

**Machine Learning for Nuclear Fission Systems:
Preliminary Investigation of an Autonomous Control
System for the MGEP**

by
Jarod Wilson

Submitted to the Department of Nuclear Science and Engineering
in partial fulfillment of the requirements for the degree of
Bachelor of Science in Nuclear Science and Engineering
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Author
Department of Nuclear Science and Engineering
May 16, 2019

Certified by.....
Benoit Forget
Professor of Nuclear Science and Engineering
Thesis Supervisor

Certified by.....
Kaichao Sun
Research Scientist and Reactor Engineer of the Nuclear Reactor
Laboratory
Thesis Supervisor

Certified by.....
Akshay Dave
Research Scientist of the Nuclear Reactor Laboratory
Thesis Supervisor

Accepted by
Michael Short
Class of '42 Career Development Associate Professor of Nuclear Science
and Engineering
Chairman, NSE Department for Undergraduate Theses

**Machine Learning for Nuclear Fission Systems:
Preliminary Investigation of an Autonomous Control System
for the MGEP**

by

Jarod Wilson

Submitted to the Department of Nuclear Science and Engineering
on May 16, 2019, in partial fulfillment of the
requirements for the degree of
Bachelor of Science in Nuclear Science and Engineering

Abstract

Commercial nuclear technology today is facing challenges due to both economic viability and concerns over safety. Next-generation reactors could potentially improve with respect to both concerns through recent advancements in computation and machine learning, through autonomous control systems which minimize human error. The MIT Graphite Exponential Pile (MGEP) has been selected as the basis of a real-world demonstration of such a system, because of its simple properties and inherent safety. This study evaluated the preliminary feasibility of an autonomous control system for the MGEP through two parallel avenues; a practical investigation of various machine learning algorithms applied to fission systems, as well as the design and fabrication of a control rod for the pile. It was found that Convolutional Neural Networks (CNNs) outperform Support Vector Regression (SVR) in predicting the MITR power-shape. Additionally, acceptable results were achieved when applying the CNN algorithm to the MGEP to predict the flux distribution of its fuel elements. Finally, it was verified that neutron detectors in the pile respond predictably to control rod insertions. Taken together, the groundwork for the further development of an autonomous control system has been laid, and the path forward is promising.

Thesis Supervisor: Benoit Forget
Title: Professor of Nuclear Science and Engineering

Thesis Supervisor: Kaichao Sun
Title: Research Scientist and Reactor Engineer of the Nuclear Reactor Laboratory

Thesis Supervisor: Akshay Dave
Title: Research Scientist of the Nuclear Reactor Laboratory

Acknowledgments

I would first like to thank Kaichao Sun from the MIT NRL; his support, guidance, and direction over the past two years allowed me to successfully complete my research on not one, but two continents. As an extension I must also thank Lin-wen Hu, who helped sponsor and oversee much of this research. I would also like to thank Professor Benoit Forget for his support, both for this project and over the course of my undergraduate career. He has served as an invaluable source of knowledge and guidance as I have progressed through MIT. This project was only possible with the help of many people; Bill McCarthy, Nick Acosta, and Yakov Ostrovsky made my experimental work possible, and Akshay Dave helped me better develop the machine-learning framework. Of course, I must also mention Professor Kord Smith, as without his efforts, no work on the Graphite Pile would have been possible. Finally I would like to thank Sara Hauptman, both for her collaborative efforts on this project, as well as her close friendship over the years.

Contents

1	Introduction	13
2	Background	15
2.1	MIT Nuclear Reactor Lab Facilities	15
2.1.1	MIT Research Reactor	15
2.1.2	MIT Graphite Exponential Pile	16
2.2	Monte Carlo Simulations	19
2.2.1	OpenMC	19
2.2.2	MCNP	19
2.3	Machine Learning Algorithms	20
2.3.1	Machine Learning Applied to Nuclear Systems	20
2.3.2	Support Vector Algorithms	21
2.3.3	Convolutional Neural Networks	22
3	Methods	25
3.1	Machine Learning Methodology	25
3.1.1	Training Data Generation	27
3.1.2	Model Generation and Evaluation	29
3.1.3	Test Cases	30
3.2	Control Rod Design, Fabrication, and Testing	32
3.2.1	Geometry	33
3.2.2	Material Choices	34
3.2.3	Construction	35

3.2.4	Verification of Detector Response	37
4	Results	41
4.1	Machine Learning Investigation	41
4.1.1	MITR: SVR and CNN Comparison	41
4.1.2	MGEP	44
4.2	Control Rod Verification for MGEP	45
5	Conclusion	55
A	Training Data Examples	57
B	SWX-237Z30 Full Material Properties	61
C	Neutron Count Rate Data	67

List of Figures

2-1	MIT Research Reactor	16
2-2	MGEP North Face	17
2-3	MGEP East Face	18
2-4	SVM Basics	21
2-5	Basic CNN Diagram	23
3-1	Machine Learning Block Diagram	26
3-2	MGEP OpenMC Modelling Source Locations	28
3-3	LIBSVM Methodology Block Diagram	29
3-4	MGEP Node Distribution.	31
3-5	Source Location for CNN Training Data.	32
3-6	MGEP Control Rod.	34
3-7	MGEP Barrier Tube.	35
3-8	MGEP Barrier Tube in Place.	36
3-9	Detector, Control Rod, and Source Locations for Control Rod Verification	38
3-10	Control Rod in Pile	39
4-1	ML Accuracy for MITR Shim Bank Movement	42
4-2	SVR Prediction for 18cm Shim Bank	47
4-3	SVR Prediction for 19cm Shim Bank	48
4-4	ML Accuracy for MITR Asymmetric Blade Insertions	49
4-5	SVR Accuracy for MITR Asymmetric Blade Insertion	49
4-6	CNN Accuracy for MITR Asymmetric Blade Insertion	50
4-7	CNN Accuracy for MGEP	51

4-8	Neutron Count Rate for Detector Positions	52
4-9	Neutron Count Rate for Detector Positions as Percentage	53
A-1	Fully Formatted MITR Training Data Example	58
A-2	Fully Formatted MGEP Training Data Example	59

List of Tables

3.1	Control Rod Assembly Dimensions	33
3.2	SWX-237Z30 Castable Shielding Material Description	37
4.1	ML Accuracy Comparison with and without Empty Fuel Channels . .	43
C.1	Count Rate Measurement Data	68

Chapter 1

Introduction

Many infamous reactor accidents, including Three Mile Island and Chernobyl, were caused in part because of errors made by their operators [1] [2]. While an extreme meltdown event is the most dangerous and rare scenario for a nuclear reactor, other more common problems stifle the competitiveness of nuclear technology for use in commercial power generation. Importantly, the presence of human variables adds a degree of uncertainty and variation which can be planned around, but ultimately lowers both operational performance and safety [3]. While humans excel at complex identification and decision-making problems, recent advancements in computational capacity coupled with advancements in Machine Learning (ML) have led to faster, more reliable control systems which can navigate such problems - the progress of self-driving vehicles serves as a clear example [4]. Of course, an autonomous nuclear reactor control system poses a significantly larger engineering challenge: it would need to be guaranteed to be able to properly respond to all possible operational scenarios without jeopardizing the safety of the reactor and surrounding environment. Even though the development and deployment of such systems are many years off, demonstration of full autonomous control for a simplified, low-risk reactor would be an important first step towards proving the feasibility of such designs.

The MIT Graphite Exponential Pile (MGEP) serves as an ideal low-risk "reactor" to demonstrate the concept of autonomous control. Seated in the Middle Lab at the MIT Nuclear Reactor Lab (NRL), the MGEP is a sub-critical graphite pile. As such, it

is a simple system made up of only graphite, various structural materials, and natural uranium as fuel. In its current state, with only naturally enriched uranium for fuel, it is physically impossible for the system to reach criticality. Importantly, this means the MGEP does not produce hazardous amounts of heat or radiation, even when fully fueled. Nonetheless, it is still a fission system, and as such can be used to demonstrate some of the key preliminary concepts of an autonomous control system. Specifically, this project seeks to verify that machine learning methods can accurately predict the power shape and flux of the MGEP in real time. Such verification is necessary, as current modelling methods based on Monte Carlo approaches generally take minutes to hours to run. While these simulations yield accurate models of static system behavior, such run-times are much too long for a real-time power control system. Additionally, because efforts to demonstrate autonomous control using the pile are in their nascent stages, some of the central elements of the system must be developed. Specifically, while no control devices are necessary to operate the MGEP, a control rod must be designed and fabricated in order to have a controllable, repeatable method to perturb the system. The functionality of such a control rod must also be confirmed after its fabrication.

Ultimately, this study will seek to investigate the preliminary functionality of an autonomous control system for the MGEP through two parallel avenues. An investigation of various machine learning algorithms and methods will allow the creation of a model which can provide accurate, real-time predictions about the pile's power shape, based on training data provided through Monte Carlo modelling. Additionally, the development of control rods with the pile along with concurrent verification of detector response with rod insertion will be done in order to lay the groundwork of the real-world physical system needed to demonstrate autonomous control. This project aims to demonstrate the preliminary feasibility of such a system and establish starting-points for future development. Ultimately this project will contribute to the demonstration of machine-learning based autonomous control of a fission system, advancing the possibility that commercial power systems of the future may be able to benefit from advances in machine learning.

Chapter 2

Background

2.1 MIT Nuclear Reactor Lab Facilities

The MIT Nuclear Reactor Laboratory (MIT-NRL), located at the heart of MIT's campus, operates a high-performance research reactor known as the MIT Reactor (MITR). The NRL seeks to provide state-of-the-art experimental facilities to MIT faculty and students, as well as the broader scientific community. Such facilities include beam ports on the MITR itself, irradiation facilities, radiochemistry labs, hot cells, and gamma activation analysis facilities [5].

2.1.1 MIT Research Reactor

The MITR is a 6MWt research reactor used for neutron activation experimentation and silicon doping. It is a light-water-cooled, heavy-water-reflected reactor which uses highly-enriched uranium for fuel [5]. The core of the reactor consists of 27 rhombic fuel-element channels, however some of the fuel channels may be kept empty or be used for experiments at any given time. The MITR is controlled by six shim blades and one fine-control regulating rod. A schematic of the upper view of the MITR core can be seen in Fig. 2-1.

In addition to its various research programs, the MITR serves as an educational and experimental tool for use by MIT faculty, students, and affiliates. The primary




Figure 2-1: Top-down schematic of the MITR core showing the fuel element channel and control mechanism locations.

advantage of the MITR's experimental facilities is the reactor's harder neutron spectrum; the maximum neutron fluxes available are 1.2×10^{14} and 6×10^{13} neutrons/cm² for fast and thermal flux, respectively. This flux is numerically close to those found in commercial power reactors, allowing for better irradiation experiments [5].

2.1.2 MIT Graphite Exponential Pile

The MIT Graphite Exponential Pile (MGEP) is a sub-critical reactor designed as a graphite pile, as the name implies. The pile is constructed primarily from two materials; graphite as the structural and moderating component, and natural uranium as the fuel. The graphite section of the pile has two components, a lattice and a pedestal. The lattice of the pile is a cube with sides roughly 91" in length. Graphite blocks no more than 4" tall and 4" wide run in perpendicular planes to form this shape. No graphite block runs the entire length of the lattice due to material availability restrictions present when the pile was constructed; the longest blocks run half the length of the pile. The pedestal also consists of perpendicular planes of graphite blocks. The pedestal sits below the lattice and serves to thermalize neutrons emitted from a neutron source which can be inserted into one of twenty-six cylindrical channels that run the length of the pedestal. The lattice also has a variety of channels. A grid of 12x12 fuel channels run horizontally through the length of the pile in the North-

South direction, and various fuel and stringer channels penetrate half the length of the pile in the East-West direction. Fuel for the pile consists of 1,288 cylindrical fuel slugs. Each slug is roughly 1" in diameter and 8" in length. Each fuel slug is simply metallic natural uranium clad in aluminum. Pictures of the pile's North and East faces can be seen in Figures 2-2 and 2-3 respectively.




Figure 2-2: The North face of the MGEF. This picture shows both the lattice and pedestal components of the pile. In the lattice, the 12x12 fuel channel matrix can be seen, as well as the foil stringer channels, which run down the center blocks of the face. The neutron source channels can be seen in the pedestal of the pile [6].



Figure 2-3: The East face of the MGEF. This picture shows the upper segment of the lattice on the East face of the pile. The detector channels as well as the foil stringer channels can be seen. The channels only run half the length of the pile [6].

Constructed in 1957, the MGEF was initially used as a source of research and study for several Masters-level theses. However, by 1976 the existence of the pile was not common knowledge, despite intermittent use through 1991. Recent efforts starting within the past two years have revamped the pile and prepared it for use once again. Importantly, MIT NSE student Micah Gale's Bachelor's thesis extensively catalogued the pile, providing an accurate resource for the pile's geometry, material makeup, and potential for experimental use. Currently, the MGEF is used as an educational tool for reactor physics classes at the undergraduate and graduate level within the MIT Department of Nuclear Science and Engineering [6].

2.2 Monte Carlo Simulations

Monte Carlo simulations are a broad class of computational algorithms used to model a variety of processes. In the most general sense, these simulations rely on repeated random sampling and statistical analysis to achieve numerical results [7]. Monte Carlo simulations are especially useful in nuclear engineering applications due to the stochastic nature of particle interactions with matter. One of the more popular examples of such simulation software is Geant4, a Monte Carlo-based toolkit developed by CERN to model particle interactions with matter [8]. Other Monte Carlo-based codes have more specificity for their use. Particularly useful for this project are OpenMC and MCNP, two neutronics-modelling codes which are often employed to model fission reactors.

2.2.1 OpenMC

OpenMC is an open source Monte Carlo code developed at MIT with a focus on fission reactor systems and reactor physics. Because of its open-source nature, OpenMC was designed with usability in mind. Inputs for OpenMC are generated through the programming language Python, allowing for faster, more intuitive, and more flexible data generation. OpenMC currently only models neutral particles (neutrons) moving through a defined model which represents a real-world experiment. As such, OpenMC is capable of modelling nuclear reactor fission systems like the MITR and MGEP [9].

2.2.2 MCNP

MCNP, or Monte Carlo N-Particle transport code, is a particle interaction modelling code developed by Los Alamos National Laboratory. MCNP is a robust, widely-used transport code which can be used to model neutron, electron, and photon transport. While MCNP has a variety of use-cases including radiation shielding, accelerator design, and medical radiation physics, it is commonly used for fission reactor modelling and design. Because of its more comprehensive validation, MCNP is typically the reference code used in rigorous reactor and neutronics modelling scenarios [10].

2.3 Machine Learning Algorithms

The past several years have seen a significant amount of excitement surrounding the field of machine learning. Significant increases in the availability of computing power spurred renewed interest in the field, and now one can readily find examples of machine learning-based analysis being applied to wider problems in almost every field. While ML has become somewhat of a buzzword, the more realistic applications of ML to solve issues are still very exciting.

At their core, ML methods serve as a way to more efficiently or effectively navigate large sets of data. They allow for the substitution of explicit domain knowledge of a problem with an adequately large data-set which possesses the desired outcome or characteristics [11]. This is accomplished through various algorithms which parse an input set of data, called the training set, and calculate patterns within it in order to generate a model. These algorithms can be highly specific or general-purpose. Despite the wide variety in algorithms and approaches, ML methodology can be separated into three broad techniques; supervised learning, unsupervised learning, and reinforcement learning. In supervised learning, the training data used contains both inputs and desired outputs, and the algorithm attempts to generate a model which best replicates the relationship between the two. In unsupervised learning, the inputs in the training set contain no desired outputs, as the goal of the algorithm is generally to locate underlying patterns or structures within the data. In between the two is reinforcement learning, where supervision occurs after the selection of an output, forming a feedback loop which evaluates and improves the performance of the model [11].

2.3.1 Machine Learning Applied to Nuclear Systems

Excitement over the potential applications of ML approaches to specific engineering challenges has not skipped over the field of nuclear science and engineering. A team at Oregon State University applied an artificial neural network to their Multi-Application Small Light Water Reactor. The team found that this approach was able to predict the reactor's behavior during core input events and loss of flow accidents

with good accuracy [12]. Similarly, a team in Croatia managed to apply a support vector machine algorithm to better solve reactor core loading patterns [13]. A joint group consisting of researchers from Tsinghua University and MIT applied support vector machines for use in a conceptual autonomous control design for a small, high-temperature salt-cooled reactor [14]. These examples not only highlight the exciting research which is applying the promise of ML techniques to nuclear technology, but they also provide two promising ML algorithms which can be investigated for use with the MGEF.

2.3.2 Support Vector Algorithms

Support Vector Machines (SVMs) are a specific set of ML algorithms developed originally by Bell Laboratories. In very basic terms, SVM algorithms serve to fit a set of data and minimize error. If a training set of data is provided in the form of $[(x_i, y_i), \dots, (x_n, y_n)]$, an ϵ -SV algorithm seeks to find a function $f(x)$ which deviates ϵ at most from the real set of values (y_1, \dots, y_n) , while remaining as flat as possible. Kernel functions allow the SVM algorithm to work and fit models in multi-dimensional, non-linear situations. From the produced function, predictions of unseen data not in the training set can be made, and evaluated for accuracy. SVMs are useful and have achieved acceptable results in both regression and classification problems [15].




Figure 2-4: The soft margin loss setting for a linear SVM. This figure shows the simplest strategy ϵ -SVMs employ to fit a set of data. Kernel functions allow the methodology to be applied to multi-dimensional cases [15].

LIBSVM

LIBSVM is an open source SVM library developed at the University of Taiwan. It employs the Sequential Minimal Optimization (SMO) algorithm for kernelized SVMs. While the library itself was developed in C++, packages now exist that make its use possible in a variety of programming languages, including Python and MATLAB [16].

2.3.3 Convolutional Neural Networks

Convolutional Neural Networks (CNNs) are a subclass of Artificial Neural Networks (ANNs). ANNs are a class of machine learning algorithms modeled off the connection and function of biological nervous systems. ANNs are comprised of a large number of computational nodes which are connected. These nodes use stochastic weighting to learn from a set of input data and optimize their output according to specified parameters. A traditional ANN structure consists of an input layer, hidden layers, and an output layer. The input layer is comprised of a chosen amount of computational nodes that feed into one or more hidden layers, each of which learns from the prior layer and assesses how a stochastic change in itself improves or worsens the model's performance. The hidden layers eventually feed into an output layer which properly formats the data to the desired product, which could be a classification or regression. A simplistic visualization of this process can be seen in Fig. 2-5. CNNs are by and large the same as ANNs, in that they consist of nodes which self-optimize through the learning process in order to generate an output. The main difference between the two is that CNNs are primarily used for pattern recognition in images [17].

Tensorflow and Keras

Tensorflow is an open-source library used for dataflow and differential math. Developed by Google, it is also widely used for machine learning applications [18]. Keras is an open-source API which aims to make the use of ML software and algorithms more accessible. Written for use with Python, Keras can run ML applications using Tensorflow, as well as CNTK and Theano [19].




Figure 2-5: The basic structure of an ANN or CNN. An input layer consisting of a set number of nodes feeds into any number of hidden layers which can be any dimension. The layers stochastically weight themselves to train to a desired parameter, before providing a model through the output layer [17].

Chapter 3

Methods

3.1 Machine Learning Methodology

Two specific ML approaches were selected for this investigation; Support Vector Machines (SVMs) and Convolutional Neural Networks (CNNs). The prior was selected because of the existing rich body of work which exists on the use of SVMs in nuclear-reactor systems. CNNs were selected for testing due to their rising prominence as an efficient and accurate approach to complex data problems, as well as the large online open-source community surrounding their support and application. Because the SVM algorithm was used to perform a regression, it will be hereafter be referred to as Support Vector Regression (SVR).

A block diagram of the overarching data-generation, model training, and model assessment process for a generalized ML framework can be seen in Fig.3-1. To summarize, training data was generated from a Monte Carlo neutronics code to model the MITR core or MGEP. These cases, which are assumed to be accurate models of the specific system, served as training data for the ML algorithms. Data generated from the Monte-Carlo codes was randomly split into training and test sets for model training; the latter is used for the ML algorithms to cross-validate as they train on the data-set. The model was then applied to make predictions about data both in and outside of the training set. Finally, the accuracy of these predictions was evaluated through regression analysis, based off of the corresponding Monte Carlo

run. The coefficient of determination (R-squared) metric was chosen to compare the "true" Monte-Carlo simulations and predictions provided by the ML algorithm. The R-Squared value is simply the square of the correlation (R) value for a true set of output values and the predicted set of values. 1 is the best R-Squared score, indicating a perfect correlation, whereas 0 indicates no correlation. The Python package used to perform this analysis can also result in negative scores, which reflect a set of prediction values arbitrarily worse at capturing the true values when compared to simply using the true values' mean.

The approach taken for this study emphasized the practical evaluation of the ML models; standard training procedures and structures were followed with little attempt for optimization for the model output. In the case of the MITR, this was to quickly assess which ML algorithm was more accurate and applicable for power shape or flux shape predictions. In the MGEP case, this was to simply gain an initial understanding of the limitations and abilities of ML models to predict the behavior of the pile under varying conditions.

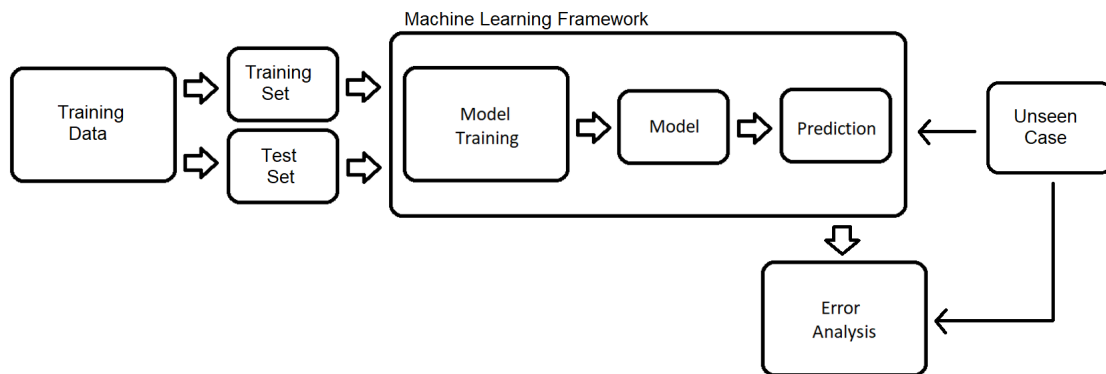


Figure 3-1: Block diagram of methodology used for the overarching data-generation, training, and prediction process for a generalized machine learning framework.

3.1.1 Training Data Generation

MITR Core Neutronic Modelling

The modelling input code for the MITR core was provided by Kaichao Sun, the Reactor Physics Analysis Engineer at the NRL. In this model, each fuel element channel acted as a cell. In addition, each fuel element channel was separated into sixteen cells vertically, thus creating a total of 432 nodes. This relatively sparse mesh was selected as the limitations of SVR methods for 3D power shape prediction were unknown, and as such a simpler model would serve as a starting point to evaluate prediction accuracy and training time.

The core was modelled under identical parameters except for the position of the shim blades. Training data from MCNP was generated for a large number of shim blade positions. Modelling scenarios were run for the shim blades between 18 cm and 30 cm, as this is the more pertinent operating range of positions for the reactor, and is also the more neutronically sensitive region of the reactor core. Modelling was done for the movement of a uniform shim bank, where all six blades are at the same position, for the 18-30cm range. Additionally, MCNP modelling was done for single blade insertion, where a single shim blade is at every 1cm increment in the 18-30cm range while the remainder of the blades are at 24cm, for each blade. This provided a total of 73 separate MCNP runs to serve as training data. The output data generated provided a normalized fission power value for each node.

MGEF Neutronic Modelling

Modelling output data for the MGEF was provided by Sara Hauptman, a fellow undergraduate student in the Department of Nuclear Science and Engineering. In this model, done in OpenMC, cells were determined by the physical geometry of the pile. Some simplifying approximations were made to reduce the total number of cells and simplify the geometry input for the Monte Carlo code. Each graphite moderating block and fuel channel block in the North-South direction was modelled as a continuous cell throughout the length of the pile. Blocks in the East-west direction

were modelled as two half-length cells, to account for the difference in channel type on either side. The fuel was also modelled as a single, continuous cylindrical cell for each fuel channel, comprised of only natural uranium.

The pile was modelled identically except for the position of the neutron source within the pedestal. Modelling was done with the source at the same height within the pile, as all of the source channels are located within the same layer. The source was modelled at a variety of positions to form a grid over the length and width of the pile, as can be seen in Fig. 3-2. The output data generated provided the neutron flux for each cell. There were a total of 650 different source locations modelled.

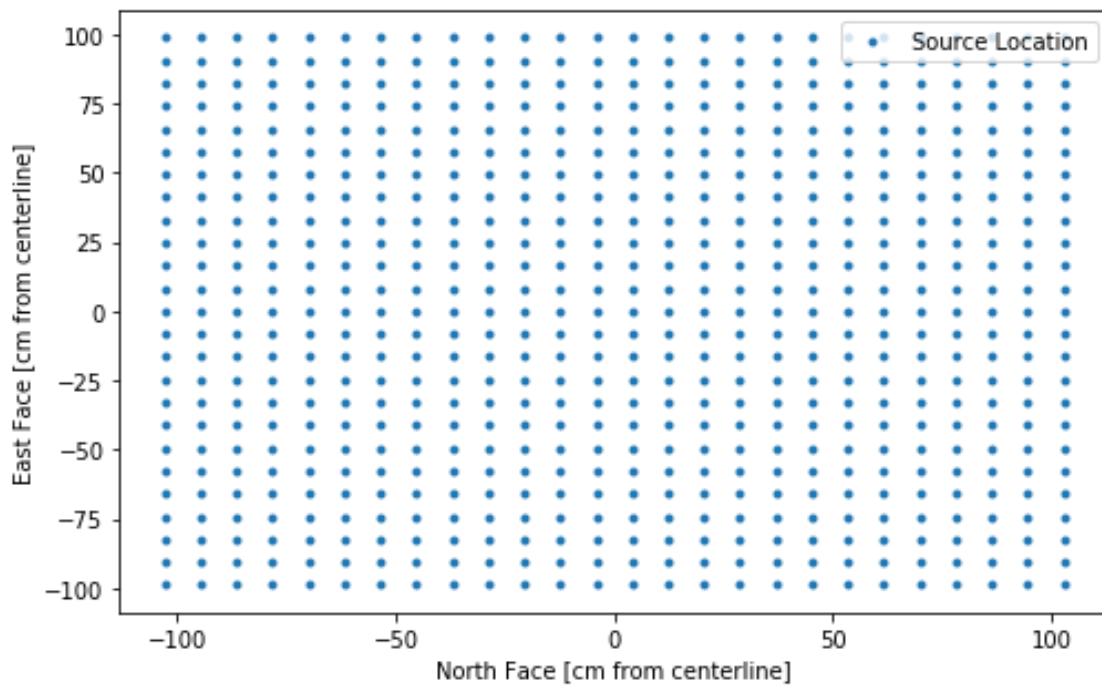


Figure 3-2: The source locations of for the OpenMC output data provided by Sara Hauptman. The height of the source for each Monte Carlo run is the same, as all source channels are located within the same layer of blocks in the pedestal. Variation in source location in the North-South direction is attained by modelling the source as further into the channel, whereas variation in the East-West direction is attained by modelling the source as in different channels within the pedestal.

3.1.2 Model Generation and Evaluation

LIBSVM

The data generated by MCNP was formatted to be used with the python binding of LIBSVM. Specifically, the 3-dimensional node power distribution was compressed to 2 dimensions, such that each node could be represented by its fuel element number and its height. This was done to allow simpler reading and manipulation of the data. An example fully-formatted run can be seen in Appendix A. To train the SVR model, the general approach described by Hsu, Chang, and Lin in *A Practical Guide to Support Vector Classification* was followed. A flowchart of this process can be seen in Fig. 3-3. The methodology of this approach was chosen as it outlines practical choices in the training strategy which tend to yield accurate results [16].




Figure 3-3: Block diagram of the LIBSVM methodology. This is the general procedure outline by Hsu, Chang, and Lin for a simple, practical application of SVRs.

Once the training data provided by MCNP modelling of the MITR core was generated, formatted, and scaled, 2-fold cross-validation was applied in conjunction with a course grid search of C and γ parameters in order to ascertain which values would yield the most accurate results. The ϵ -regression option was used, with an ϵ value of 0.01 selected. This grid search was done simply by enumerating over a range of values for each parameter, evaluating the model's performance for each step. After the entire range was tested, the best C and γ values were used. Once the ideal parameters were obtained, a model was trained using the selected training data. This model could then be used to make predictions on data outside of the training set.

Keras and Tensorflow

The data generated by OpenMC was formatted to be used with the python binding of Keras using the Tensorflow back-end. For the MITR training data, like the LIBSVM formatting, the 3-dimensional nodal power distribution was compressed to 2 dimensions, such that each node could be represented by its fuel element number and its height. This was done to allow simpler reading and manipulation of the data. For the MGEP, because the geometry of the pile and its model are not uniform in the number or location of cells from layer to layer, a 2-dimensional nodal structure was imposed, as can be seen in Fig. 3-4. Each node was assigned to have the neutron flux value of the fuel cell in the volume it occupied. As such, the model was essentially trained on and able to make predictions for the neutron flux distribution for fuel elements within the pile. This approach for modelling the pile was taken as it was the simplest scenario with the fewest amount of data points. This allowed for rapid application and testing of the CNN algorithm for predicting the pile's flux profile. For both the MITR and MGEP, using a 2-dimensional matrix structure for the data closely mimics that of a picture, which essentially forms a 2-dimensional matrix of RGB pixel values. This similarity was useful, as CNNs are primarily used for image identification and classification. An example fully-formatted run can be seen in Appendix A. For both the MITR and MGEP cases, a dense model was trained. This model consisted of one dense input layer, one dense activation layer using the relu activation setting, seven dense layers, and one output layer. The model used the adam optimizer, the mean squared error loss function and the accuracy metric.

3.1.3 Test Cases

MITR

ML modelling of the MITR core was split into two test cases, both of which trained an SVR and CNN model. Both ML algorithms were used to ascertain which one would be better-suited for use on the MGEP data, in terms of speed and accuracy.

The first test case used modelling data which represented uniform movement of




Figure 3-4: Visualization of nodes imposed on the MGEP neutronics modelling data. This was done in order to be able to format the MGEP data to be compatible with the CNN ML algorithm. Each red sphere represents one node which was attributed a position and flux value, based on the fuel element whose volume the node center occupied. As such, the model training and testing for the MGEP case was essentially evaluating the neutron flux distribution of the fuel elements.

the entire shim bank (all six shim blades at once). Only runs from even steps in the 18-30cm range were used. As such, there were only 7 MCNP modelling runs used for training data in the SVR and CNN models. Predictions were then made for shim bank heights for every integer step over the entire range.

The second test case used the full set of training data generated for the MITR, as described in Section 3.1.1. Again, the training data selected used runs from the even steps within the 18-30cm range, for each shim blade insertion. Predictions were then made for shim blade heights for every integer step, for every blade.

MGEP

Only one test case was evaluated for the MGEP, using only the CNN algorithm. This is because the CNN algorithm was initially found to be far superior to the SVR algorithm, both in terms of speed and accuracy. Half of the output data was

provided was used to train the model. Predictions were then made for the flux distribution based off the source location for every run. The selection of source locations used solely for prediction evaluation as well as those used for training and prediction evaluation can be seen in Fig. 3-5.

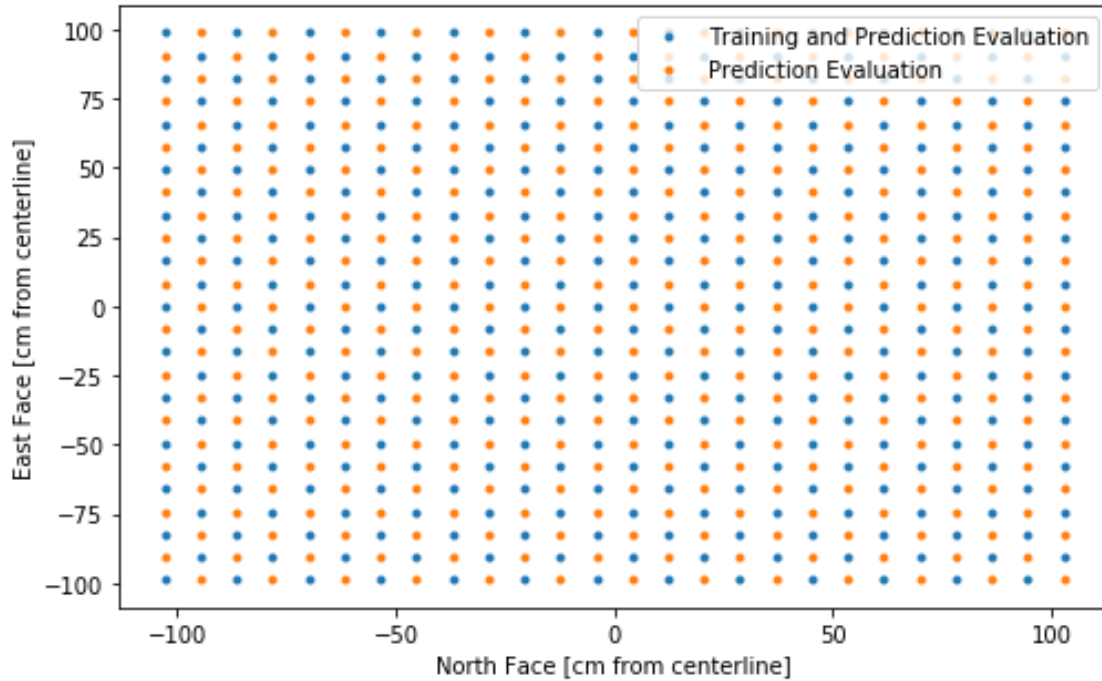


Figure 3-5: Figure showing source locations selected for CNN model training and prediction. Those in blue served as the training set of data used to train a model for the MGEP case. Predictions were then made for the flux distribution for every source location - both those shown in blue and orange.

3.2 Control Rod Design, Fabrication, and Testing

Because the ultimate goal of this project and the research that will come after it is to build a system which can model the pile and manipulate it in real time, a control mechanism is necessary. Currently, there are no control rods that exist for the pile, and only the neutron source type/location and fuel loading pattern can change the flux profile and power shape of the pile. Changing either of these is a time-intensive process, however. As such, a movable control rod was designed, built, and verified for work with the MGEP.

Table 3.1: Control Rod Assembly Dimensions

Part	Length [in]	Outer Diameter [in]	Inner Diameter [in]	Thickness [in]
Barrier Tube	45	1.25	1.12	0.065
Flange	3	N/A	1.12	0.125
CR Tube with Cap	46	1	0.87	0.065
Handle	4.5	0.75	N/A	N/A

3.2.1 Geometry

The required geometry of the control rod assembly was determined by the geometry and characteristics of the MGEP. Primarily, the control rod had to be large enough that there would be sufficient neutron absorbing material within it to alter the behavior of the pile. As such, the rod was designed to fit within the pile's fuel channels, as these are its largest penetrations. Additionally, there was concern over the rod contaminating the pile with absorbing material, or the removal of the rod causing the graphite of the inner fuel channel to wear away over time. As such, it was decided that a barrier tube would need to be created as part of the control rod assembly.

The control rod assembly consists of two extremely simple components. The barrier tube is a 1.25" OD open tube of length 45". The thickness of the tube wall is 0.065". A 3" square flange of thickness 0.125" was welded to one end of the tube to create a stop for the tube. The flange has one through-hole in each corner, which allows for the possibility of tapping into the graphite and fixing the barrier tube to the pile with screws more permanently. The control rod itself is simply a 1" OD open tube of length 45.5", also with a wall thickness of 0.065". The tube is threaded on both ends to allow a handle and cap to be affixed to it. It was decided that a castable shielding material would be filled into this tube to act as the absorber material for the control rod. With the cap and handle attached, the length of the control rod that was filled with shielding material is roughly 45". The length of the control rod with the cap attached is 46", and the length of the control rod assembly with the cap and handle attached is 50.5". Pictures of the control rod and barrier tube can be seen in Figures 3-6, 3-7, and 3-8. A summary of the various part dimensions can be seen in Table 3.1.




Figure 3-6: The control rod for the MGEP. This rod was filled with castable shielding material which will act as a neutron absorber when inserted into the pile.

3.2.2 Material Choices

For the tubing, multipurpose Al-6061 aluminum was selected as the material choice primarily because it is cheap and easily machinable, and comes readily available in a variety of tube sizes. Aluminum as a cladding material for the control rod and barrier tube was also advantageous as the cladding for the fuel slugs used in the pile is also aluminum. As such, there was no fear of the aluminum changing the behavior of the pile in any unanticipated manner.

For the neutron absorbing material, Shieldwerx SWX-237Z30 High Temperature Boron-Silicone castable shielding was chosen. This material was chosen because of the freedom it provided as a simple yet effective option for control rod fabrication. The SWX-237Z30 mix comes in a five-part set, where four parts are liquid, and one is powder. To prepare the castable shielding only required simple hand mixing. It



Figure 3-7: Barrier tube for the MGEP. This tube is designed to act as a stationary buffer between the control rod and the graphite of the fuel channel it is inserted into, to prevent wear on the inner channel and potential contamination of the pile.

could then be directly poured into the control rod tubing. A summary of the nominal elemental analysis of the shielding material can be seen in Table 3.2. The full material properties and specifications as provided by Shieldwerx can be seen in Appendix B.

3.2.3 Construction

After initial concepts for the control rod assembly were created, Yakov Ostrovsky, a Research Engineer at the NRL, was consulted for assistance with the final design and part ordering. The tubing was then sent to the MIT Central Machine shop for preparation, which included cutting the tubing to the proper lengths, welding a flange to the barrier tube, threading either ends of the control rod tube, and fabricating the cap and handle of the control rod tube. Mixing and casting of the SWX-237Z30 mix



Figure 3-8: The barrier tube for the MGEP inserted into a detector channel on the East face of the pile.

Table 3.2: SWX-237Z30 Castable Shielding Material Description

Element	Percent by Weight
Oxygen	23.39
Aluminum	5.67
Silicon	17.57
Carbon	19.08
Hydrogen	3.27
Boron	30.51
Sodium	0.04
Iron	0.18
Zinc	0.1
Nitrogen	0.17

was done in the MIT NRL receiving garage.

The mass of the control rod was measured before and after the shielding material was cast and cured. The total mass of material poured into the control rod is 701.1g. As the volume of the inner space of the control rod is 438cm³, the density of the shielding material in the control rod is 1.60g/cm³. The stated density of the cured shielding material by Shieldwerx is 1.59g/cm³, as can be seen in Appendix B. The small discrepancy in these values can be attributed to error in mass measurement, and as such it can be determined that the shielding material was properly mixed and cast.

3.2.4 Verification of Detector Response

A simple test was done to verify that the control has sufficient neutron absorbing material to alter the behavior of the pile as well as to verify that the perturbations it causes are detectable by the He-3 neutron detectors currently set up for use with the pile. The detectors themselves are Quaesta Neuchrometers. They were inserted to the center of the E2, E5, E8, E11, G1, H2, H5, H8, and H11 fuel channels, as shown in Fig. 3-9. These locations were selected to test detector response at a variety of locations, both above and below the control rod insertion location, as well as varying distances from it. The count-rate was measured for a detector at each of these locations, for the control rod fully inserted (46"), inserted halfway (23"), and fully removed (0").

The control rod itself was inserted into the central channel of the East face, as can be seen in Fig. 3-10. It is important to note that the control rod is only half the length of the pile, so fully inserted in this context refers to the control rod. Full insertion of the control rod only penetrates to half the length of the pile. A ^{137}Cs source was placed in the central position of the pedestal for this experiment.

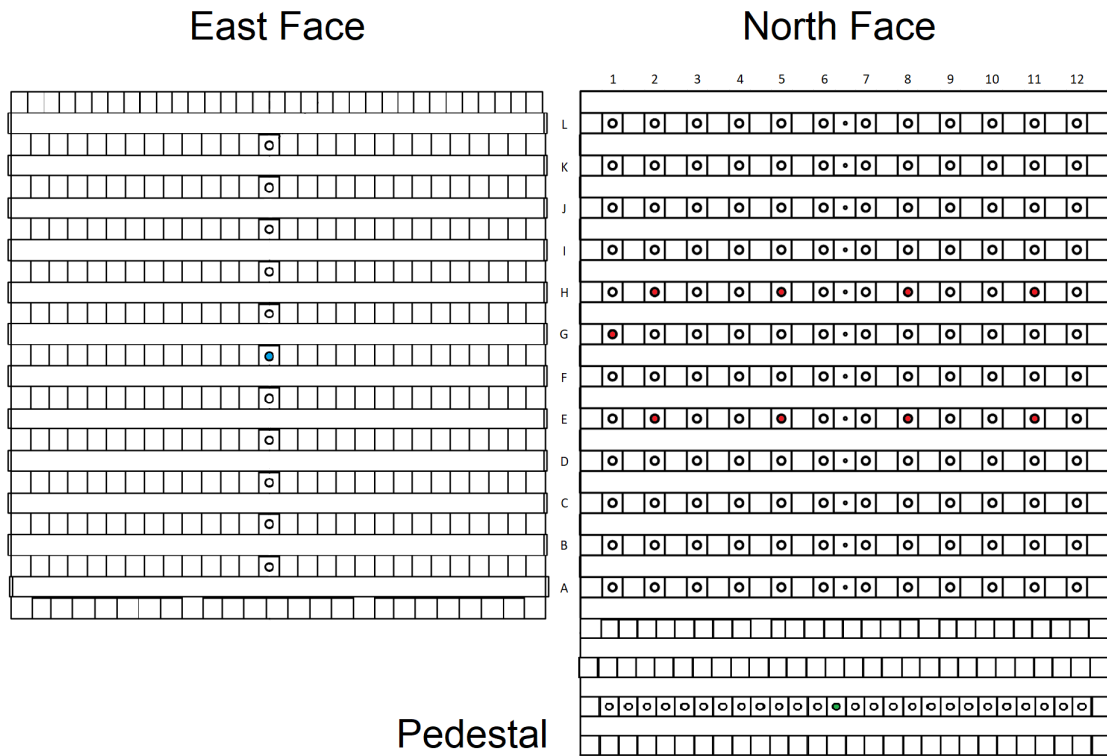


Figure 3-9: The chosen neutron detector locations, control rod location, and source location for the verification of detector response to control rod insertion. The selected fuel channels that the detectors were inserted to the center of in the North face are denoted in red, the channel used for rod insertion in the East face in blue, and the source location in the pedestal in green.



Figure 3-10: A picture showing the control rod fully inserted into the MGEP's East face, in the central detector channel which penetrates half the length of the pile.

Chapter 4

Results

4.1 Machine Learning Investigation

4.1.1 MITR: SVR and CNN Comparison

Error analysis for the test case comparison of SVR vs. CNN performance for the simple case of uniform shim bank movement can be seen in Fig. 4-1. For shim bank positions which were included in the training set, performance between the two models was comparable, However, for shim bank positions outside of the training set, while the performance of the CNN model stayed relatively constant, performance of the SVR model depreciated significantly. A comparison of the MCNP output and SVR prediction for the bank at 18cm and 19cm can be seen in Figures 4-2 and 4-3 respectively. Here the poor performance of the SVR model can be seen more clearly; for the case of the bank at 19cm, which was outside of the training set, the model struggles to predict the empty fuel element cells with value zero. Its prediction for every other cell is also significantly off. These discrepancies explain the negative R-Squared values for the odd shim bank positions. It appears as though for a small training data set, the SVR algorithm struggles to make predictions for cases outside of that set. For these results, it is likely that training the SVR model with the empty fuel channels included resulted in a smearing effect, reducing the overall power of the core. Training the model on data which does not include these cells would likely

result in better accuracy. It is also worth noting that the SVR model took 3.2hr to train, whereas the CNN model took 22sec. Both models were able to make individual predictions in 1ms or less, making both practical for real-time use. In this simple test case, the CNN model outperformed the SVR model in terms of accuracy and speed.

To see how much the poor prediction of the empty channels dominated R-Squared analysis, this data was removed from the MCNP runs and ML predictions. Another R-Squared analysis was done on this set of data. The results can be seen in Table 4.1. As can be seen, there are mixed results. For the CNN predictions, the accuracy was boosted over the whole range. For the SVR predictions, accuracy was higher for shim bank positions within the training set, but the same or slightly lower for those outside it. As such, poor performance of the SVR algorithm cannot be solely attributed to its inability to capture the zero channels.




Figure 4-1: The R-Squared score of an SVR and CNN prediction against the MCNP run which modelled the corresponding shim bank height, for every 1cm step of the shim bank.

Error analysis for the test case comparison of SVR vs. CNN performance for the full set of training data generated can be seen in Figures 4-4, 4-5, and 4-6. Predictions

Table 4.1: Comparison of the R-Squared scores for CNN and SVR prediction of uniform shim bank movement, with and without empty channels. This analysis was done to see the effect the empty channels had on the overall R-Squared score, with the hypothesis being that the SVR’s inability to predict those values as zero disproportionately affected the model’s accuracy for the rest of the distribution.

Bank Position [cm]	Empty Fuel Channels Included		Empty Fuel Channels Removed	
	SVR	CNN	SVR	CNN
18	0.809920304	0.762868608	0.9939931	0.936397724
19	-0.367481628	0.81282417	-0.45100018	0.997542787
20	0.809737368	0.814007926	0.993768587	0.998951311
21	-0.336253071	0.81339756	-0.412674223	0.99824177
22	0.809887624	0.812960711	0.993952993	0.997721854
23	-0.300541156	0.812650826	-0.368845964	0.997337835
24	0.810939554	0.813190932	0.995243999	0.998008976
25	-0.267954796	0.812588117	-0.328853613	0.997221098
26	0.811119938	0.813620568	0.995465378	0.998496948
27	-0.268282703	0.813357898	-0.329256044	0.998255609
28	0.811102267	0.81344184	0.995443691	0.998264152
29	-0.267302231	0.812676346	-0.328052738	0.997324682
30	0.810737145	0.811565925	0.994995587	0.996027834

were made for each shim blade at every integer step in the range of 18-30cm, while the remaining blades were at 24cm. There is a clear divergence in the performance of the two models. While there is a high degree of variation in the performance of the SVR model over the full range of heights, accuracy from blade to blade at the same height is highly consistent. The CNN model performs consistently, both across the range of heights, and blade to blade. Additionally, predictions from the CNN model have a significantly higher R-squared score when compared to those from the SVR model. In this case, the SVR model took 5.4hr to train, whereas the CNN model took 89sec. Again, both models were able to make individual predictions in 1ms or less, making both practical for real-time use. As can be seen, the CNN model outperformed the SVR model significantly, both in terms of accuracy and speed.

Because this study mainly seeks to gain an understanding of the preliminary performance of common ML algorithms for use in fission systems, the reasoning behind the large discrepancy in performance between the SVR and CNN model seen here cannot be fully explained. As noted before, it is possible that the formatting of data into 2-dimensional matrices complimented the performance of the CNN algorithm, which is typically used in image classification. It is also important to note that minimal efforts were taken to optimize the performance of either model. It is possible that such efforts could limit the discrepancies in performance, or boost the overall accuracy of both models. While optimization could make the SVR's accuracy comparable to the CNN's, the drastic difference in training time most likely cannot be reduced, making the CNN preferable on that basis alone.

4.1.2 MGEP

Error analysis for the performance of the CNN algorithm for predicting the flux distribution of the fuel elements based on source location in the MGEP can be seen in Fig. 4-7. As can be seen, the CNN model predicts the fluxes of the fuel elements quite well for the majority of source locations. However, near the boundaries of the source location grid, especially in the corners, the accuracy of the model decreases significantly. This is most likely due to a variety of factors. Importantly, because on

the edge cases the model has less training data which demonstrates the pattern of the flux shape for a source location and those that surround it, it inherently has less information through which the model can train to predict that area. As such, this degradation of performance on the edges is expected. Nonetheless, neutron sources are rarely inserted into positions far from the center of the pedestal. Because of this, the model performs best in the most relevant area of source location. Additionally, the average R-squared score for all source locations is 0.78. While the model is not highly accurate overall, over the entire range of source locations it attains respectable accuracy in its predictions. It is also highly accurate, with an R-Squared score of 0.95 or higher, for a large number of predictions, in the area where source location is more pertinent. As such, this preliminary investigation reveals that a CNN-based ML approach has high potential to predict the MGEP's behavior. While this test case was relatively simplistic, further work to improve accuracy and predict the flux or power shape for more cells within the pile would likely achieve favorable results. Such improvement could likely be gained by better optimization of the model training parameters.

4.2 Control Rod Verification for MGEP

The results of the neutron detector measurements for various detector locations can be seen in Figures 4-8 and 4-9. As can be seen, the insertion of the control rod does have an appreciable and statistically significant effect on the measured count rate for various detector positions within the pile, wherein the insertion of the control rod lowers the neutrons counted due to it absorbing them. Additionally, a local and global effect for the MGEP system can be seen. For detectors in positions closer to the rod and above it, namely positions G1, H2, and H5, the reduction in count rate is most significant, with a 23.3, 16.9, and 13.2 percent reduction in count rate from the rod full out to full in respectively. This is due to the rod shielding the detector from the neutron source in the pile. The count rate reduction for detector positions E2 and E5 (7.0% and 5.0% respectively) is apparent despite their proximity

to the source, as they are also close to the control rod insertion point. For detector positions H8 and H11 the global effect of the control rod insertion can be seen in a 6.0% and 3.6% respective count rate reduction, as in these positions the detectors were no longer shielded from the neutron source by the control rod. Overall it has been confirmed that inserting the control rod does have the expected, measureable effects on neutron flux within the pile. Such effects are more pronounced closer to the rod insertion point, and the position of the rod and detector relative to the neutron source is important. Nonetheless, a control rod assembly which perturbs the pile and is compatible with the current detector setup has been manufactured and is ready for use. The full set of count rate data measured can be seen in Appendix C.

Vertical Node Number	Fuel Element Label																MCNP	SVR Prediction									
	A-1	A-2	A-3	B-1	B-2	B-3	B-4	B-5	B-6	B-7	B-8	B-9	C-1	C-2	C-3	C-4			C-5	C-6	C-7	C-8	C-9	C-10	C-11	C-12	C-13
1	0	19906.3	0	15972.2	14973.9	0	13746.3	15165.7	0	15906	14695.8	0	9002.93	9973.75	8146.22	10111.6	10286	9501.66	9801.49	8340.33	10095	10128.1	9299.76	9151.54	8133.02	9601.96	9825.42
2	0	17658.8	0	14245.4	13243.9	0	13956.9	13349.9	0	14208.8	13044.4	0	8514.34	8585.6	7920.08	9017.85	9571.03	8946.82	9012.64	8006.78	9183.01	9527.65	8794.78	8580.42	7827.96	8826.24	9152.48
3	0	19893.2	0	16632.2	15347.4	0	16737.4	15435.7	0	16697.7	15127.5	0	10308.8	10165.1	9717.67	11007.7	11557.1	11138	10668.2	10360.1	11151.4	11753.1	10588	10212	9522.2	10364.5	10733.9
4	0	22796.5	0	19299.5	17914.3	0	19027	17772.9	0	19208.3	17390.3	0	11967	12033.8	11599.2	12658.2	13753.6	13208.1	12584.6	11739	12806.3	13516.9	12442.8	12120.9	11384.9	12124.2	12804.2
5	0	25409.4	0	21213.6	20183.2	0	21523	20366.4	0	21191.3	19548.8	0	13622.3	13965.4	13181.2	14491.3	15701.9	15087.3	14419.8	13253.3	14341	15168.7	14283.5	13697.8	12666.4	13891.4	14635.2
6	0	27217.6	0	23444.4	21728.2	0	23464.5	21950	0	23645.5	21510.7	0	15533.4	15570.2	14641.9	16003.3	17072.3	16642.5	16136.4	14935.1	15708.7	16640.9	16033.2	15094.8	14155.7	15176.3	15896
7	0	29213.3	0	24763.7	23569.5	0	25425.1	23766.5	0	24914.6	23039	0	16725.8	16802.2	16212	17375.4	18807.4	17962.9	17397.4	16215.5	17201.1	18141.2	17349.4	16688.4	15477.4	16512.9	17203.6
8	0	29629.9	0	25973.9	24813.6	0	26080.2	24960.4	0	25959.6	24015.8	0	17480.5	18080.5	18047.1	18790.8	19666.5	18991.8	18485.2	17815.3	18501.6	19432.8	18381.2	17678.4	17157.1	17661.7	17922.7
9	0	29717.5	0	26320.4	25308.1	0	26714.1	25996.9	0	26518.5	25151	0	18792	19355.2	20054.9	19771.6	20664.4	20548.9	20054.4	20317.4	19601	20970.6	20234.4	18674.1	19489.8	18977.1	19348.7
10	0	29841.4	0	26896.2	25669.3	0	27328.4	26808.7	0	27384	26289.5	0	20846.1	21290.9	22885.2	22039.5	23280.1	22788.5	22390.8	21363.9	22976.2	22262.2	21063.7	22111.6	20818.5	21665.7	22157.7
11	0	30144.9	0	27204.4	26008.7	0	27323.7	26835.8	0	27384	26289.5	0	23196.2	23797.2	24908.3	24635.4	25048.1	24842.1	24308.2	25775.6	24629.6	24921.8	23377.8	24605	23309.3	23752	24312.6
12	0	29364.1	0	26732.8	26333.9	0	26990.5	26267.8	0	26445.3	25486.4	0	24313.4	24505.9	26235.8	25547	25984.8	25642.1	25900.3	26435.1	24934.6	23373.3	24425.5	24442.8	25021.9	24604.5	24832.6
13	0	28549.3	0	26394.7	25221	0	26971.9	25563.2	0	26958	24704.4	0	24213.4	24054	23857.1	24986.2	26287.5	25440	24998.1	25703.4	24647.3	25199.5	24654.3	24091.3	23505.2	23785.4	25189.7
14	0	27664.3	0	25533.9	23868.8	0	25278.8	23860.9	0	25308.6	23591.7	0	23437.1	23976.3	24070	23337.6	25210	24160.7	23093	24197.5	23226	24871.2	23438	22306.4	23586.9	22167.2	24710.7
15	0	25604	0	23580.1	21582.6	0	23144.7	21941.1	0	23308.6	21351.7	0	21719.2	20542.2	22113.9	21263	23520.4	22003.5	20660.1	22206.6	20964.9	23169.9	21505.4	19989.7	21463.8	20695.2	23035.5
16	0	31996.1	0	28834.1	25977.8	0	27691.4	26210.2	0	27976.9	25689.5	0	25955	24546	25715.3	25178.1	27510.2	25600.4	24246	25507.8	25258.4	26622.5	24885.5	23551	24710.9	24793.1	27194.9
1	2918.22	14658.1	2923.94	12034.4	11780.3	2915.6	12036.6	11401.5	2916.21	12113.8	11242	2920.35	7786.1	8497.62	7470.11	8813.31	8680.38	8654.33	8313.4	7906.38	8469.83	8887.32	8479.54	8018.34	7587.54	8235.51	8780.77
2	2533.29	13468.5	2534.59	11374.7	10710.1	2538.48	11246.4	10637	2541.61	11388.5	10431.4	2536.23	7944.21	8015.76	7292.91	8315.33	8270.17	8274.71	7774.08	7694.91	8241.96	8565.7	8141.6	7781.17	7488.15	7933.89	8235.98
3	2920.49	14384.5	2925.42	12350.8	11854.7	2924.53	12302.6	11731.4	2919.02	12511	11465	2924.47	8654.54	9125.3	8308.22	9323.58	9414.47	9406.37	8921.99	8766.67	9048.58	9542.17	9092.74	8673.81	8435.5	8778.22	9063.37
4	2913.81	15948.6	2912.45	13773.9	13386.6	2915.5	13983.5	13319.4	2918.26	14087.4	12881.8	2921.4	9717.66	10215.8	9334.69	10485.3	10752.2	10685.1	10097.4	9803.35	10041.4	10900.1	10335.7	9795.76	9543.14	9831.78	10375.7
5	2904.75	17355.3	2907.57	15169.8	14777.6	2937.17	15337.9	14458.2	2936.85	15323.3	14149.4	2932.47	10588.6	11192.2	10356	11422.5	11868.7	11772.2	11067	10817.9	10971	11926.7	11226.5	10673.7	10426.6	10677.5	11337.6
6	2922.63	18778.8	2922.78	16344.8	15982	2917.21	16553.5	15628.9	2943.27	16615.7	15174.4	2924.16	11395	12142.9	11727.8	12325.1	12791.7	12699.4	11943	11796.8	11931.3	12877.4	12140.3	11594.8	11303.7	11591.3	12206.2
7	2914.67	19732.5	2921.66	17277.7	16835.6	2930.43	17575.1	16849.2	2951.1	17579.8	16348	2918.69	12241	12919.3	12129.3	13153.4	13558.5	13581.1	12786.8	12612.3	12744.3	13627.8	12905.3	12489.6	12167.6	13461.8	13905.4
8	2918.81	19883.5	2920.73	17934.4	17641.4	2901.35	18107	17633.6	2908.58	18132.3	17022.6	2920.61	13028.1	13724.7	13138.6	13894.5	14198.2	14270.2	13506.7	13304.7	13444.4	14311.4	13584.5	13333.4	13110.7	13159.5	13678.5
9	2870.25	19975.9	2919.16	18267.4	18133.9	2918.48	18401.9	18154.9	2909.59	18428.1	17440.8	2920.53	13914	14379.7	14347.1	14706.6	15084.6	15166.9	14558.7	14716.5	14201.3	15186.1	14483	14185.4	14182.7	13933.4	14279.5
10	2918.45	20258.1	2906.07	18373.8	18333.7	2906.6	18640	18661	2934.89	18776.2	17810.7	2922.54	15102.7	15822.6	16001.5	15865.2	16156.1	16258.2	15651.5	15464.6	16254.1	15597.9	15384.1	15645	15292.5	15452.4	16000.3
11	2919.13	20204.4	2919.22	18563.5	18468.6	2920.12	18644	18796.4	2934.55	18826.9	17974.8	2921.7	16233.2	17141.3	17500.3	17990.4	17440.3	17507	17161.8	18076.7	18338.9	17325.6	16644.3	16770.8	17054.6	16600.3	16612.5
12	2735.46	20138.7	2839.78	18688.2	18408.5	2805.35	18478.8	18526.9	2939.65	18701.1	17737.9	2923.51	16925.6	17841.6	17931.5	17906.4	17906.4	18838.2	17253.3	17533.3	15464.6	16254.1	15597.9	15384.1	15645	15292.5	15452.4
13	2780.98	19878.2	2716.81	18457.7	17855.2	2718.66	18264.5	17887	2836.37	18486	17291.4	2915.01	16990.4	17990.5	17433.8	17958.9	17978.5	17865	17191	18173.6	1704	17478.7	17137.7	17037.6	17213.2	17110	17704.2
14	2918.2	18596.7	2915.67	17491.7	16975.2	2926.89	17963.7	16721.9	2922.8	17784.3	16413.1	2922.8	16699.1	16184.6	16752.7	17275.9	17282.2	16236	17341.1	16033.4	17126.6	16498.3	15929	16290.2	16111.5	17078.9	17078.9
15	2539.91	18254	2539.82	16790	15475	2541.42	16504.9	15719.4	2555.07	16888.4	15126.1	2526.41	15287.4	15025.2	15353.3	15566.4	16322.5	16041	14795.9	16056.4	15064.7	16481.4	15575.2	14771.3	15470	15132	16588.8
16	2918.86	21357	3523.86	19622.6	18371	2923.13	18959.6	18132.8	2922.36	19126.1	17765.1	2920.53	14679.4	17591.1	17755.6	17753	18830.5	18148.9	16972.1	17925	17032.2	18603.3	17588.7	16569.3	17505.7	17330.8	18632.1

Figure 4-3: The MCNP output and SVR prediction for the fission power distribution for a shim bank at the 19cm position.

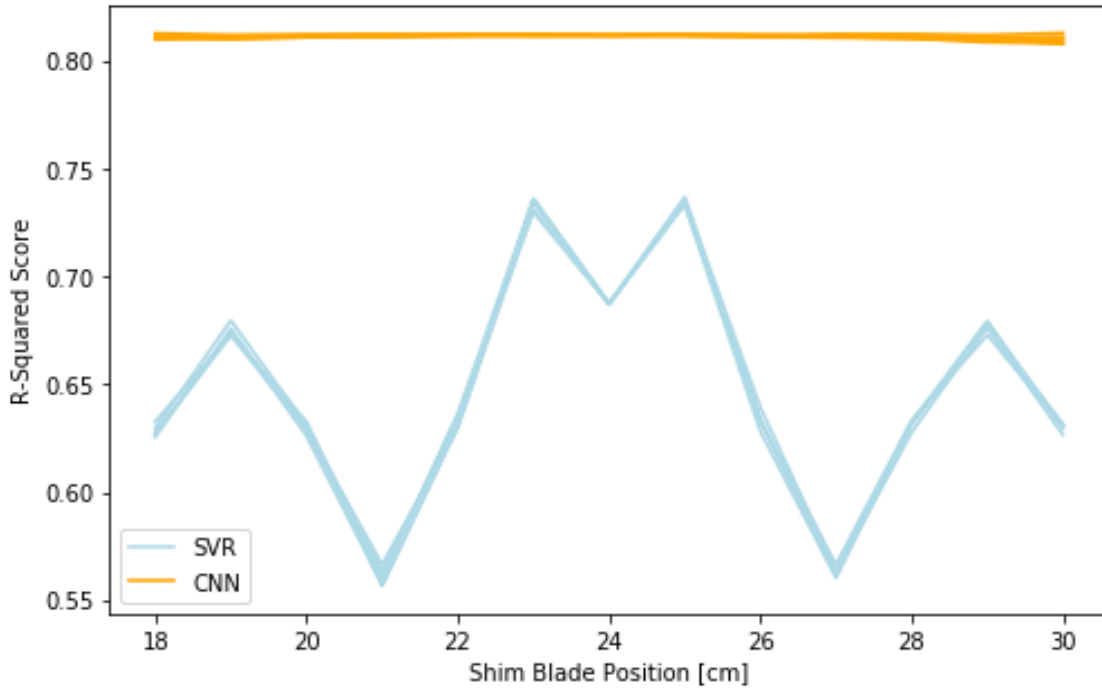


Figure 4-4: The R-Squared score of an SVR and CNN prediction against the MCNP run which modelled the corresponding shim blade height, for every 1cm step of each blade while the remaining blades are at 24cm.

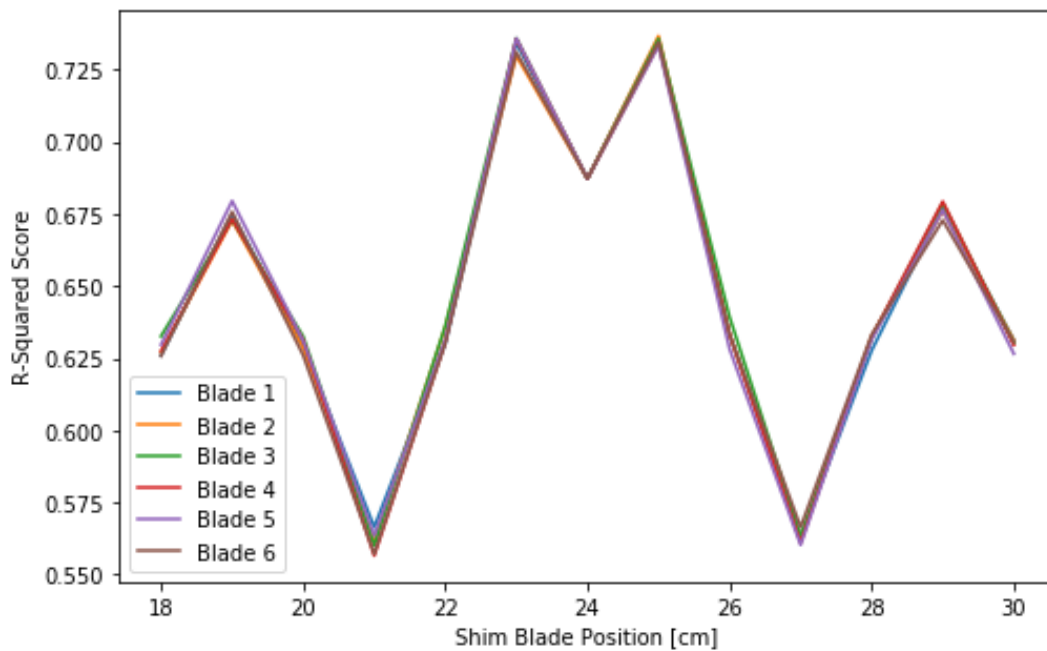


Figure 4-5: The R-Squared score of an SVR prediction against the MCNP run which modelled the corresponding shim blade height, for every 1cm step of each blade while the remaining blades are at 24cm.

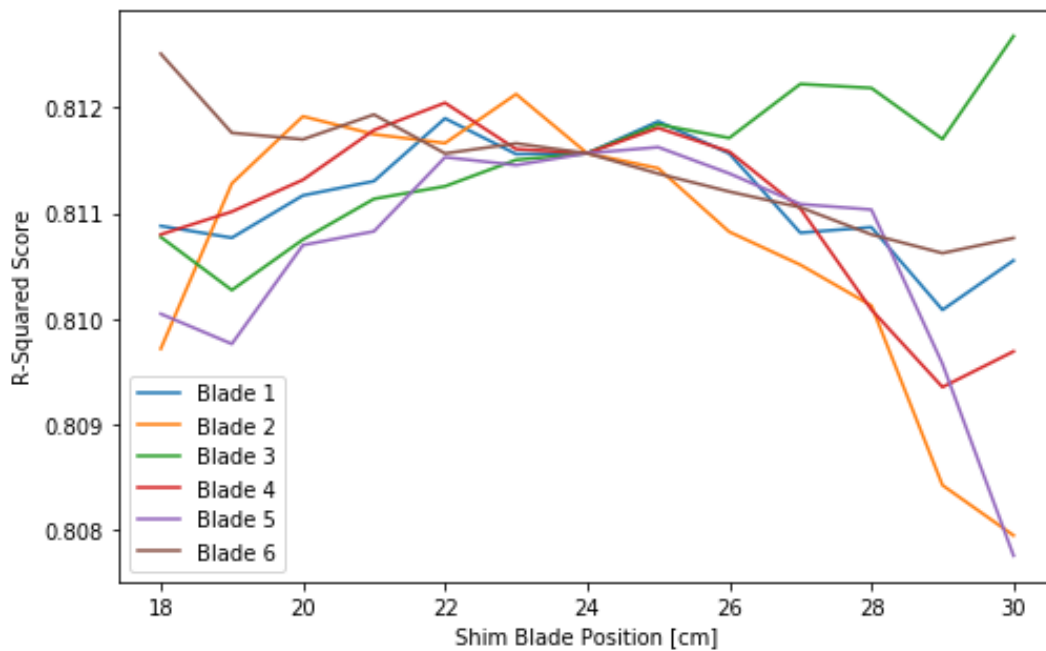


Figure 4-6: The R-Squared score of a CNN prediction against the MCNP run which modelled the corresponding shim blade height, for every 1cm step of each blade while the remaining blades are at 24cm.

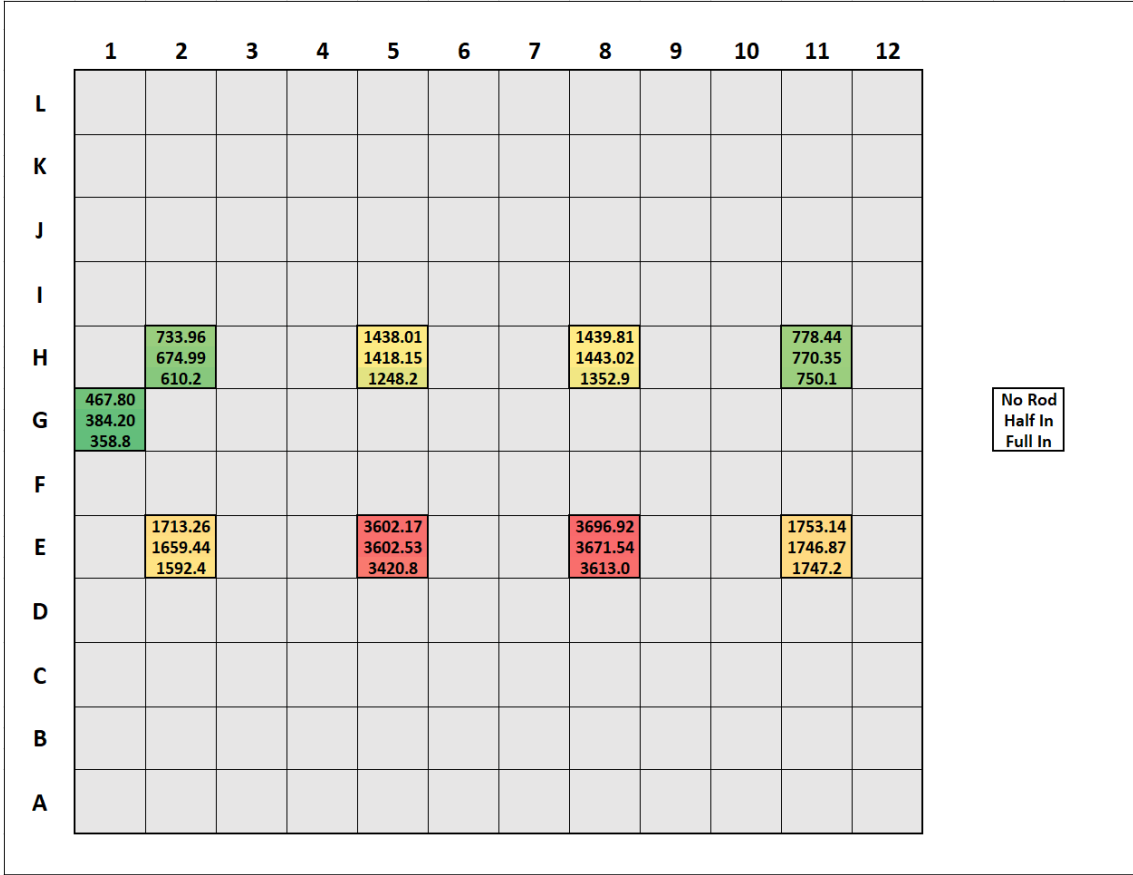


Figure 4-8: The measured neutron count rates for every detector position, as measured in counts/s. A count rate was measured for every detector position with the control rod fully out (0"), halfway inserted (23"), and fully inserted (46") into the pile. Here every square represents one of the fuel channels on the pile's North face. Measurements were only taken for nine of these positions, as denoted by the ones with data. The color scale denotes the magnitude of the count rate, with red being higher, green being lower.




Figure 4-9: Count rate as a percentage for every detector position. For every position, the count rate measured with the rod fully out (0") is used as the reference, and as such is denoted as 100%. The count rates for the halfway inserted rod (23") and fully inserted rod (46") are given as a percentage of this reference, for every detector position. Here every square represents one of the fuel channels on the pile's North face. Measurements were only taken for nine of these positions, as denoted by the ones with data. The color scale denotes the percentage change from the reference, with red being higher, green being lower

Chapter 5

Conclusion

This project sought to investigate the feasibility of using machine learning frameworks to develop an autonomous control system for the MGEP. Overall the findings are favorable, and suggest that this task is possible. As such, this project will lay the groundwork for the future development of any such system

For the investigation of ML methods, two key findings were achieved. First, both CNN and SVR algorithms seem capable of predicting the response of a fission system to control mechanism insertion fairly well, however the CNN framework outperforms the SVR framework both in terms of speed and accuracy. While there were variations in the SVR accuracy which must be investigated further, the CNN models remained highly consistent with R-Squared values of 0.8 or higher. Additionally, because the SVR models took hours to train as compared to the CNN models which took minutes, the latter approach seems favorable for further development and experimentation. Second, when applying a CNN algorithm to predict the flux distribution of the MGEP based on its source location, high accuracy can be attained in the model's predictive ability, especially in the more pertinent central source locations of the pile pedestal. The average R-Squared score over the full set of source locations was 0.78, however for much of the central locations it was 0.95 or higher.

For the experimental aspect of control rod design and fabrication, the most important tasks have been accomplished. Primarily, a control rod assembly consisting of a boron-containing control rod and barrier tube has been fully fabricated and as-

sembled. It has been experimentally verified that the insertion of this control rod does have a measurable and predictable effect on detector response within the pile, as reductions in count rate as large as 23% from the rod removed to fully inserted were measured. Measured reductions in count rate varied by detector location, however global reductions on the scale of 3-6% were seen.

Future work for ML approaches can seek to improve the models through testing to find optimal training parameters, to achieve better speed and accuracy. Specifically, further work is necessary to better assess the performance of both methods, especially SVR, to optimize them and increase accuracy. It is likely that in the case of the MITR cases, removing empty fuel elements from the ML training set would greatly increase model performance. Additionally, because a control rod now exists and has been verified to work with the current pile setup, both Monte Carlo and ML modelling can be done to test the accuracy of each at predicting the behavior of the pile with control rod insertion.

Future work with the control rod apparatus could seek to perform more advanced experiments using the rod to better understand its effect on the flux and power profile of the pile. Additionally, only static measurements of the neutron count rate were taken for the experiment conducted in this study. Future work could develop a script which measures and records a time variant effect with rod insertion. Finally, a simple but effective tool now exists to alter pile behavior. If an autonomous control system is to be designed for the pile, the control rod will serve as an important tool to perturb the pile and react to feedback.

Appendix A

Training Data Examples

pos	4.116161	0	-41.275	0.387368	0.395394	0.348144	0.2683	0.19133	0.116556	0.052929
0.043571	0.096865	0.160107	0.238657	0.323584	0.387368	0.395394	0.2683	0.19133	0.116556	0.052929
0.040676	0.088127	0.138435	0.194597	0.247005	0.283436	0.285312	0.21467	0.159194	0.100712	0.04551
0.03536	0.078259	0.118112	0.152997	0.185153	0.209054	0.217	0.166792	0.128705	0.08433	0.038301
0.02834	0.065027	0.096263	0.123847	0.147006	0.158219	0.165623	0.128946	0.102063	0.067327	0.032359
0.02202	0.048974	0.073844	0.092226	0.108039	0.121978	0.124713	0.099824	0.082319	0.055132	0.024864
0.01764	0.037866	0.056034	0.072302	0.085793	0.091494	0.089324	0.075918	0.060002	0.039755	0.020202
0.013969	0.030645	0.043027	0.051419	0.061783	0.068486	0.070545	0.057255	0.046477	0.030708	0.01397
0.012558	0.02221	0.034386	0.042646	0.050068	0.054061	0.048569	0.042103	0.033501	0.024044	0.012215
0.007555	0.014638	0.023195	0.030085	0.033639	0.037776	0.040909	0.03537	0.023775	0.017633	0.008312
0.0056	0.011885	0.014582	0.021057	0.024359	0.024473	0.026727	0.020765	0.017132	0.011159	0.005196
0.00322	0.005599	0.009924	0.01205	0.013846	0.015981	0.015237	0.013774	0.010439	0.006908	0.002946
0.001327	0.002634	0.004104	0.005282	0.007524	0.007705	0.007372	0.006535	0.005187	0.002832	0.0013

Figure A-2: An example of a fully-formatted training data input for modelling and prediction of the MGEP. The 3-dimensional power distribution was compressed to 2 dimensions; each cell represents a fuel rod, with rows corresponding to the rows of fuel channels in the pile. This example is for the OpenMC data generated for a source location (4.11, 0, -41.275). 650 unique source positions were modelled to generate the training data for the MGEP case.

Appendix B

SWX-237Z30 Full Material Properties



SWX-237Z30 High Temperature Boron-Silicone Technical Information

SWX-237Z30 Boron-Silicone is a fire- and heat-resistant castable neutron shielding material. It is a modified form of our standard SWX-237 where the boron content has been increased to nominally 30% by weight.

Applications: Neutron Shielding

Typical Uses: Fire Safety, Criticality, High Temperatures, Source Storage

Shielding Effectiveness For Various Types of Radiation:

Thermal Neutrons -	excellent
Fast Neutrons -	good
Gammas -	fair
Capture Gammas -	good

Temperature Limit: 400°F (205°C)

Machinability: Fair

Forms and Sizes of Shielding: Pre-Cast to a wide range of geometries

Radiation Resistance (structural and shielding integrity):

Accumulated Gamma Radiation Exposure Limit: 1.0×10^{10} Rad
Accumulated Neutron Radiation Exposure Limit: 5.0×10^{18} n / cm²

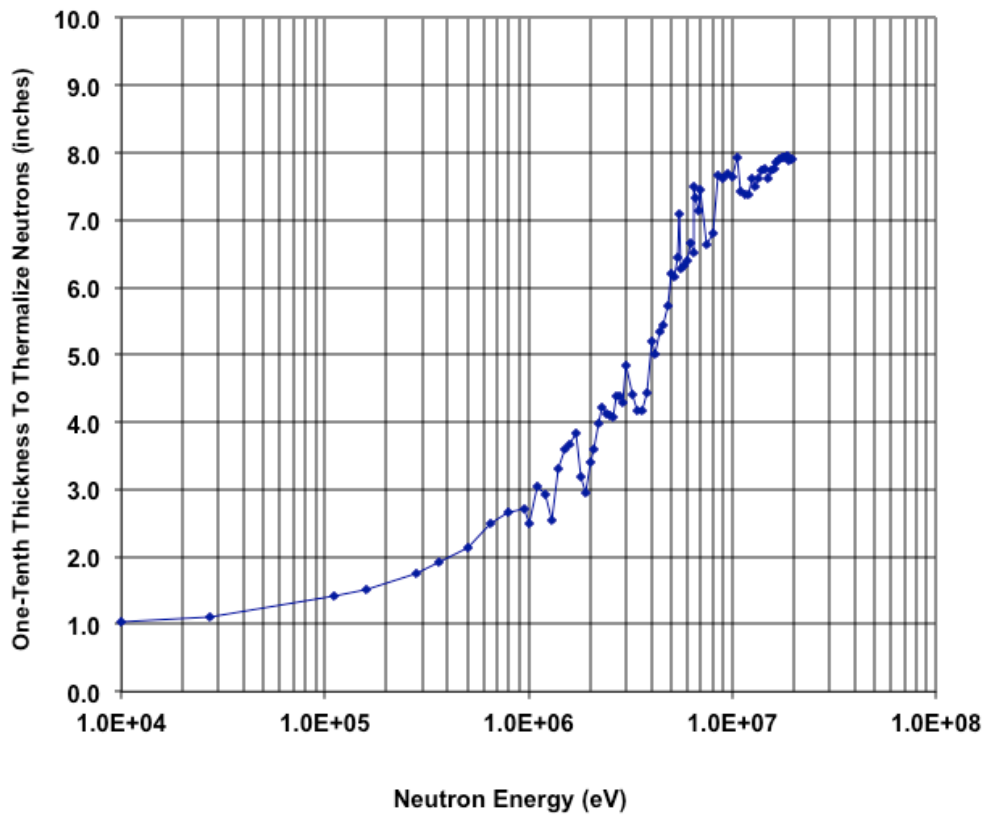
Density: 1.59 gram / cm³ (99 pounds / cubic foot)

Thermal Neutron One-Tenth Thickness: 0.045 inches

Thickness of SWX-237Z30 required to reduce incident flux of thermal neutrons by a factor of 10 (exit thermal flux is 10% of incident thermal flux)


Neutron Moderation Characteristics

Thickness of SWX-237Z30 required to reduce to thermal 90% of an incident neutron flux, as a function of initial neutron energy (exit epithermal flux is 10% of incident neutron flux)



Gamma Attenuation Characteristics

Thickness of SWX-237Z30 required to reduce an incident gamma flux by a factor of 10, as a function of incident gamma energy (exit gamma flux is 10% of incident gamma flux)





SWX-237Z30 High Temperature Boron-Silicone

Nominal Elemental Analysis

<u>Element</u>	<u>Percent by Weight</u>	<u>Number of atoms/cc</u>
Oxygen	23.39	1.40×10^{22}
Aluminum	5.67	2.01×10^{21}
Silicon	17.57	5.97×10^{21}
Carbon	19.08	1.52×10^{22}
Hydrogen	3.27	3.10×10^{22}
Boron	30.51	2.70×10^{22}
Sodium	0.04	1.67×10^{19}
Iron	0.18	3.10×10^{19}
Zinc	0.10	1.43×10^{19}
Nitrogen	0.17	1.19×10^{20}
Density	1.59 gram/cm ³	

Appendix C

Neutron Count Rate Data

Table C.1: Detailed data for count rate measurements for various detector positions, for the control rod removed, inserted half-length, and fully-inserted.

Fuel Channel	Control Rod Insertion [in]	Count Rate [counts/s]	1σ [counts/s]
E2	0	1713.255	3.118275
E2	23	1659.437	4.508974
E2	46	1592.394	4.040749
E5	0	3602.175	6.008239
E5	23	3602.526	7.822378
E5	46	3420.756	5.077239
E8	0	3696.92	5.997283
E8	23	3671.54	5.947902
E8	46	3612.987	6.141305
E11	0	1753.139	4.578462
E11	23	1746.872	4.543345
E11	46	1747.209	4.869777
H2	0	733.9644	1.916356
H2	23	674.99	2.059361
H2	46	610.2467	1.798917
H5	0	1438.007	3.985844
H5	23	1418.154	4.095695
H5	46	1248.189	3.19545
H8	0	1439.805	3.966589
H8	23	1443.025	3.907801
H8	46	1352.908	3.909957
H11	0	778.4376	3.174669
H11	23	770.3486	2.722174
H11	46	750.1441	2.57839
G1	0	467.8042	1.45728
G1	23	384.199	1.257445
G1	46	358.8123	1.32457

Bibliography

- [1] Backgrounder on the Three Mile Island Accident. Backgrounder, United States Nuclear Regulatory Commission, June 2018. Available at: <https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/3mile-isle.html>.
- [2] Backgrounder on Chernobyl Nuclear Power Plant Accident. Backgrounder, United States Nuclear Regulatory Commission, August 2018. Available at: <https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/chernobyl-bg.html>.
- [3] Backgrounder on Improvements Resulting From Davis-Besse Incident. Backgrounder, United States Nuclear Regulatory Commission, June 2018. Available at: <https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/davis-besse-improv.html>.
- [4] The Autonomous Car Industry Emerges. *Trends Magazine*, pages 29–36, 2017.
- [5] Nuclear Reactor Laboratory, Report to the President. Technical report, Massachusetts Institute of Technology Nuclear Reactor Laboratory, 2016. Available at: <http://web.mit.edu/annualreports/pres16/2016.18.12.pdf>.
- [6] Micah Gale. Developing Modern Graphite Exponential Pile Experiments to Augment Reactor Physics Education. Bachelor’s Thesis, Massachusetts Institute of Technology, Department of Nuclear Science and Engineering, June 2018.
- [7] Samik Raychaudhuri. Introduction to Monte Carlo Simulation. Miami, FL, December 2008. IEEE.
- [8] Geant4. Available at: <https://geant4.web.cern.ch/>.
- [9] OpenMC. Available at: <http://web.mit.edu/smharper/www/index.html>.
- [10] A General Monte Carlo N-Particle (MCNP) Transport Code. Available at: <https://mcnp.lanl.gov/>.
- [11] Osvaldo Simeone. A Very Brief Introduction to Machine Learning With Applications to Communication Systems. Technical report, IEEE, August 2018. Available at: <https://arxiv.org/abs/1808.02342>.

- [12] Mario Gomez Fernandez, Akira Tokuhiko, Kent Welter, and Wu Qiao. Nuclear Energy System’s Behavior and Decision-Making Using Machine Learning. *Nuclear Engineering and Design*, 324, 2017.
- [13] Kresimir Trontl, Dubravko Pevec, and Tomislav Smuc. Machine Learning of the Reactor Core Loading Pattern Critical Parameters. *Science & Technology of Nuclear Installations*, 2008.
- [14] Kaichao Sun, Yunyun Zeng, Liu Jingquan, and Lin-Wen Hu. Machine Learning Based System Performance Prediction Model for Reactor Control. *Annals of Nuclear Energy*, 3113:270–278, November 2017.
- [15] Alex J. Smola and Bernhard Scholkopf. A Tutorial on Support Vector Regression. *Statistics and Computing*, 14:199–222, 2004.
- [16] Chih-Chung Chang and Chih-Jen Lin. LIBSVM – A Library for Support Vector Machines, July 2018. Available at: <https://www.csie.ntu.edu.tw/~cjlin/libsvm/>.
- [17] Keiron O’Shea and Ryan Nash. An Introduction to Convolutional Neural Networks, November 2015. Available at: <https://arxiv.org/abs/1511.08458>.
- [18] Tensorflow. Available at: <https://www.tensorflow.org/>.
- [19] Keras: The Python Deep Learning Library. Available at: <https://keras.io/>.