted nuclear material being undetected.

The F1 score (the harmonic mean of the precision and recall values) confirms that the use of the neutron flux gradient is advantageous and that model N+G_x+G_y performs better than the other considered models. Despite the magnitude and direction of the gradient having a more immediate physical interpretation, the x and y components are proven to be more beneficial to the ANN for the reconstruction of the diversion patterns. The two Cartesian components are a more primary representation of the gradient and they contain independent information from

each other, meanwhile the magnitude and direction are both derived from the two components and thus may be more correlated with each other.

Again, the three ANN models cannot fully reconstruct any of the diversion patterns. However, the majority of the predictions, especially from model $N+G_x+G_y$, are close to the real diversion patterns.

An example referred to as A is shown in Figure 3.5.5. The diversion pattern has three fuel pins replaced by dummy pins at each corner of the fuel assembly. Such a scenario is challenging to detect because the amount of replaced nuclear material is small, the affected locations are relatively far from the guide tubes (where detectors could be placed), and the pattern is symmetrical. The two models trained with the gradient of the neutron flux can detect the lack of fuel pins at the correct positions. The diversion pattern is better retrieved from the thermal neutron flux and the two Cartesian components of the gradient than from the thermal neutron flux and the magnitude and direction of the gradient. The use of only the thermal neutron flux fails to detect any of the replaced fuel pins and the assembly is predicted as intact.




Figure 3.5.5: Example A of a diversion case from the dataset and how it was reconstructed by the three ANN models.

A comparison between the simulated signatures of an intact fuel assembly and example A is shown in Figure 3.5.6. In the two cases, the thermal neutron flux in the guide tubes is affected only negligibly by the diversion, while the gradient has significant deviations, which are stronger in the guide tubes closest to the locations of the missing pins. Therefore, the use of the gradient leads to an improved performance of the machine learning algorithm.

On the other hand, there exist a few cases in which the predictions did not resemble the real diversion pattern at all. Example B shown in Figure 3.5.7, is representative of these cases. In the assembly, two rows of fuel pins replaced by dummy pins next to the upper edge cause a significant disruption in the distribution of the thermal neutron flux and its gradient within the system, which can be easily seen in the simulated measurements inside the guide tubes, see Figure 3.5.8. However, model $N+G_x+G_y$ (which processes the thermal neutron flux together with the two components of the gradient and provides the best predictions) fails to reconstruct the pattern correctly. The reason for these anomalies is that example B (as well as the other few cases with the same issue) have no common features with the other configurations available in the dataset, therefore the network is not sufficiently trained to characterize them well.




Figure 3.5.6: Simulated responses of the thermal neutron flux and the gradient of the neutron flux from an intact fuel assembly and the diversion pattern from example A.




Figure 3.5.7: Example B of a diversion case (left) and its reconstruction via model $N + G_x + G_y$ (right).




Figure 3.5.8: Simulated responses of the thermal neutron flux and the gradient of the neutron flux from example B.

To investigate this issue, the dataset is expanded by adding two new diversion scenarios that resemble example B, see Figure 3.5.9. Then, model $N+G_x+G_y$ is trained with the updated dataset and tested on example B again, see Figure 3.5.10. The predictions are significantly improved since the model has learned from two similar cases.




Figure 3.5.9: Additional diversion patterns for the updated dataset.




Figure 3.5.10: Example B (left) and its reconstruction via updated model $N+G_x+G_y$ (right).

Chapter 4

Conclusions

A summary of the research and recommendations for future work are given in sections 4.1 and 4.2, respectively. Ethical reflections on the research carried out in this PhD project are presented in section 4.3.

4.1 Summary of the research

The work presented in this thesis investigates a non-intrusive methodology that can enhance safeguards inspections and the detection of partial defects in Spent Nuclear Fuel (SNF) assemblies from Pressurized Water Reactors (PWRs). The methodology relies on the algorithmic processing of SNF characteristic signatures via Artificial Neural Networks (ANNs) to detect whether or not nuclear materials have been diverted from the assembly under inspection and provide information regarding the extent and exact location of the diversion (if any). A novel aspect of the study is related to the possibility of using the gradient of the thermal neutron flux, a physical quantity that has not been used in safeguards inspections before, as a characteristic signature that could improve the detection and localisation of diverted fuel pins.

To enable the measurements of the thermal neutron flux and its gradient within a SNF assembly, the design of a miniaturized detector, referred to as the Neutron Gradient Detector (NGD), was proposed in this thesis and its performance was evaluated using the Serpent Monte-Carlo code. The detector is based on four thin LiF/ZnS(Ag) optical fiber-mounted neutron scintillators arranged in an aluminium matrix according to a rectangular pattern. The detector can be used to estimate the two components of the gradient of the scalar neutron flux, from the difference between the measurements provided by the diagonally-opposite pairs of scintillators, and thus provide an estimation of the magnitude and direction of the gradient vector in the measurement position. (Section 2.1 and Paper I)

The detector was first modelled and its performance simulated in a hypothetical setup with a ^{252}Cf neutron source in a water tank. A quantitative analysis showed that determination of the gradient vector was feasible with the proposed design and that the presence of the detector and associated shielding effects do not introduce a significant distortion of the flux. It was also demonstrated that a calibration process can be performed in order to provide a correct estimation of the gradient despite having scintillators with different sensitivities to thermal neutrons. In

addition, results showed that the presence of minor uncertainties in the placement of the detector and its orientation during the calibration process have minor influence on the calculated gradient. (Section 2.2 and Paper I)

Monte-Carlo calculations were performed with the gradient detector being placed in the 25 empty guide tube positions of a 17x17 PWR SNF assembly model. The gradient was calculated for an intact version of the assembly with all the fuel pins present and for diverted cases with a number of the fuel pins being replaced by stainless steel dummy pins. The comparisons between the intact case (used as reference) and the diversion scenarios showed that relevant information can be retrieved from the gradient which can help to localize the region of diversion. (Section 2.3 and Paper II)

Two LiF-ZnS(Ag) optical fiber-based neutron scintillators, on which the concept of the NGD is based, were experimentally characterized. Their sensitivity to thermal neutrons and their relative efficiencies were estimated with a neutron source in the hot-cell laboratory at Chalmers University of Technology and in the BR1 research reactor at the Belgian Nuclear Research Centre (SCK CEN). The sensitivity of these type of scintillators to gamma radiation was also investigated at the LNK facility at SCK CEN using ^{137}Cs and ^{60}Co sources with different dose rates (up to 192 Gy/h). Results showed that the scintillators can be sensitive to gamma rays starting from a dose rate of ~ 70 Gy/h. The contribution from gamma rays can be however discarded by setting a proper discrimination threshold. (Section 2.4 and Paper III)

For solving the inverse-task of identifying partial defects in SNF assemblies from measurements of characteristic signatures, the application of machine learning models based on Artificial Neural Networks (ANNs) were studied. Two tasks were approached with the purpose of identifying, at different levels of detail, whether fuel pins has been removed or replaced by dummy pins or not. As a first step, the ANN models were trained and tested using a dataset of simulated responses of the Partial Defect Tester (PDET) for both intact fuel assemblies and diversion scenarios with different values of Burn-Up (BU), Initial Enrichment (IE) and Cooling Time (CT). (Section 3.2)

For the first task, SNF assemblies are categorized into different classes defined by the percentage of replaced fuel pins. For this purpose, the ANN processes the PDET information for one fuel assembly, estimates the probability of the fuel assembly to be in each of the prescribed classes, and assigns the fuel assembly to the class with the highest probability. The ANN was initially trained using only the neutron flux responses from the PDET dataset and its performance was compared with other machine learning algorithms, i.e., Decision Tress (DTs) and k-Nearest Neighbors (kNNs). The ANN model showed the best classification accuracy (90.8%) compared to the other models. (Subsection 3.3.1 and Paper IV)

The ANN was then further optimized to process both the neutron flux and gamma emission rates from the PDET and it was able to reach a classification accuracy of 96.6%. For the misclassifications, the relative differences between the probabilities of the predicted and the true class is below 5% with few exceptions. In addition, the majority of the misclassifications occur between the class of intact fuel assemblies and the class with the smallest partial defects (i.e., less than 10% of replaced fuel pins). (Subsection 3.3.2 and Paper V)

The second task was to identify the presence or absence of each individual fuel pin inside an assembly and thus to retrieve the exact arrangement of the fuel pins from the analysis of the PDET responses. An ANN model was developed in which it estimates the probability of the

fuel pins to be replaced and labels the fuel pins with probability higher than the threshold value of 0.5 as missing, otherwise as present. The model was able to fully reconstructs 57.5% of the fuel assemblies and with a high level of confidence. A substantial majority of the misclassified fuel pins consists of false negatives, i.e., replaced fuel is diagnosed as present. The estimated probability of these fuel pins is below 0.5 and the distribution has a large peak close to zero. The aspect of many diverted fuel pins misclassified with probabilities close to zero is important because a high level of confidence is associated with the error and thus it becomes difficult to assess how trustworthy the results are. This bias can be related to the existing imbalance of labels in the dataset used for the training process (the intact fuel pins are 78.4% of the total). In addition, a large fraction of the misclassified fuel pins has probabilities around the threshold value of 0.5, so the wrong labelling of the fuel pins is assigned with a lower confidence. (Section 3.4 and Paper V)

In order to investigate whether or not the gradient of the neutron flux can be advantageous for identifying and localising diverted fuel pins in SNF assemblies, a new dataset was generated using Monte-Carlo simulations. The dataset includes values of the neutron flux and its gradient in the guide tubes of an intact assembly and 107 assemblies with different patterns of replaced fuel pins. (Subsection 3.5.1)

An ANN capable of processing the neutron flux and the gradient from the new dataset was developed with the same functionality as in the second task, i.e, identifying the presence or absence of each fuel pin inside an assembly. First, a set of models were trained using all possible combinations of the available input features, i.e., the neutron flux (N), the magnitude of the gradient (G_m), the direction of the gradient vector (G_d). The performance of each model was quantified in terms of the pin-accuracy, the precision, the recall and the F1 score. The model based on the neutron flux and the magnitude of its gradient ($N+G_m$) was found to have the best performance in all four metrics. Any model that used information related to the gradient (magnitude, direction, or both) leads to an improvement in the recall value which reflects a better identification of replaced fuel pins. However, results also showed that using the direction of the gradient vector as an input feature had a negative effect on the precision of the ANN. (Subsection 3.5.2 and Paper VI)

A second version of the gradient dataset was created in which the two Cartesian components of the gradient are used instead of the magnitude and direction. The training process was further optimized and a comparison was performed between three different models, i.e., a model for the analysis of only neutron flux measurements, a model which combines the neutron flux and the magnitude and direction of its gradient, and a model for the neutron flux together with the two Cartesian components of its gradient. The results again confirmed that the information about the neutron flux gradient is advantageous for the detection of patterns of replaced fuel pins within the assembly. In addition, the two Cartesian components of the gradient were more effective in training the ANN than the magnitude and direction of the gradient and led to better predictions. (Subsection 3.5.3 and Paper VII)

The models trained using the new dataset with the gradient measurements were not able to fully reconstruct any of the diversion patterns. This is expected due to the small size of the dataset. Nevertheless, models that were trained using either the magnitude of the gradient or the two Cartesian components in addition to the scalar neutron flux were able to provide results that are close to the real assembly configurations and thus can generalize to some extent the mapping

from the measured signatures to the patterns of replaced fuel pins. It was also observed that some diversion patterns cannot be reconstructed at all because they have no common features with the rest of the dataset. In order to overcome these issues, the dataset needs to be expanded and the diversion patterns included in it need to be more systematically selected, see section 4.2.

4.2 Future work

Future work may be recommended in three main areas as follows.

4.2.1 Neutron gradient detector

A prototype of the gradient detector is currently under construction at Chalmers University of Technology. The conceptual design of the detector described in section 2.1 and shown in Figure 2.1.1, is slightly modified. The main difference is in the scintillators. The optical fiber-based scintillators considered in the conceptual design and experimentally tested (section 2.4), were hand-made at KURNS and consist of a small volume of a mixture of a neutron converter (LiF) and scintillation material (ZnS(Ag)), glued on top of a thin light guiding fiber (about 1 mm in diameter). For the prototype, boron loaded plastic scintillators are chosen instead. Additional differences in the prototype is that each of the four scintillators will be coupled with four light guiding fibers instead of one and a silicon PM-array is used for the readout of light from the detector instead of individual PM-tubes. The use of boron loaded plastic scintillators as well as using a silicon PM-array was recommended to us by Prof. John Mattingly of North Carolina State University.

The reason for the change of the scintillators is that the LiF/ZnS(Ag) material is not transparent and the scintillation light can only escape from close to the surface, hence its efficiency cannot be improved by increasing its volume. Boron loaded plastic scintillators on the other hand are transparent, which allows their efficiency to be proportional to their size. Although the radial dimensions are limited by the measurement space accessible inside the nuclear fuel assembly (the guide tubes are about 1 cm in diameter), the volume can be increased via the axial length to a certain extent where the axial variation of the neutron flux still remains negligible. The boron loaded scintillators are also commercially available so that a regular supply of them is viable.

The reason for using four fibers per scintillator instead of one is to maximize the light output from the scintillators. The increased number of light guiding fibers (16 in total) makes it more practical to use a PM-array as a photo-multiplication device instead of a separate PM-tube for each fiber.

Once the construction of the prototype of the detector is finalized, its performance will have to be assessed both experimentally and via Monte-Carlo simulations. In the simulations, the generation and propagation of scintillation light might also be considered.

4.2.2 Artificial Neural Networks (ANNs)

ANNs are complex mathematical models that are usually referred to as "black boxes". The interpretation of the results and understanding the reasons of the specific output obtained from the network is not an easy task. In this PhD project, studies were performed to gain insights

into how the models work. However, future research can investigate techniques such as feature importance analysis, activation visualization, integrated gradients, and partial dependence plots in order to increase the explainability of the ANN models and to study the complex connections between inputs and outputs. For example, these techniques can help to clarify aspects such as the negative effect of the direction of the gradient vector on the precision of the network and the effect of having fuel assemblies with different values of BU, IE and CT in the dataset.

The explainability and transparency of the ANN models are crucial for building trust, ensuring accountability, and addressing ethical concerns associated with their deployment in nuclear safeguards applications.

4.2.3 Training dataset

The size and design of the dataset used to train and test any machine learning model are crucial. As observed in section 3.5, the ANN models using the neutron flux gradient were not able to fully reconstruct any of the diversion patterns because of the small size of the training dataset. Future work can be carried out to expand the dataset by including additional configurations of diverted fuel assemblies. A dataset that includes all possible scenarios, is however not feasible because of the limitation of computational resources. Then, sampling techniques that allow for better representations of the space of diversion scenarios may be beneficial.

The two synthetic datasets that were used to train and test the ANN models relied on Monte-Carlo simulations, see section 3.2 and subsection 3.5.1. Monte-Carlo simulations are computationally expensive. Future work could explore the use of lower-order simulation methods for the generation of training data.

4.3 Ethical considerations

While Artificial Intelligence (AI) offers numerous benefits and potential applications in the nuclear field, it is always essential to address the ethical aspects associated with such a technology. In the current context, the use of Artificial Neural Networks (ANNs) in the area of nuclear safeguards introduces several ethical considerations that need to be highlighted at an early stage such as data privacy, latent biases, explainability, proliferation risks, and the role of human oversight.

4.3.1 Data privacy

The collection and utilization of sensitive data, i.e., information related to nuclear facilities and materials, raise ethical concerns regarding privacy and security. Certain protocols, including data anonymization, encryption, and robust cybersecurity measures, might be essential to protect such confidential information. Ensuring data privacy can help to prevent any unauthorized access or misuse of the data, as well as to protect the individuals or organizations providing the data.

4.3.2 Latent bias

ANNs are susceptible to biases from the training data which can be difficult to identify. Such hidden biases may have unintended consequences, for example, as discussed in this thesis, ANNs

may inadvertently have false positive results (intact fuel pins being detected as missing) or false negative results (missing fuel pins being detected as intact), which could have undesired consequences for nuclear non-proliferation efforts. Rigorous testing for bias, diversification of training data, and ongoing monitoring are important to rectify biases and ensure fairness.

4.3.3 Explainability

The complexity of ANNs and their black-box nature poses challenges in understanding and interpreting their predictions, which might raise questions regarding their accountability. Incorporating explainability features, such as interpretable model architectures, visualizations, and transparent reporting mechanisms, is thus important to ensure accountability and trust.

4.3.4 Proliferation risks

ANN models for nuclear safeguards may be counter-trained for proliferation purposes. For example, the algorithms can be used to recognize diversion patterns that are not easily detectable by inspectors and thus facilitates the illicit removal of sensitive nuclear material. Robust cybersecurity measures to safeguard against unauthorized access, coupled with international collaboration and adherence to non-proliferation agreements, are essential to mitigate these risks.

4.3.5 Human oversight

Over-reliance on ANN models without human oversight raises concerns about excluding human judgment. An approach, where human experts play a central role in decision-making, validation, and oversight, is thus crucial for balancing the capabilities of ANNs with human expertise. As mentioned in the objectives (section 1.3), the ANN models developed in this thesis aim to facilitate the job of the safeguard inspectors and not to replace it.

Bibliography

- [1] J. Wagemans, G. Vittiglio, E. Malambu, and H. A. Abderrahim, “The BR1 reactor: a versatile irradiation facility for fundamental research and industrial applications,” in *The 1st International Conference on Advancements in Nuclear Instrumentation, Measurement Methods and their Applications, Marseilles, France, June, 2009*.
- [2] *Nuclear power reactors in the world*. Reference data series No.2, International Atomic Energy Agency (IAEA), 2023.
- [3] D. Fischer, “History of the international atomic energy agency,” 1997.
- [4] A. F. Woolf, *Arms Control and Nonproliferation: A Catalog of Treaties and Agreements*. DIANE Publishing, 2010.
- [5] H. Yoo, J.-h. Seo, N.-Y. Lee, J.-h. Lee, M.-s. Koh, and S.-H. Ahn, “Methodology for evaluating proliferation resistance of nuclear systems and its case study,” *Progress in Nuclear Energy*, vol. 100, pp. 309–315, 2017.
- [6] *Nuclear Material Accountancy and Control in Support of Nuclear Security A WINS International Best Practice Guide*. World Institute of Nuclear Security (WINS), 2016.
- [7] *Safeguards Glossary*. International Atomic Energy Agency (IAEA), 2022.
- [8] C. Poinssot, S. Bourg, N. Ouvrier, N. Combernoux, C. Rostaing, M. Vargas-Gonzalez, and J. Bruno, “Assessment of the environmental footprint of nuclear energy systems. comparison between closed and open fuel cycles,” *Energy*, vol. 69, pp. 199–211, 2014.
- [9] E. Branger, S. Grape, and P. Jansson, “Partial defect detection using the DCVD and a segmented region-of-interest,” *Journal of Instrumentation*, july 2020.
- [10] S. Vaccaro, I. Gauld, J. Hu, P. De Baere, J. Peterson, P. Schwalbach, A. Smejkal, A. Tomanin, A. Sjöland, S. Tobin, and D. Wiarda, “Advancing the fork detector for quantitative spent nuclear fuel verification,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 888, pp. 202–217, 2018.
- [11] M. Mayorov, T. White, A. Lebrun, J. Brutscher, J. Keubler, A. Birnbaum, V. Ivanov, T. Honkamäa, P. Peura, and J. Dahlberg, “Gamma emission tomography for the inspection of spent nuclear fuel,” in *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pp. 1–2, 2017.

- [12] *Safeguards Techniques and Equipment, International Nuclear Verification Series, No. 1 (Rev. 2)*. International Atomic Energy Agency (IAEA), 2011.
- [13] Perret, G., Rochman, D., Vasiliev, A., and Ferroukhi, H., “Neutron emission measurements of PWR spent fuel segments and prelim validation of depletion calculations,” *EPJ Web Conf.*, vol. 247, 2021.
- [14] I. Pázsit, “On the possible use of the neutron current in core monitoring and noise diagnostics,” *Annals of Nuclear Energy*, vol. 24, no. 15, pp. 1257–1270, 1997.
- [15] V. Arzhanov, I. Pázsit, and N. S. Garis, “Localization of a vibrating control rod pin in pressurized water reactors using the neutron flux and current noise,” *Nuclear Technology*, vol. 131, no. 2, pp. 239–251, 2000.
- [16] S. Avdic, P. Lindén, and I. Pázsit, “Measurement of the neutron current and its use for the localisation of a neutron source,” *Nucl. Instr. Meth.*, vol. A 457, no. 3, pp. 607–616, 2001.
- [17] P. Lindén, J. K. H. Karlsson, B. Dahl, I. Pázsit, and G. Por, “Localisation of a neutron source using measurements and calculation of the neutron flux and its gradient,” *Nucl. Instr. Meth.*, vol. A 438, no. 2-3, pp. 345–355, 1999.
- [18] C. Mori, T. Osada, K. Yanagida, T. Aoyama, A. Uritani, H. Miyahara, Y. Yamane, K. Kobayashi, C. Ichihara, and S. Shiroya, “Simple and Quick Measurement of Neutron Flux Distribution by Using an Optical Fiber with Scintillator,” *Journal of Nuclear Science and Technology*, vol. 31, no. 3, pp. 248–249, 1994.
- [19] K. Watanabe, M. Yamanaka, T. Endo, and C. H. Pyeon, “Real-time subcriticality monitoring system based on a highly sensitive optical fiber detector in an accelerator-driven system at the Kyoto University Critical Assembly,” *Journal of Nuclear Science and Technology*, vol. 57, no. 2, pp. 136–144, 2020.
- [20] F. Vitullo, V. Lamirand, J.-B. Mosset, P. Frajtag, O. Pakari, G. Perret, and A. Pautz, “A mm^3 fiber-coupled scintillator for in-core thermal neutron detection in crocus,” *IEEE Transactions on Nuclear Science*, vol. 67, no. 4, pp. 625–635, 2020.
- [21] E. L. Brunetto, F. Vitullo, V. Lamirand, K. Ambrozic, D. Godat, M. Buck, G. Pohlner, J. Starflinger, and A. Pautz, “High resolution measurements with miniature neutron scintillators in the SUR-100 zero power reactor,” *EPJ Web Conf.*, vol. 253, p. 04029, 2021.
- [22] J. Leppänen, “The Serpent Monte Carlo code: Status, development and applications in 2013,” *Annals of Nuclear Energy*, vol. 82, pp. 142–150, 2015.
- [23] G. F. Knoll, *Radiation detection and measurement*. John Wiley & Sons, 2010.
- [24] T. Natsume, H. Doi, F. Ruddy, J. Seidel, and A. Dulloo, “Spent fuel monitoring with silicon carbide semiconductor neutron/gamma detectors,” in *Reactor Dosimetry: 12th International Symposium*, ASTM International, 2007.
- [25] I. Pázsit and M. Kitamura, *Advances in Nuclear Science and Technology*, book section The Role of Neural Networks in Reactor Diagnostics and Control, pp. 95–130. Boston, MA: Springer US, 1996.

- [26] R. Rossa, A. Borella, and N. Giani, “Comparison of machine learning models for the detection of partial defects in spent nuclear fuel,” *Annals of Nuclear Energy*, vol. 147, p. 107680, 2020.
- [27] R. Rossa, A. Borella, and K. van der Meer, “Comparison of the SINRD and PDET detectors for the detection of fuel pins diversion in PWR fuel assemblies,” in *Proc. INMM-59 Annual Meeting*, 1 October 2018.
- [28] S. Grape, E. Branger, Zs. Elter, and L. Pöder Balkeståhl, “Determination of spent nuclear fuel parameters using modelled signatures from non-destructive assay and random forest regression,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 969, p. 163979, 2020.
- [29] M. Durbin and A. Lintereur, “Implementation of machine learning algorithms for detecting missing radioactive material,” *Journal of Radioanalytical and Nuclear Chemistry*, vol. 324, no. 3, pp. 1455–1461, 2020.
- [30] N. Shoman, B. Cipiti, T. Grimes, and B. Wilson, “Advances in machine learning for safeguarding a PUREX reprocessing facility,” in *INMM 61st annual meeting*, 2020.
- [31] I. Pázsit, N. S. Garis, and O. Glöckler, “On the Neutron Noise Diagnostics of Pressurized Water Reactor Control Rod Vibrations —IV: Application of Neural Networks,” *Nuclear Science and Engineering*, vol. 124, no. 1, pp. 167–177, 1996.
- [32] A. M. Durrant, G. Leontidis, S. Kollias, A. Torres, C. Montalvo, A. Mylonakis, C. Demaziere, and P. Vinai, “Detection and localisation of multiple in-core perturbations with neutron noise-based self-supervised domain adaptation,” *Proc. Int. Conf. MC 2021, American Nuclear Society*, 2021.
- [33] T. Tasakos, G. Ioannou, V. Verma, G. Alexandridis, A. Dokhane, and A. Stafylopatis, “Deep learning-based anomaly detection in nuclear reactor cores,” in *Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2021)*, October, 2021.
- [34] A. K. Jain, J. Mao, and K. M. Mohiuddin, “Artificial neural networks: A tutorial,” *Computer*, vol. 29, no. 3, pp. 31–44, 1996.
- [35] A. Abraham, “Artificial neural networks,” *Handbook of measuring system design*, 2005.
- [36] S. Sharma, S. Sharma, and A. Athaiya, “Activation functions in neural networks,” *towards data science*, vol. 6, no. 12, pp. 310–316, 2017.
- [37] K. Janocha and W. M. Czarnecki, “On loss functions for deep neural networks in classification,” *arXiv preprint arXiv:1702.05659*, 2017.
- [38] S. Vani and T. M. Rao, “An experimental approach towards the performance assessment of various optimizers on convolutional neural network,” in *2019 3rd international conference on trends in electronics and informatics (ICOEI)*, Tirunelveli, India, IEEE, April, 2019.
- [39] C. J. Werner, “MCNP6.2 Release Notes,” 2018.
- [40] S. Sitaraman and Y. Ham, “Characterization of a safeguards verification methodology to detect pin diversion from pressurized water reactor (PWR) spent fuel assemblies using

- Monte Carlo techniques,” in *Proceedings of the INMM 48th Annual Meeting, Tucson, AZ*, 2007.
- [41] Y. S. Ham and S. Sitaraman, “Partial defect tester: A novel approach to detect partial defects in pressurized water reactor spent fuel,” *Nuclear Technology*, vol. 175, no. 2, pp. 401–418, 2011.
- [42] Y. Ham, P. Kerr, S. Sitaraman, R. Swan, R. Rossa, and H. Liljenfeldt, “Partial defect verification of spent fuel assemblies by PDET: Principle and field testing in interim spent fuel storage facility (CLAB) in Sweden,” in *2015 4th International Conference on Advancements in Nuclear Instrumentation Measurement Methods and their Applications (ANIMMA)*, pp. 1–7, IEEE, 2015.
- [43] M. Abadi *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015. Software available from tensorflow.org.
- [44] F. Chollet *et al.*, “Keras.” <https://keras.io>, 2015.
- [45] T. M. Mitchell, *Machine Learning*. 1997.
- [46] M. Iggane, A. Ennaji, D. Mammass, and M. El Yassa, “Self-training using a k-nearest neighbor as a base classifier reinforced by support vector machines,” *International Journal of Computer Applications*, vol. 56, 2012.